

©Copyright 2025

Kartik Krishna

Coherent Structures and Optimal Control Theory

Kartik Krishna

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

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Steven L. Brunton, Chair

Zhuoyuan Song, Co-Chair

Xu Chen

Program Authorized to Offer Degree:
Dept. of Mechanical Engineering

University of Washington

Abstract

Coherent Structures and Optimal Control Theory

Kartik Krishna

Chair of the Supervisory Committee:

James B. Morrison Endowed Career Development Professor Steven L. Brunton
Department of Mechanical Engineering

Coherent structures are persistent, large-scale spatiotemporal features in fluid flow fields. Optimal control theory is a branch of mathematics highly relevant to engineering for manipulating the behavior of dynamical systems. In this work, we explore the interplay and connections that arise between these two concepts in two different contexts. In the first context, we are motivated by the problem of path planning of ocean drifters moving within flow fields. Here, we are interested in how the background coherent structures characterized by the Lyapunov exponents of the flow impact and shape energy efficient trajectories of the robot moving within. In the second context, we investigate a strategy for controlling the dynamics of swirling vortex structures in flow fields. In this case, we consider “coherent structures” of the flow field characterized by Koopman eigenfunctions. In particular, we investigate how formulating control in terms of these eigenfunctions changes the vortex dynamics. Given that there has been much recent work in data-driven computation of these coherent structures, we anticipate this work to be of considerable interest to scientists and engineers.

TABLE OF CONTENTS

	Page
List of Figures	iii
Chapter 1: Introduction	1
1.1 Motivating Problem	1
1.2 Overview and Literature Review	2
1.3 Background	4
1.4 Model Problem	10
Chapter 2: Interplay of sensor dynamics with Lagrangian coherent structures .	14
2.1 Introduction	14
2.2 Double Gyre Results	16
2.3 Advanced Test Cases	25
2.4 Discussion	28
Chapter 3: Finite-time Lyapunov exponent analysis of controlled agents in flow fields	34
3.1 Introduction	34
3.2 Methodology for Control FTLE (cFTLE)	36
3.3 Results	38
3.4 Example from Reinforcement Learning	46
3.5 Discussion	49
Chapter 4: Control of point vortex dynamics using Koopman eigenfunctions . .	51
4.1 Introduction	51
4.2 Methodology	54
4.3 Single invariant control results	59
4.4 Multi-invariant control results	64
4.5 Mixing	68
Chapter 5: Conclusion	71

5.1	Control within a flow field	71
5.2	Control of a flow field	73
Appendix A:	Supplementary data - Chapter 2	97
Appendix B:	Supplementary data - Chapter 3	109

LIST OF FIGURES

Figure Number		Page
2.1	<p>Overview of the proposed methodology for analyzing the connections between finite-horizon energy optimal trajectories and the FTLE field. A self-propelling agent is controlled to transit from a starting location to a goal location through a finite-time horizon energy-optimal trajectory in a time-varying double gyre flow field. The resulting agent trajectory, along with the finite-horizon predicted trajectories at each time step, are shown and color-coded based on instantaneous energy expenditure (top left). The trajectory history (solid) and the future forecast trajectory bundle (dashed) at an example time instant are shown (top right); the instantaneous FTLE ridges are also shown below these with blue indicating the repelling LCS and red indicating the attracting LCS. As can be observed from the snapshots taken at four particular times (bottom), the energy expenditure along the planned trajectory, given by the color of the dashed line), and the shape of the finite-horizon trajectory depend on the evolution of the local FTLE ridges. . . .</p>	15
2.2	<p>Dependency of the resulting sensor trajectories on the ratio between the penalties on energy expenditure (R) and state error (Q), for a fixed time horizon $T_H = 4$. The trajectories are color-coded by instantaneous energy expenditure. There are three costs shown in the top right figure: $J = J_u + J_e$, $J_e = \sum Q(x - x_g)^T(x - x_g)\Delta t$, and $J_u = \sum Ru^T u\Delta t$. Here, J generally increases with R/Q. The trajectories undergo several qualitative changes (bifurcations) as R/Q is increased, forming different types of periodic orbits shown on the bottom of the figure. An example of this is when R/Q is changed from 25 to 26, we observe a major change in the shape of the final orbit around the goal, as opposed to the minor change from $R/Q = 15$ to 25, where the final eye-shaped orbit only gradually increases in size. The formation of these orbits are dependent on the background flow FTLE, as can be seen from the inset of case $R/Q = 2$. We observe that as the LCS move, the stable and unstable LCS intersect at a point, whose location changes every instant, and the sensor moves in a manner in which it is right on top of this intersection point for most of the time when $R/Q = 2$ (top left box). . .</p>	17

2.3	Different instances of a trajectory for different time horizons while keeping $R/Q = 2$ fixed. By varying the time horizon, T_H , we see that the extra loop in Figure 2.2 is due to a sensitivity of the planned path with respect to T_H , where the agent becomes stuck in the right gyre for lower T_H . We can see here with reference to Figure A.1 how as the time horizon increases, the agent moves less aggressively and improves its timing with the gyre oscillations to reach the goal sooner without taking an extra loop. Another interesting aspect of this plot can be seen at $t_3 = 4$ and $t_4 = 6$ where the multiple agents line up to mirror the movement of the red LCS, further confirming the strong correlation finite-horizon optimal trajectories and the FTLE ridges.	18
2.4	Influence of the FTLE ridges on the energy expenditure both along prediction horizon (right top) and instantaneously (right bottom) under parameters $T_H = 4.0$, $R/Q = 100$, and $\Delta t = 0.1$. An agent's motion along with the predicted forecast trajectory (dashed line) are shown on the left, together with the repelling (blue) LCS and the attracting (red) LCS. There is a correlation between the spike in both the instantaneous energy spent (right top) and the cost along the forecast trajectory (right bottom) with movement across an FTLE ridge. Here, unlike in Figure 2.2, the summations $J_e = \sum Q(x - x_g)^T(x - x_g)\Delta t$ and $J_u = \sum Ru^T u\Delta t$ are only along the forward dashed line in the plots under 'snapshots' and not along the entire trajectory as in Figure 2.2	20
2.5	The mobile sensor settle on periodic orbits around the goal state (left) and the magnitude of the Fourier transform of the instantaneous energy spent by the mobile sensor (right). We observe that the time series of the energy spent is periodic with frequencies at integer multiples of the double gyre oscillation frequency, which correspond to the peaks in the right plot.	21
2.6	Variations of agent trajectory under different double gyre oscillation frequencies. The frequency of the periodic orbits depends on the double gyre oscillation frequency.	21
2.7	Multiple simulations were carried out at each R/Q ratio spaced logarithmically, from 0 to 100, time horizon ranging from 1 to 10, and gyre frequency ranging from $2\pi/4$ to $2\pi/14$. The data presented here is in the form of scatter plots for each gyre frequency with each color representing the Pareto optimal tradeoff curve between the total energy spent along each trajectory and the sum of deviations from target along the trajectory. The trajectories shown in the bottom row correspond to the highlighted purple circles (1,2,3) in the Pareto optimal corresponding to $\omega_4 = 2\pi/10$	22

2.8	As the sensor moves in the double gyre flow field, it is constantly taking control actions $u = [u_x, u_y]$, where, u_x, u_y are the x and y-components of its actuation respectively. At each instant, the sensor is also moving over a background double gyre flow velocity vector whose components, v_x and v_y , are given by (1.15). The top row of histograms are of the magnitude of control actions $\ u\ $ taken (in red), against the magnitude of the background current velocity $\ v(x_s, y_s, t)\ $ (in grey), where, x_s, y_s are the sensor coordinates at time t . The second row shows the heading angle of the sensor (the orientation of dx/dt), plotted in red, against the orientation of the background flow field velocity vector, plotted in grey. The corresponding trajectories are shown on the bottom.	23
2.9	MPC trajectories with varying R/Q ratio, starting from different initial conditions with the same goal location. We find that eventually, the trajectories converge to similar periodic orbits thereby using the LCS in similar ways to orbit around the goal location despite having different transients.	24
2.10	MPC trajectories to different goal locations with varying R/Q ratio for the same initial condition. We find that in the case of placing the goal near the middle to bottom region (top left and top middle plots) of double gyre, the sensor is able to form stable orbits near the attracting LCS, where the sensor moves with the base of the attracting LCS. However, placing the goal near the repelling LCS causes more difficulty for MPC in forming small stable orbits (top right). The bottom plots show that it is possible to form small periodic orbits in the corners of the double gyre flow field when the goal is placed close to them.	25
2.11	MPC trajectories formed by the sensor in an ABC flow field as a function of the R/Q ratio. We observe that for a time horizon of 5, we are able to find several cases of periodic orbits loitering close to the goal location. In these cases, the initial position is $x_{\text{start}} = [\pi/2, 1, 6]$ and the goal is $x_{\text{goal}} = [5, 2, 1]$.	27
2.12	This scatter plot shows the performance of MPC in the ABC flow field similar to Figure 2.7. We observe that MPC trajectories with short time horizons are able to reach the goal state and increasing the time horizon has a benefit of decreasing the actuation energy usage. We also observe that the inflection points and breaking off points correspond to bifurcations in the trajectories.	27

2.13	This figure shows the use of MPC to plan trajectories in the Gulf of Mexico dataset. We have chosen the region highlighted in dashed dotted lines in the top left plot. The bottom six plots show the trajectory generated for $R/Q = 1$ (full trajectory can be seen in the rightmost plot in Figure 2.14) and $u_x, u_y \leq 2$ km/hr in color shading from yellow to purple to highlight how energy is spent along the path. The units on the x and y axis are longitude and latitude respectively. The red cross in the figures show start location, $\mathbf{x}_{\text{start}} = [-85.5, 19.8]$, the green cross shows the goal location, $\mathbf{x}_{\text{goal}} = [-83.7, 18.9]$. The black dot shows the instantaneous sensor location. Viewing the six figures in sequence we observe that the sensor moves across the blue repelling LCS from $t = 6$ to $t = 8$, where t is the time in days. We observe a spike in the instantaneous energy spent and a slow drop as we move away from the repelling LCS (as seen in the top right plot of $u(t)$). We then observe that the sensor synchronizes with the attracting LCS to move towards the goal location.	29
2.14	This figure shows the change in MPC trajectories in the Gulf of Mexico as the R/Q ratio is varied. We see an aperiodic loitering state near the goal location. Given that the time horizon in this case $T_H = 9.6$ hours, which is relatively short compared to the total trajectory time of 100 days, we can see that even with short time horizons the MPC trajectories are capable of making it to the goal location	30
3.1	This figure outlines the methodology of cFTLE which takes Lagrangian controlled trajectories as input and outputs LCS curves which highlight interesting transport features of the control law which are generally not visualizable by simply plotting all the trajectories or by plotting the control vector field. The left most figure shows that in order to evaluate our control algorithm, we generate the simulation from a mesh grid of initial agent positions at $t = 0$. We end the simulations at $t = T_A$ (as shows in the middle figure). Finally we use the FTLE algorithm on this data to generate “cFTLE” ridges which highlight important features in the domain.	35
3.2	Control FTLE ridges separate dynamically different regions of agents moving in the flow field. We show two distinct patches of agents - a purple and a yellow patch. Both patches start initially on different sides of the cFTLE ridge. The time evolution shows that cFTLE ridges separate the patches that undergo different paths and accrue largely different short-term state error cost despite starting close to each other.	38

3.3	cFTLE boundaries enclose regions in the domain which accrue similar cost. Despite the fact that the purple and the yellow patch start far from each other, they meet in short time since they are in the same lobe. This can be viewed particularly at $t = 3.6$ on, where the yellow and purple patch are within the same bounded cFTLE region. These structures are not apparent when viewing the control law or the flow field alone. An application of visualizing such regions in the domain may aid in solving rendezvous problems in the ocean where two distinct groups of agents may need to meet in finite time	40
3.4	The plot (a) shows the control FTLE ridges for double gyre setup at an R/Q ratio of 80 on the left, with a time of advection $T_A = 15, T_H = 3$. The agents are tracking the goal set point at $\mathbf{x}_{goal} = [0.5, 0.5]$. On the right (b), we show the accumulated state error J_F for a meshgrid of agents in the domain for the same R/Q and T_H . We see that the cFTLE ridges are located where there is a drastic change in color on the right figure. This indicates that the cFTLE ridges at a particular T_A partition the domain by the cost spent (i.e., large change in J_F).	41
3.5	On the left, we plot the instantaneous energy spent by an agent starting from a meshgrid of initial conditions moving to a goal location on the left ($\mathbf{x}_{goal} = [0.5, 0.5]$) at $t = 0$. On the right, we plot the accumulated state error over the time horizon of MPC ($T_H = 4.5, R/Q = 15$). In both plots, the streams are the control law generated by all the MPC agents, and the cFTLE ridges are shown where the time of advection is equal to the time horizon of the MPC, $T_A = T_H$. We observe here that the cFTLE ridge overlaps with regions where actuation energy expenditure peaks locally. Also, similar to Figure 3.4, the cFTLE divides the domain where the cost gradient is the largest. 41	41
3.6	This figure shows the cFTLE ridges for four different R/Q ratios of 20, 40, 60, 80, and how the barriers shrink as the aggressiveness of the MPC increases. Here, the time of advection for cFTLE $T_A = 15$, and the time horizon of MPC is $T_H = 3$. We observe that as the aggressiveness of the control decreases, the cFTLE approaches the passive FTLE.	44
3.7	This figure highlights the change in cFTLE ridges as a function of the goal location. Here, $R/Q = 15, T_A = T_H = 4.5$. The streamplot on both plots show the control law. The color in the background shows the regions where most energy is spent (brighter yellow shows larger energy spent and darker shows less). On the left plot, the goal location is set to the right gyre, and on the right plot, the goal location is set to the left gyre. We observe how the fluxes generated by the control law moves the passive FTLE ridge. A left flux moves the cFTLE ridge to the right, and vice versa. This figure continues to highlight the connection between the terms in the cost function and cFTLE ridge as in Figure 3.4.	45

3.8	This figure shows the change in cFTLE field for time horizon values $T_H = \{2, 3, 4, 5\}$, using an advection time of $T_A = 15$ and $R/Q = 50$ at $t_0 = 0$. Unlike in Figure 3.6, the cFTLE ridge changes due to more intelligent use of control and not because of sheer use of greater effort. The cFTLE ridges exhibits more structures, and the curves have more branches due to several ridges coming close to each other and collapsing on each other.	46
3.9	This figure shows the range of policies that can be generated at time $t = 0$. The policies considered in this chapter vary with time t . The left and middle policies are generated using MPC with a time horizon of 1.0 and 3.0, respectively. The policy on the right is generated by using DDPG and can be considered the “infinite-horizon” case. All policies have an R/Q ratio of 70. From left to right, there is an increase in spatial complexity of the policy as it begins to use more information of the flow field. This spatial complexity in turn influences the cFTLE ridges computed, which can be seen in Figure 3.8 and Figure 3.10.	47
3.10	This example shows the repelling cFTLE field plotted for a policy generated using DDPG with $R/Q = 70$. Much like Figure 3.2, we use two patches to highlight transport mechanisms. We also use the same initial patch locations for comparison.	47
4.1	Uncontrolled vortex dynamics: The key dynamic states of a 4-vortex system are highlighted: periodic (or single-frequency periodic), quasi-periodic (or multi-frequency periodic), and chaotic (broadband frequency). These dynamic states differ by the angle the vortices initially make with the origin [18]. The control demonstrated in this work aims at switching between these different states. From left to right, we show the mixing of tracer particles (starting from the initial distribution to the left), vortex trajectories and the Fourier transform (FFT) applied to the x-component of a single vortex to characterize occurring frequencies. A broad spectrum implies chaos, whereas sharp peaks at discrete frequencies can imply quasiperiodicity or periodicity.	52
4.2	Impact of reference tracking of the four invariants of the vortex system on the vortices’ trajectories using an actuator in a fixed position. From top row to bottom: (1) Different dynamic regimes appear depending on a higher or lower reference value for the Hamiltonian. (2) Increasing angular impulse translates the vortices away from the origin. Negative A shows that vortices merge closer in a similar strategy to increasing the Hamiltonian. (3-4) Change in linear impulse manifests as translation of the vortex configuration.	60

4.3	MPC control signals (circulation) corresponding to the different cases of invariant control displayed in Fig.2. Red curves are associated with a higher reference value, while blue curves represent a lower reference value in the objective function. From left to right: (1) We can observe here that the red curve is oscillatory with increasing frequency in time while the blue curve decreases the frequency in time. (2) While an increase in angular momentum results in a step function, a decrease yields consistent periodic behavior similar to the Hamiltonian but without convergence. (3) In contrast, convergence can be quickly achieved using a step function for both increasing and decreasing linear impulse.	63
4.4	This figure summarizes the resulting dynamics from multi-invariant control. The first row shows the vortex trajectories with the reference change in angular impulse. The second row is for the change in Hamiltonian	65
4.5	Schematic diagram for transitions: On the left, we show how multi-invariant control can be used to move between periodic and quasi-periodic regimes as seen in Figure 4.1. Here, the average radius/variance of the vortex distribution and the center of vorticity or the mean can both be controlled. When the Hamiltonian is increased under these constraints, the vortex merging is constrained to states along the dotted circle, which exhibit quasi-periodic behavior. Transitioning between different dynamic regimes is very limited using single-invariant control, whereas multi-invariant control enables one to reach a more diverse set of states. On the right, we show that a similar transition is possible from periodic to chaotic behavior by shifting the centre of vorticity to the right, thereby causing an asymmetry in the vortex configuration. These particular configurations have been exhaustively studied independently in [18]. The key difference between these configurations is the angles the vortices make with the origin.	66
4.6	This plot shows the power spectral density of the x-component of a single vortex for both single- and multi-invariant Hamiltonian control. This information can be used to classify different dynamic regimes as shown in Figure 1. We see that in single-invariant control, increasing the Hamiltonian yields periodic behavior with a single frequency, whereas decreasing the Hamiltonian modifies the dynamics towards chaotic behavior with broadband frequency characteristics. However, for multi-invariant control, this chaotic transition does not occur. Instead, when increasing the Hamiltonian quasi-periodicity with multiple frequencies can be observed. This is due to enforcing symmetry in the case of multi-invariant control.	67

4.7	This plot shows the comparison between single invariant and multi invariant Hamiltonian control in terms of mixing. The plots are scatter plots of actuator effort $\int u(t)dt$ over time on the x -axis and the KL-Divergence on the y -axis. We observe that enforcing symmetry decreases the KL-divergence in the case of decreasing Hamiltonian. This is due to vortex dynamics being restricted to periodic behavior instead of chaotic. In the case of increasing Hamiltonian, the mixing improves due to a transition to quasiperiodicity. We also observe that multi-invariant control uses more energy than single-invariant control.	68
A.1	In Figure 2.2, we observe that for $T_H = 4, R/Q = 2$, the trajectory makes a loop in the right gyre before ultimately reaching the goal. This is not the case for $R/Q = 3$ which is more greedy in spending energy. This is somewhat counter-intuitive, since we expect that spending more energy should drive the mobile sensor more rapidly to the goal. In this plot, we show the difference in how the control energy is spent comparing $R/Q = 2$ and $R/Q = 3$. We observe that $R/Q = 2$ takes longer to reach the goal because good performance is very sensitive to timing in the gyre. We see in the above figure that the aggressive control (grey trajectory) moves too far ahead the non aggressive control (purple trajectory) to make use of the gyre dynamics to move directly to the goal state.	98
A.2	Trajectories for various R/Q at a time horizon of $T_H = 1$	99
A.3	Trajectories for various R/Q at a time horizon of $T_H = 2$	100
A.4	Trajectories for various R/Q at a time horizon of $T_H = 3$	101
A.5	Trajectories for various R/Q at a time horizon of $T_H = 4$	102
A.6	Trajectories for various R/Q at a time horizon of $T_H = 5$	103
A.7	Trajectories for various R/Q at a time horizon of $T_H = 6$	104
A.8	Trajectories for various R/Q at a time horizon of $T_H = 7$	105
A.9	Trajectories for various R/Q at a time horizon of $T_H = 12$	106
A.10	Trajectories for various R/Q at a time horizon of $T_H = 12$ with warm start.	107

A.11	As the sensor moves in the double gyre flow field, it is constantly taking control actions $u = [u_x, u_y]$, where, u_x, u_y are the x and y-components of its actuation respectively. At each instant, the sensor is also moving over a background double gyre current vector whose components v_x, v_y are given by (1.15). The top row is a histogram of the x-component of control actions u_x taken (in red), against the x-component of the background current velocity $v_x(x_s, y_s, t)$ (in black), where, x_s, y_s are the sensor coordinates at time t . The second row is a similar plot for the y-component. We observe that the gyre takes values beyond the actuation capacity of the sensor, which highlights the under-actuated nature of the problem. Also, at low R/Q ratios, the distribution of control actions follows a distribution with two peaks at $+/- 0.1$, which corresponds to a situation similar to bang-bang control. As we increase the R/Q ratio, the distribution of control actions move to a single peak centered around zero corresponding to the use of very little control effort when compared to the background velocity.	108
B.1	This figure shows two patches of agents - one purple, and one yellow starting on opposite sides of a repelling cFTLE ridge computed with $T_A = 15$ and evolving through the unsteady flow field with control. We also plot the attracting cFTLE ridges to show the full picture of high the cFTLE ridges govern the dynamics. We see that the attracting cFTLE ridges highlight curves of long term attraction	110
B.2	We plot the passive FTLE ridges in the uncontrolled double gyre on the left, and the deformed ridges under the action of MPC at $T_H = 3.0, R/Q = 80$ on the right. The blue curves are repelling structures, and the red curves are attracting structures.	110

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to my advisors, Steve Brunton and Dr. Zhuoyuan Song. I am deeply appreciative of Steve's unwavering belief in me and his generous support across all my endeavors (from research pursuits to spontaneous karaoke). His encouragement has been invaluable throughout my journey.

I am equally grateful to Zhuoyuan for the many insightful discussions about navigating the academic landscape. His guidance and perspectives on this path have been truly beneficial. My thanks also go to my committee members Mehran Mesbahi and Xu Chen for all the helpful advice and comments on my thesis. I also thank Michelle DiBenedetto for the many conversations on turbulence and for consistently offering an open and supportive ear during my time as a graduate student.

My sincere thanks also go to my early collaborators, Aditya Nair and Eurika Kaiser, for the enjoyable experience of working on my master's project. I especially want to acknowledge Aditya for the numerous thought-provoking conversations and the unique perspectives he shared during my master's studies.

The successful completion of my PhD would not have been possible without the constant support of my parents and brother. I am profoundly grateful for their unwavering help and encouragement throughout this demanding process.

I also extend my heartfelt thanks to all the cherished friends who have remained a vital part of my life and provided an essential emotional support system. Your friendship has been a constant source of strength.

Finally, I would like to thank my guitar and harmonica.

Chapter 1

INTRODUCTION

Broadly speaking, this work investigates the use of control theory in the field of fluid dynamics. Fluid flow fields typically display spatiotemporal coherence in their dynamics that are colloquially termed *coherent structures*. While intuitively, the notion of a coherent structure can be imagined (e.g. swirling vortex patterns in fluids), there exist several mathematical definitions of coherent structures (eg. exact coherent structures, Lagrangian coherent structures, vorticity, Koopman modes etc.).

In this work, we investigate two different notions of coherent structures from dynamical systems theory. The first is the use of Lagrangian coherent structures for path planning *within* fluid flow fields. This will be the focus of majority of the text (Chapters 1-3). The second context discussed in Chapter 4 investigates the use of Koopman eigenfunctions for the control *of* a fluid flow (vortex dynamics).

In the current chapter we outline the motivating problem, literature review, background on Lagrangian coherent structures, optimal control theory and the model problem relevant to and investigated in Chapters 2 and 3. Chapter 4 makes use of the optimal control ideas developed in this chapter; however, since the motivating problem and background for Koopman eigenfunctions and vortex dynamics are quite different from the material in other chapters, they are discussed in Chapter 4, making it more self-contained. In Chapter 5 we summarize the ideas of the dissertation and provide commentary on future work.

1.1 Motivating Problem

The ability to generate energy-efficient trajectories that take advantage of the inherent motions of a background flow field has significant implications for monitoring large bodies of water with intelligent mobile sensors [11, 135, 190], furthering our understanding of the climate and natural ecosystems [35, 50, 194]. Developments in this area also present economic

opportunities for cost reduction in industries that rely heavily on maritime transport and shipping. Self-powered mobile sensors typically have complex performance tradeoffs, limiting size, weight, and power (SWAP). Further, most mobile sensors will only have partial and imperfect information about the ambient flow field, resulting in a finite-horizon predictive window to make decisions about its trajectory. Improving the generation of energy-efficient trajectories that intelligently leverage the flow field to *go with the flow* may have significant benefits in extending the duration and reach of these mobile sensing platforms. This work provides an extensive analysis of trajectories generated through a finite-horizon model predictive control (MPC) optimization of a mobile sensor in a time-varying background flow across a wide range of system parameters. Further, we establish connections between the control performance and efficiency with the alignment of these trajectories along coherent structures in the background flow.

1.2 Overview and Literature Review

For mobile agents with actuation capabilities, there exists a vast literature on various algorithms for trajectory generation in dynamic fluid environments. For example, graph search algorithms and stochastic optimization have been investigated for path planning [88, 133, 170]. Assimilating in-situ observations obtained by mobile sensors in an adaptive fashion into ocean models has also been explored, for example with mixed integer programming algorithms [96, 192]. Coordinated control of ocean gliders for adaptive ocean sensing has been exhaustively studied in Monterey bay [16, 48, 94, 95]. Algorithms inspired from computational fluid dynamics have also been used to explore coordinated control of swarms in flow fields [99–101, 163, 164].

There has been relatively little work in developing a deep understanding of the connection between the dynamics of the flow field and the nature of the optimal trajectories within the flow fields, with a few notable exceptions [70, 73, 151, 193]. A key challenge in exploring this connection is the complexity of fluid flow fields, which typically involve the existence of multiple scales in space and time.

To understand the complexity of fluids, techniques from dynamical systems are often

employed. Lagrangian coherent structures (LCS) have emerged as a robust and principled approach to uncover invariant manifolds that mediate the transport of material in unsteady fluid flows [65, 66, 68, 154, 156, 157, 171]. Specifically, LCS define the transport barriers in a flow field where passive drifters are attracted to or repelled by. There has been considerable work in the development of algorithms to accurately and efficiently compute these structures from data [24, 31, 47, 67, 68, 98, 152, 153, 156, 176].

In dynamical systems, Lyapunov exponents provide a measure of the sensitivity of the trajectory to initial conditions. In chaotic vector fields, two almost identical initial conditions can lead to flows that diverge exponentially in finite time. The finite-time Lyapunov exponent quantifies this stretching over a fixed, finite-time horizon, resulting in a scalar field over a domain of interest highlighting the most sensitive regions to perturbations in the initial conditions [31, 59, 68, 98, 156]. Moreover, the FTLE has also been used in fluid dynamics applications to compute Lagrangian coherent structures (LCS), which are finite-time analogues of invariant manifolds that mediate the transport of material in unsteady fluid flows [59, 65, 66, 68, 154, 156, 157, 171]. The LCS, and consequently the FTLE, define transport barriers in a flow field where passive drifters are attracted to or repelled from. The FTLE method has been successfully applied to bio-propulsion [189], medicine [49, 158], the spread of microbes [177], and the study of aerodynamics [140, 141], among other domains.

The ideas from both trajectory generation and the theory of LCS have been related in the past [70, 73, 151, 193]. A predecessor of this was the planning of space missions using invariant manifolds [84]. In the context of ocean transport, Inanc, Shadden, and Marsden [73] showed that the optimal trajectories of autonomous agents generated using a receding-horizon optimal control algorithm overlap with Lagrangian coherent structures. Moreover, Senatore and Ross [151] exploited this idea further to generate energy optimal paths by controlling the agents to track the background LCS. Recent papers have further explored the connections between optimal control and LCS [89, 90, 132] in the context of path planning in the ocean. However, there is still a need to better understand how the prediction horizon and relative cost of actuation in the autonomous agent optimization relate to the use of coherent structures in the unsteady background flow. These ideas will

be the main focus of Chapter 2.

Moreover, when planning trajectories, many applications aim at achieving certain objectives ranging from reaching a static goal location to maintaining certain connectivity of a multi-agent sensor network for part of or the entire mission [163, 164]. In such applications, it is still imperative to carefully choose the deployment locations since the agent’s ability to reach certain regions is largely determined by its actuation limits and the background flow dynamics. For example, it might be impossible for two groups of agents that are dominated by close-by, but different flow structures, to rendezvous. Furthermore, tuning the hyperparameters of an on-board control strategy to obtain the best performance is a challenging task. The ability to summarize and visualize the dependence of the control performance on the control hyperparameters may aid in this process. To address these challenges, in Chapter 3, we investigate the use of finite-time Lyapunov exponents (FTLEs) to quantify the performance and sensitivity of planning and control algorithms. In other words, we compute the FTLE of the controlled system.

Among the existing control methods, model predictive control [34, 55] and reinforcement learning [17, 33, 62, 134, 173] are among the most useful paradigms in modern control [25], and analyzing these control laws and *policies* with FTLE will be the focus of Chapter 3. This is in contrast with Chapter 2 where we analyze individual trajectories. We also discover a mathematical connection between optimal control and FTLEs of the controlled flow (which are different from the original flow field). This study is therefore valuable for strategies to deploy and track sensing platforms such as weather balloons, ocean floats, underwater vehicles, ocean gliders, etc.

1.3 Background

In this section, we introduce the mathematical details of an approach, each for analyzing and generating fuel-efficient trajectories for a mobile sensor in an unsteady background flow: FTLE fields and MPC. First, we introduce the computation of FTLE fields [66, 68, 156] for passive tracer particles to extract Lagrangian coherent structures from a time-varying flow field. This method is particularly important to characterize the uncontrolled behavior

of drifters in terms of finite-time attraction and repulsion behaviors. Next, we introduce the preliminaries of finite-horizon MPC, which is an online control optimization algorithm that optimizes a cost function defined over a finite-time prediction horizon. These concepts will be used throughout this dissertation

We will use MPC for trajectory optimization of a mobile sensor in an unsteady background flow. MPC is a natural choice, since the mobile sensor will have limited actuation authority, and information about the flow field will only be approximate as it is limited to a finite-time horizon. The reader might correctly observe that MPC is used extensively in this work to establish connections between energy-optimal trajectories and the FTLE. However, these connections are not unique to MPC; policies generated through other formulations of optimal control can also be related to the FTLE, as we will show with the case of reinforcement learning in Chapter 3. With this in mind, we broadly introduce the reader to optimal control theory and notions of the value function in this section.

1.3.1 Finite-Time Lyapunov Exponents

The original FTLE [66, 68, 156] primarily addressed dynamical systems of the form

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{f}(\mathbf{x}(t), t), \quad (1.1)$$

where the function $\mathbf{f}(\mathbf{x}(t), t) : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ represents the system dynamics (in our case, the vector field within which our drifter is moving in). $t \in \mathbb{R}$ represents time, and $\mathbf{x}(t) \in \mathbb{R}^n$ is the state of the system (for our purposes, the position of the drifter/sensor in a flow field). We consider $n = 2$ or 3 , depending on the dimension of space. An FTLE algorithm takes a dynamical system as input and generates a scalar variable that can be used in the computation of separatrices, known as Lagrangian coherent structures (LCS) in the flow field. These separatrices demarcate the boundaries between different regions in a domain of interest where passive tracers remain trapped [80]. This shall be seen later in Chapter 3, particularly highlighted in Figure 3.2. In the past, FTLE analysis has mostly been performed using passive drifters to study time-varying vector fields, such as ocean flows [13, 14, 122] and pollution transport models [92]. More broadly, FTLE has also been

used to compute coherent structures for a wide range of other flows [51, 59, 108, 124, 125, 139].

The FTLE field for the dynamical system $\mathbf{f}(\mathbf{x}(t), t)$ can be computed as follows. First, we initialize a grid of passive drifter particles at time t_0 and numerically integrate them through $\mathbf{f}(\mathbf{x}(t), t)$ for a fixed amount of time (i.e., the advection time) $T_A \in \mathbb{R}$, resulting in a flow map $\Phi_{t_0}^{t_0+T_A} : \mathbb{R}^n \rightarrow \mathbb{R}^n$:

$$\Phi_{t_0}^{t_0+T_A} : \mathbf{x}(t_0) \mapsto \mathbf{x}(t_0) + \int_{t_0}^{t_0+T_A} \mathbf{f}(\mathbf{x}(\tau), \tau) d\tau. \quad (1.2)$$

The operator $\Phi_{t_0}^{t_0+T_A}$ maps an initial state $\mathbf{x}(t_0)$ to a final state $\mathbf{x}(t_0 + T_A)$ by the differential equation flow induced in phase space.

Next, a Jacobian matrix of partial derivatives of the flow map, $\mathbf{D}\Phi_{t_0}^{t_0+T_A}$, is computed using finite differences for each drifter on the grid, represented by the grid node indices $i, j \in \mathbb{Z}^+$, such that

$$\begin{aligned} \left(\mathbf{D}\Phi_{t_0}^{t_0+T_A} \right)_{i,j} &= \begin{bmatrix} \frac{\Delta x_i(t_0+T_A)}{\Delta x_i(t_0)} & \frac{\Delta x_j(t_0+T_A)}{\Delta y_j(t_0)} \\ \frac{\Delta y_i(t_0+T_A)}{\Delta x_i(t_0)} & \frac{\Delta y_j(t_0+T_A)}{\Delta y_j(t_0)} \end{bmatrix} \\ &= \begin{bmatrix} \frac{x_{i+1,j}(t_0+T_A) - x_{i-1,j}(t_0+T_A)}{x_{i+1,j}(t_0) - x_{i-1,j}(t_0)} & \frac{x_{i,j+1}(t_0+T_A) - x_{i,j-1}(t_0+T_A)}{y_{i,j+1}(t_0) - y_{i,j-1}(t_0)} \\ \frac{y_{i+1,j}(t_0+T_A) - y_{i-1,j}(t_0+T_A)}{x_{i+1,j}(t_0) - x_{i-1,j}(t_0)} & \frac{y_{i,j+1}(t_0+T_A) - y_{i,j-1}(t_0+T_A)}{y_{i,j+1}(t_0) - y_{i,j-1}(t_0)} \end{bmatrix}, \end{aligned} \quad (1.3)$$

where $x, y \in \mathbb{R}$ are the horizontal and vertical components of the position vector $\mathbf{x}(t)$. The key idea here is that nearby initial conditions that rapidly separate from each other in finite advection time are highlighted as regions of large FTLE, representing repelling coherent structures. This flow map Jacobian is used to compute the Cauchy-Green deformation tensor

$$\Delta_{i,j} = \left(\mathbf{D}\Phi_{t_0}^{t_0+T_A} \right)^* \mathbf{D}\Phi_{t_0}^{t_0+T_A}, \quad (1.4)$$

where $*$ represents the matrix transpose. Finally, the largest eigenvalue λ_{\max} of $\Delta_{i,j}$ for each drifter i, j is used to compute the FTLE field:

$$\sigma_{i,j} = \frac{1}{|T_A|} \ln \sqrt{(\lambda_{\max})_{i,j}}. \quad (1.5)$$

Alternatively, $\sigma_{i,j}$ can be viewed as the maximum singular value from the singular value decomposition (SVD) of $\mathbf{D}\Phi_{t_0}^{t_0+T_A}$. It is important to note that for unsteady flow fields, the FTLE field will also vary in time, so that at each new time step a new grid of drifters must be reinitialized and advected through the flow. This procedure is typically quite expensive to compute, although there are algorithms to speed up the calculations [31, 98].

The ridges of the computed FTLE field can be extracted to visualize manifolds in the domain, known as Lagrangian coherent structures. This requires an additional step of computing the Hessian of $\sigma_{i,j}$ for ridge extraction. FTLE based on drifter particles integrated forward in time ($T_A > 0$) result in coherent structures that repel particles. Similarly, FTLE based on particles integrated backward in time ($T_A < 0$) results in attracting coherent structures. These can be seen in Figure 2.1 and Figure 3.1 as red and blue curves, where the red curves are attracting and the blue are repelling.

The FTLE field can be used to determine the LCS of an unsteady vector field [66, 68, 156]. The LCS are curves or surfaces in the domain where nearby trajectories $\mathbf{x}(t)$ are strongly attracted to or repelled from, making them time-varying analogues of stable and unstable invariant manifolds in dynamical systems theory [71]. FTLE fields and the resulting LCS are related to almost invariant sets from statistical dynamical systems [42, 52–54]. In particular, LCS act as separatrices in the flow, segmenting different regions where passive tracers remain trapped [80]. FTLE and LCS have also been extensively used to analyze ocean flows [13, 14, 122], for example, to model the spread of pollution [92]. More broadly, FTLE has also been used to compute coherent structures for a wide range of other flows [51, 59, 108, 124, 125, 139]. In this work, we will use FTLE fields generated from passive particles to investigate the trajectories of active mobile sensors, to understand how and when these sensors exploit structures in the flow field for energy-efficient transport.

1.3.2 Model Predictive Control

The dynamics of mobile sensors operating in real environments are often strongly nonlinear and subject to hardware constraints, time delays, non-minimum phase dynamics, instability, and restrictions on actuation capability. These limitations make the use of tra-

ditional linear control approaches challenging, motivating the powerful model predictive control optimization [34, 55, 78, 110] described here. In this work, we use MPC to generate trajectories for a mobile sensor in an unsteady background flow and investigate how these trajectories vary with the optimization parameters.

In general, the dynamics of a nonlinear system with actuation $\mathbf{u} \in \mathbb{R}^m$ can be written as

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{g}(\mathbf{x}(t), \mathbf{u}(t), t), \quad (1.6)$$

where \mathbf{g} is the vector field with control.

MPC is a powerful method for calculating the actuation \mathbf{u} by formulating an iterative optimization problem that minimizes a cost function over a finite-time horizon. The controller enacts this optimized actuation policy for a short time, often for a single time step, and then the optimization problem is recomputed and initialized at the current state. In this way, MPC is quite robust to model uncertainty and disturbances, as the optimization is continuously being reinitialized as new information is available about how the system actually responds to the actuation. Computing over a finite-time horizon might also make MPC more flexible and faster than a global optimization technique, especially for chaotic systems, which may result in stiff long-time optimizations. These benefits make MPC more versatile and widely used over other traditional trajectory generation algorithms. Finally, the FTLE and MPC computations are both performed over a finite time horizon, suggesting the potential for a connection between the outputs of the two algorithms.

Typically, the optimization cost for MPC can be formulated as

$$J = \mathbf{e}(t_0 + T_H)^T \mathbf{Q}_2 \mathbf{e}(t_0 + T_H) + \int_{t_0}^{t_0 + T_H} [\mathbf{e}(\tau)^T \mathbf{Q}_1 \mathbf{e}(\tau) + \mathbf{u}(\tau)^T \mathbf{R} \mathbf{u}(\tau)] d\tau, \quad (1.7)$$

subject to the system dynamics in (4.13) and control constraints imposed by physical limitations:

$$\mathbf{u}_{\min} \leq \mathbf{u}(t) \leq \mathbf{u}_{\max}. \quad (1.8)$$

Here, \mathbf{u}_{\min} and \mathbf{u}_{\max} are the minimum and maximum values the components of \mathbf{u} can take, respectively. For example, the actuators may be unable to produce thrusts beyond a certain

value. The state error is given by $\mathbf{e}(t) = \mathbf{x}(t) - \mathbf{x}_{\text{goal}}$. The finite-time horizon over which we forecast our model for the optimization is $T_H \in \mathbb{R}^+$; this term is similar to T_A , the advection time used to calculate FTLE. $\mathbf{R} \in \mathbb{R}^{m \times m}$ is a positive definite matrix that quantifies the penalty on actuation effort, and $\mathbf{Q}_1 \in \mathbb{R}^{n \times n}$ and $\mathbf{Q}_2 \in \mathbb{R}^{n \times n}$ are positive semi-definite matrices that quantify the penalty on deviations of the state from the goal throughout the trajectory and at the final time step, respectively. For computational purposes, (4.12) is often discretized. The sampling time step is Δt , the discretization of $d\tau$. In our work, the Δt serves as the “replanning frequency”, or, the time after which the a re-optimization of a trajectory is carried out. It is possible to improve the computational speed and convergence of the algorithm with a *warm start*, which uses the trajectory computed in a previous instance as the initial guess for the trajectory in the next instance [72].

1.3.3 Optimal Control Theory

Going back to Equation (1.6),

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{g}(\mathbf{x}(t), \mathbf{u}(t), t), \quad (1.9)$$

where $\mathbf{g}(\mathbf{x}(t), \mathbf{u}, t) : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R} \rightarrow \mathbb{R}^n$, $t \in \mathbb{R}$ represents time, and $\mathbf{x}(t) \in \mathbb{R}^n$ is the state of the agent. The agent is also able to apply actuation (such as propulsion) that is modeled as $\mathbf{u} \in \mathbb{R}^m$ as explained later in Section 1.4.

In optimal control, one seeks to find a control policy $\mathbf{u}(\mathbf{x}, t)$ that minimizes a cost function J . Intuitively, this function often penalizes two terms, the distance from the goal we wish to drive the system states toward and the energy spent. Therefore, when the control is found which minimizes this function, the agent moves toward the goal in an energy efficient manner. In our work, we primarily consider a linear–quadratic regulator (LQR) type quadratic cost function given by

$$J_{LQR}(\mathbf{x}_0, t) = \int_{t_0}^{t_0+T_H} \left[\mathbf{e}(\mathbf{x}_0, \tau)^T \mathbf{Q} \mathbf{e}(\mathbf{x}_0, \tau) + \mathbf{u}(\tau)^T \mathbf{R} \mathbf{u}(\tau) \right] d\tau, \quad (1.10)$$

where $\mathbf{x}_0 \in \mathbb{R}^n$ is the initial spatial location of an agent, $t_0 \in \mathbb{R}$ is the initial time, $\mathbf{e}(\mathbf{x}_0, t) \triangleq$

$\mathbf{x}(\mathbf{x}_0, t) - \mathbf{x}_{\text{goal}}$ is the distance from goal or the state tracking error starting from an initial condition \mathbf{x}_0 , $\mathbf{u} \in \mathbb{R}^m$ is the control vector, $T_H \in \mathbb{R}$ is the time horizon, $\mathbf{Q} \in \mathbb{R}^{n \times n}$ is the state penalty matrix, and $\mathbf{R} \in \mathbb{R}^{m \times m}$ is a matrix of control penalty.

Functions of the form described by (1.10) are *value functions* [173], which can also be viewed theoretically as solutions to the Hamilton-Jacobi Bellman (HJB) equations [25]. These functions assign a scalar value to each initial condition in the domain corresponding to the future cost that will be accrued from that particular initial condition following the optimal policy. The field of deep reinforcement learning often involves estimating these functions using neural networks from limited data of the agent moving through the domain or from having limited-to-no knowledge of the governing dynamics. Analytically, some useful mathematical relations have been established [167, 179] when using cost functions of the above form, coupled with control-affine or kinematic models highlighted in Eq. (3.1). One major result is

$$\mathbf{u}^*(\mathbf{x}_0, t) = -\mathbf{R}^{-1} \nabla J_{LQR}(\mathbf{x}_0, t), \quad (1.11)$$

which directly relates the spatial gradient of the cost function in the domain to the *optimal* control law $\mathbf{u}^*(\mathbf{x}_0, t)$ that minimizes Eq. (1.10). We will use this result later to establish connections to FTLE. When $\mathbf{u}(\mathbf{x}, t)$ is plotted spatially, it can be viewed as a *policy*. To end this section, we note that MPC (the previous section) is actually a *numerical method* for computing the theoretical ideas discussed in this section.

1.4 Model Problem

We now discuss the models used to simulate the agent dynamics and the unsteady flow field the mobile sensor operates within. We also provide specific parameters that are used for all numerical experiments.

1.4.1 Sensor Dynamics

In a two-dimensional setting, a simple kinematic model for the dynamics of the mobile sensor is given by adding the velocity due to actuation, $\mathbf{u}(t)$, to the background flow velocity

$\mathbf{v}(\mathbf{x}(t), t) : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^2$:

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{v}(\mathbf{x}(t), t) + \mathbf{u}(t). \quad (1.12)$$

The state $\mathbf{x}(t) = [x, y] \in \mathbb{R}^2$ is the position vector. The key assumption in this model is that, without control, the velocity of the sensor, $d\mathbf{x}/dt$, matches the velocity of the background fluid flow. Thus, the uncontrolled mobile sensor can be considered as a passive Lagrangian drifter, and (1.12) degenerates to (1.1) when $\mathbf{u}(t) = 0$ all all times. Moreover, it assumes that the sensor can generate its own relative velocity $\mathbf{u}(t) = [u_x, u_y] \in \mathbb{R}^2$ in addition to the flow-induced velocity. It is possible to develop more sophisticated models for the mobile sensor dynamics that include inertial and rotational dynamics; in Zhang et al. [193], it was shown that trajectories based on such models also show strong correlation with the presence of background LCS.

1.4.2 Double Gyre Flow Field

As a running example through this text, we will investigate the motion of the mobile sensor above in the unsteady double gyre flow field described here. The double gyre flow is an analytically defined, periodic vector field that is often used to study mixing and coherent structures related to those found in geophysical circulations. In particular, the double gyre represents a typical large-scale ocean circulation phenomenon often observed in the northern mid-latitude ocean basins. This circulation is quite dominant and is persistent, consisting of sub-polar and sub-tropical gyres. As a major type of ocean circulation, several main features of the double gyre phenomena have been identified through analyzing observational data and numerical simulations [74, 165, 166].

The double gyre velocity field is derived from the stream function

$$\phi(x, y, t) = A \sin(\pi f(x, t)) \sin(\pi y), \quad (1.13)$$

where the time dependency is introduced by

$$f(x, t) = a(t)x^2 + b(t)x, \quad (1.14)$$

with time dependent coefficients

$$a(t) = \epsilon \sin(\omega t) \quad \text{and} \quad b(t) = 1 - 2\epsilon \sin(\omega t).$$

This flow is defined on a nondimensionalized domain of $[0, 2] \times [0, 1]$, where $A, \epsilon, \omega, x, y, t \in \mathbb{R}$. Here, ϵ dictates the magnitude of oscillation in the x -direction, ω is the angular oscillation frequency, and A controls the velocity magnitude. Unless stated otherwise, the parameters used for the double gyre flow field are as in Shadden et al. [156], where $A = 0.1$, $\epsilon = 0.25$, and $\omega = 2\pi/10$. The resulting velocity field is given by

$$\mathbf{v}(x, y, t) = \begin{bmatrix} -\frac{\partial \phi}{\partial y} \\ \frac{\partial \phi}{\partial x} \end{bmatrix} = \begin{bmatrix} -\pi A \sin(\pi f(x, t)) \cos(\pi y) \\ \pi A \cos(\pi f(x, t)) \sin(\pi y) \end{bmatrix}. \quad (1.15)$$

1.4.3 Specific Control Objective

By combining the mobile sensor model and the double gyre flow field, the dynamics of the sensor are given by,

$$\frac{d}{dt} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} -\pi A \sin(\pi f(x, t)) \cos(\pi y) \\ \pi A \cos(\pi f(x, t)) \sin(\pi y) \end{bmatrix} + \begin{bmatrix} u_x \\ u_y \end{bmatrix}. \quad (1.16)$$

The objective in a lot of this thesis is to move a mobile sensor from a starting location at coordinates $\mathbf{x}_{\text{start}} = [2, 1]$ to a goal location at $\mathbf{x}_{\text{goal}} = [0.5, 0.5]$. The cost function is given by

$$J = \int_{t_0}^{t_0+T_H} [\mathbf{e}(\tau)^T \mathbf{Q} \mathbf{e}(\tau) + \mathbf{u}(\tau)^T \mathbf{R} \mathbf{u}(\tau)] d\tau, \quad (1.17)$$

where $\mathbf{e}(t) \triangleq \mathbf{x}(t) - \mathbf{x}_{\text{goal}}$ is the state tracking error. This cost function is subject to constraints on the actuation of the mobile sensor

$$|u_x| \leq 0.1 \quad \text{and} \quad |u_y| \leq 0.1,$$

which ensure that the maximum sensor velocity is significantly smaller than the largest background flow field velocity, $\pi A \approx 0.314$. This constraint is imposed to model the lim-

ited actuation available in real world scenarios. The discretized time step across all double gyre simulations is $\Delta t = 0.1$. Each simulation was run for 800 steps. In Chapter 3, for each agent, the control algorithm computes $\mathbf{u}(\mathbf{x}, t)$ numerically for a grid of initial conditions in the domain. To solve the resulting optimization problems of MPC, we use the CasADi [2] and MPCTools [138] packages, which use an interior point filter-line search algorithm (IPOPT). The discretized time step used for RK4 methods and for direct multiple shooting collocation in MPC across all double gyre simulations is $\Delta t = 0.1$. In later advanced examples with reinforcement learning, we use the stable baselines package [130].

A relevant aspect of MPC is that it computes the optimization over a finite time horizon T_H , which is similar to how FTLE performs its computations over a time of advection T_A . From Eq. (1.10), in our case, $\mathbf{Q} = Q\mathbf{I}_{2 \times 2}$ and $\mathbf{R} = R\mathbf{I}_{2 \times 2}$, where \mathbf{I} is the identity matrix. We combine these two parameters into one parameter, namely the R/Q , the specific values for which will be elaborated on as they appear in the text.

Chapter 2

INTERPLAY OF SENSOR DYNAMICS WITH LAGRANGIAN COHERENT STRUCTURES

2.1 Introduction

In this chapter (drawing the results from [87]), we investigate the explicit connections between finite-horizon energy-optimal trajectories of a mobile sensor and the underlying background flow dynamics. We specifically analyze how key parameters of the MPC-based optimization affect how the resulting autonomous agent trajectory utilizes unsteady fluid coherent structures for energy-efficient transport. This analysis is performed primarily on the double gyre flow field, which is a testbed to understand mixing and transport in the ocean. Subsequently, we also further verify our observations on advanced flow fields including the Arnold–Beltrami–Childress (ABC) flow, a three-dimensional incompressible analytical flow field, and unsteady flow data from the Gulf of Mexico.

A summary of our methodology is shown in Figure 2.1. The choice of MPC is particularly relevant in this work, as both the FTLE and MPC rely on finite-time horizons in their computations. To explore this connection, we perform an exhaustive search through several of the trajectory optimization parameters that are important to practitioners, including the prediction time horizon and step size for MPC, the relative cost of actuation versus state tracking error, and the maximum agent velocity. We find that there are strong correlations between the presence of background FTLE ridges and the actuation energy expenditure at the corresponding locations along the trajectory.

This chapter is organized as follows. In Section 1.3, the core methodology of MPC and FTLE are discussed. MPC will be the primary optimization algorithm used to generate trajectories, and these will be analyzed using FTLE fields. Section 1.4 describes models for the mobile sensor dynamics and actuation, along with the dynamics of the unsteady double gyre background flow field. The main results are presented in Section 2.2, including in-

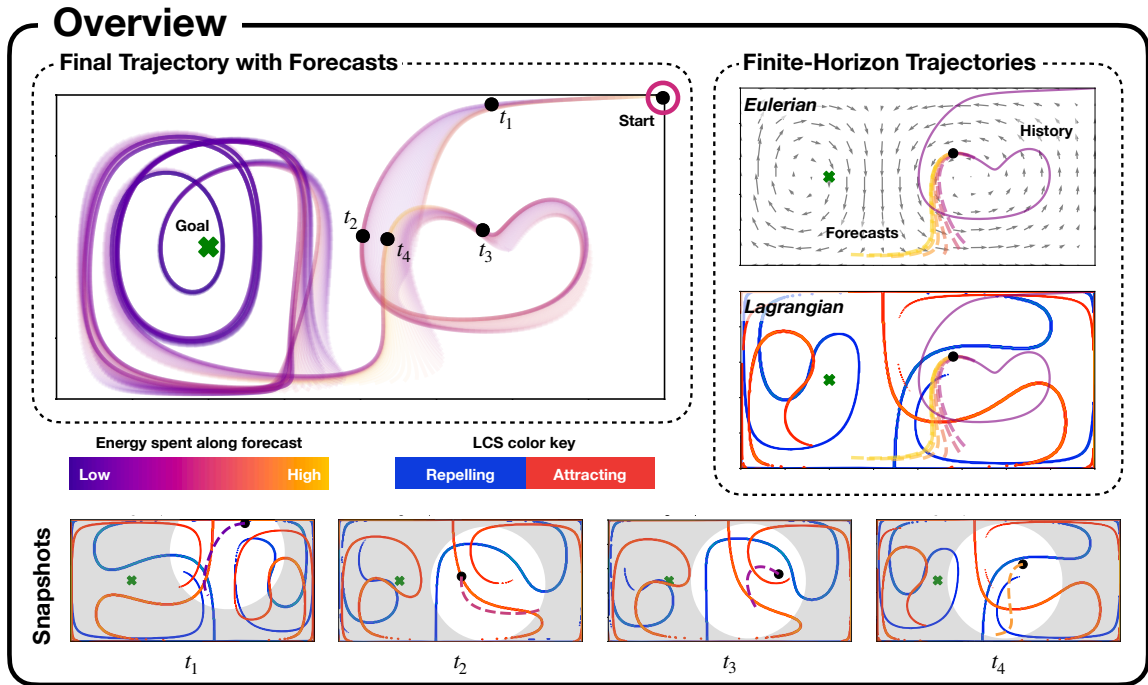


Figure 2.1: Overview of the proposed methodology for analyzing the connections between finite-horizon energy optimal trajectories and the FTLE field. A self-propelling agent is controlled to transit from a starting location to a goal location through a finite-time horizon energy-optimal trajectory in a time-varying double gyre flow field. The resulting agent trajectory, along with the finite-horizon predicted trajectories at each time step, are shown and color-coded based on instantaneous energy expenditure (top left). The trajectory history (solid) and the future forecast trajectory bundle (dashed) at an example time instant are shown (top right); the instantaneous FTLE ridges are also shown below these with blue indicating the repelling LCS and red indicating the attracting LCS. As can be observed from the snapshots taken at four particular times (bottom), the energy expenditure along the planned trajectory, given by the color of the dashed line), and the shape of the finite-horizon trajectory depend on the evolution of the local FTLE ridges.

depth analysis of trajectories generated across a wide range of system parameters. In particular, the time horizon of the MPC optimization, the relative cost of actuation versus state tracking error, and the frequency of the background flow oscillation are all investigated. In Section 2.3, we highlight the use of MPC on other flow fields and demonstrate our observations on the ABC and Gulf of Mexico flow fields. Section 2.4 provides a summary of results and a discussion of limitations with suggestions for future work. Appendix A also

provides additional plots and analysis of the data that was not presented in the main text.

2.2 Double Gyre Results

In this section, we examine energy-efficient trajectories for an active mobile sensor generated using MPC across a range of hyperparameters, including the prediction horizon, penalty weights on the state error and control effort, and the double gyre oscillation frequency. The MPC results for this thorough parameter sweep are presented in Section 2.2.4 and summarized in Figure 2.7. Our goal is to understand the sensitivity of the trajectory to parameters and to uncover performance tradeoffs, for example with the time horizon of optimization. We find a large *sweet spot* where effective, energy-efficient trajectories are generated. Further, we establish connections between the efficient mobile sensor trajectories and the Lagrangian coherent structures of the underlying flow field.

2.2.1 Trajectories with Different Relative Actuation Cost, R/Q

Figure 2.2 shows the effect of varying the ratio of control effort penalty R to the state error penalty Q on the trajectories, for a fixed time horizon of $T_H = 4$; similar plots for a range of time horizons from $T_H = 1$ to $T_H = 12$ are shown in Figures A.2–A.10 in the Appendix. The ratio R/Q quantifies the relative cost of actuation, and varying this parameter is important to understand performance tradeoffs when the mobile sensor has a limited actuation budget. As R/Q is increased, corresponding to actuation being more expensive, the agent actuates less, and the state tracking error increases. This increase in state tracking error tends to correspond to larger steady-state limit cycles about the goal state. The weighted actuation cost $J_u = \sum Ru^T u \Delta t$ increases with R/Q , as we fix $Q = 1$ and increase R ; however, the unweighted actuation $\sum u^T u \Delta t$ decreases with R/Q . Importantly, the trend of cost versus R/Q is not strictly monotonic, and there are discontinuous jumps corresponding to bifurcations in the orbit; the non-monotonic behavior and bifurcations are more pronounced for other T_H in the Appendix. For small R/Q values such as $R/Q = 2$ and $R/Q = 3$, the agent moves around the goal state in a tight orbit, and this orbit continuously expands as R/Q increases, as shown for $R/Q = 15$. However, between $R/Q = 25$

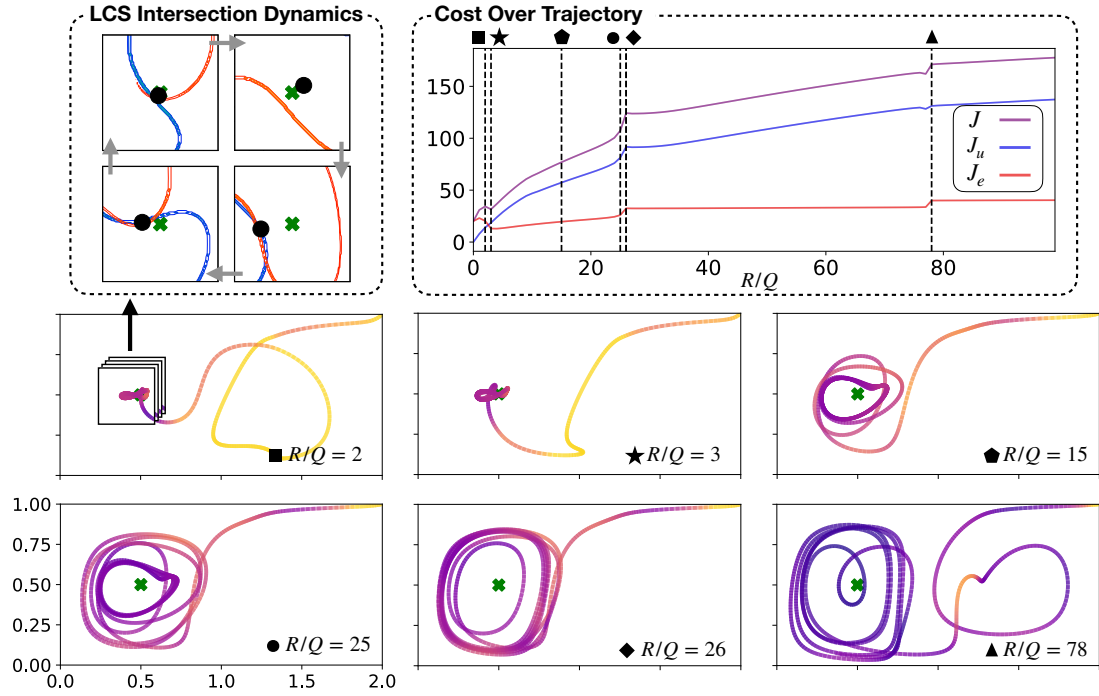


Figure 2.2: Dependency of the resulting sensor trajectories on the ratio between the penalties on energy expenditure (R) and state error (Q), for a fixed time horizon $T_H = 4$. The trajectories are color-coded by instantaneous energy expenditure. There are three costs shown in the top right figure: $J = J_u + J_e$, $J_e = \sum Q(x - x_g)^T(x - x_g)\Delta t$, and $J_u = \sum Ru^T u\Delta t$. Here, J generally increases with R/Q . The trajectories undergo several qualitative changes (bifurcations) as R/Q is increased, forming different types of periodic orbits shown on the bottom of the figure. An example of this is when R/Q is changed from 25 to 26, we observe a major change in the shape of the final orbit around the goal, as opposed to the minor change from $R/Q = 15$ to 25, where the final eye-shaped orbit only gradually increases in size. The formation of these orbits are dependent on the background flow FTLE, as can be seen from the inset of case $R/Q = 2$. We observe that as the LCS move, the stable and unstable LCS intersect at a point, whose location changes every instant, and the sensor moves in a manner in which it is right on top of this intersection point for most of the time when $R/Q = 2$ (top left box).

and $R/Q = 26$ the trajectory undergoes a rapid qualitative change, where the radius of the orbit around the goal state jumps.

It is interesting to note in Figure 2.2 that the $R/Q = 2$ agent has an initial loop in the right basin, while the $R/Q = 3$ agent does not. This behavior is counter-intuitive, as the $R/Q = 2$ agent should expend control more freely, and thus more aggressively seek the

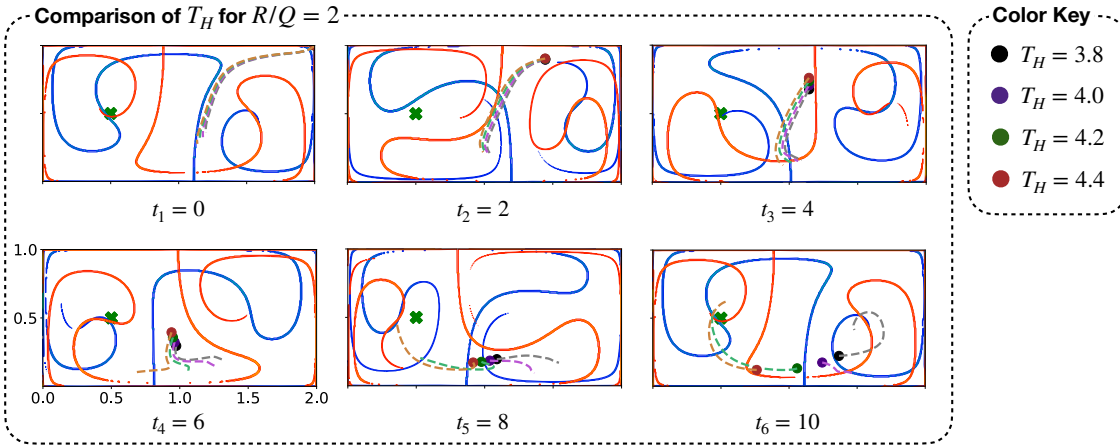


Figure 2.3: Different instances of a trajectory for different time horizons while keeping $R/Q = 2$ fixed. By varying the time horizon, T_H , we see that the extra loop in Figure 2.2 is due to a sensitivity of the planned path with respect to T_H , where the agent becomes stuck in the right gyre for lower T_H . We can see here with reference to Figure A.1 how as the time horizon increases, the agent moves less aggressively and improves its timing with the gyre oscillations to reach the goal sooner without taking an extra loop. Another interesting aspect of this plot can be seen at $t_3 = 4$ and $t_4 = 6$ where the multiple agents line up to mirror the movement of the red LCS, further confirming the strong correlation finite-horizon optimal trajectories and the FTLE ridges.

goal state. As shown in Figure A.1 in the Appendix, the more aggressive agent does move away from the starting state faster initially; however, it becomes trapped on the side of a repelling LCS farther away from the goal location and must make an entire orbit around the right gyre before approaching the goal state. The maximum agent velocity is smaller than the maximum gyre velocity, so even the most aggressive agents are unable to break out of the right gyre without precise timing. This type of bifurcation also occurs for fixed $R/Q = 2$ by varying the time horizon, as in Figure 2.3. In this case, the behavior is more consistent with intuition, as the longer time horizon trajectories avoid being trapped in the right gyre.

Previous work [73] suggests that low-energy trajectories tend to coincide with the LCS of the background flow. In our example, even for $R/Q = 2$ and $R/Q = 3$, the mobile agent can be seen aligning with and exploiting the coherent structures. For example, in the top left of Figure 2.2, the $R/Q = 2$ sensor moves along on the intersection of the attracting and

repelling LCS as it orbits the goal state. In the next section, we will see that the agent also precisely times its actuation before and after crossing a repelling LCS to take advantage of the background drift.

2.2.2 *Instantaneous Energy v.s. FTLE Ridge*

Given the existing connection of low-energy trajectories and FTLE ridges, we are interested in how the energy is utilized along a trajectory. Figure 2.4 shows how the agent ‘schedules’ an increase in actuation to cross a repelling (blue) FTLE ridge. After crossing, the agent decreases its actuation, as it is naturally repelled from the blue ridge and attracted by the red ridge into the left basin. Similar timing and utilization of the FTLE ridges is observed for a wide range of time horizons and control aggressiveness. This occurs in the double gyre only across the *repelling* hyperbolic LCS near the middle of the domain, as opposed to the shear-driven LCS near the goal location. This can be observed across the all gyre trajectories, where the yellow (high energy) regions in the trajectory are in the middle of the domain.

2.2.3 *Periodic Orbits*

We observe that controlled trajectories often form periodic orbits around the goal state, as seen in Figures 2.2, 2.5, and 2.6. Because the background flow field is periodic, the agent would require constant actuation to stay fixed at the goal state. Instead, the agent trajectory tends to form a periodic orbit around the goal, balancing state tracking error and control expenditure. Typically, this orbit is larger for agents with a tighter energy budget (i.e., for larger R/Q). Many past studies have focused on trajectory planning where the final state is fixed at the goal. However, given the constantly evolving background flow field and its dominant effect on mobile sensor dynamics, it is important in practice, to consider the cases where the final state cannot be fixed. Figure 2.6 also indicates that the shape of the final periodic orbit depends on the frequency of the double gyre oscillation, with the frequency of the agent orbit synchronizing with the gyre frequency. In the other example flows observed below, similar periodic orbits are observed, where the agent *loiters* around

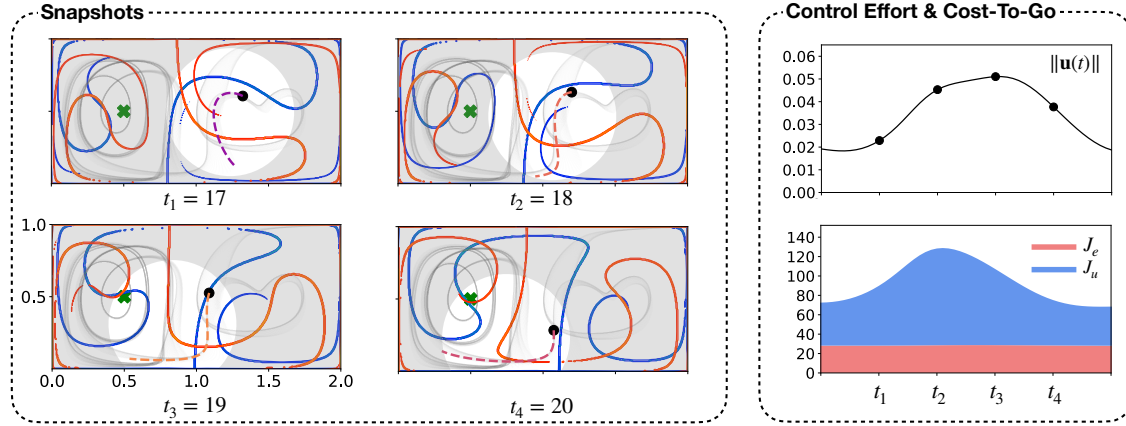


Figure 2.4: Influence of the FTLE ridges on the energy expenditure both along prediction horizon (right top) and instantaneously (right bottom) under parameters $T_H = 4.0$, $R/Q = 100$, and $\Delta t = 0.1$. An agent's motion along with the predicted forecast trajectory (dashed line) are shown on the left, together with the repelling (blue) LCS and the attracting (red) LCS. There is a correlation between the spike in both the instantaneous energy spent (right top) and the cost along the forecast trajectory (right bottom) with movement across an FTLE ridge. Here, unlike in Figure 2.2, the summations $J_e = \sum Q(x - x_g)^T(x - x_g)\Delta t$ and $J_u = \sum Ru^T u\Delta t$ are only along the forward dashed line in the plots under 'snapshots' and not along the entire trajectory as in Figure 2.2

the goal state. It will be interesting to investigate these orbits in more detail, including the classes of flows they exist in, and the conditions under which they bifurcate.

2.2.4 MPC Parameter Sweep

We now present an exhaustive sweep through two of the most critical parameters for MPC, the prediction horizon T_H and the cost function penalty ratio R/Q , for different gyre oscillation frequency ω . The first two parameters are related to the power and prediction capability of the mobile sensor, and the third parameter characterizes the unsteadiness of the background flow. We perform a full parameter sweep for the time horizon ($T_H \in [0, 10]$) and the cost penalty ratio $R/Q \in [0, 100]$, for the double gyre frequency $\omega \in [\pi/6, \pi/3]$. For each parameter value, we compute the state tracking error and the (unweighted) actuation energy expenditure, integrated along the entire trajectory.

Figure 2.7 shows the results from the MPC parameter sweep. For all time horizons




Figure 2.5: The mobile sensor settle on periodic orbits around the goal state (left) and the magnitude of the Fourier transform of the instantaneous energy spent by the mobile sensor (right). We observe that the time series of the energy spent is periodic with frequencies at integer multiples of the double gyre oscillation frequency, which correspond to the peaks in the right plot.




Figure 2.6: Variations of agent trajectory under different double gyre oscillation frequencies. The frequency of the periodic orbits depends on the double gyre oscillation frequency.

and gyre frequencies, we observe that the trajectories sweep out a Pareto front in control expenditure versus state tracking error as R/Q is varied logarithmically from 0 to 100. The bottom row of Figure 2.7 shows three representative trajectories along the Pareto front. As R/Q is increased, there is often a sharp drop in control cost with a relatively small increase in state tracking error, suggesting that there are energy-efficient trajectories that achieve relatively good tracking performance. However, we observe a break point in this monotonic trend, beyond which increasing R/Q results in rapid deterioration of the state error with relatively little decrease in control cost. This break point corresponds to the scenario where the motion of the sensor is dominated by the background flow, and the

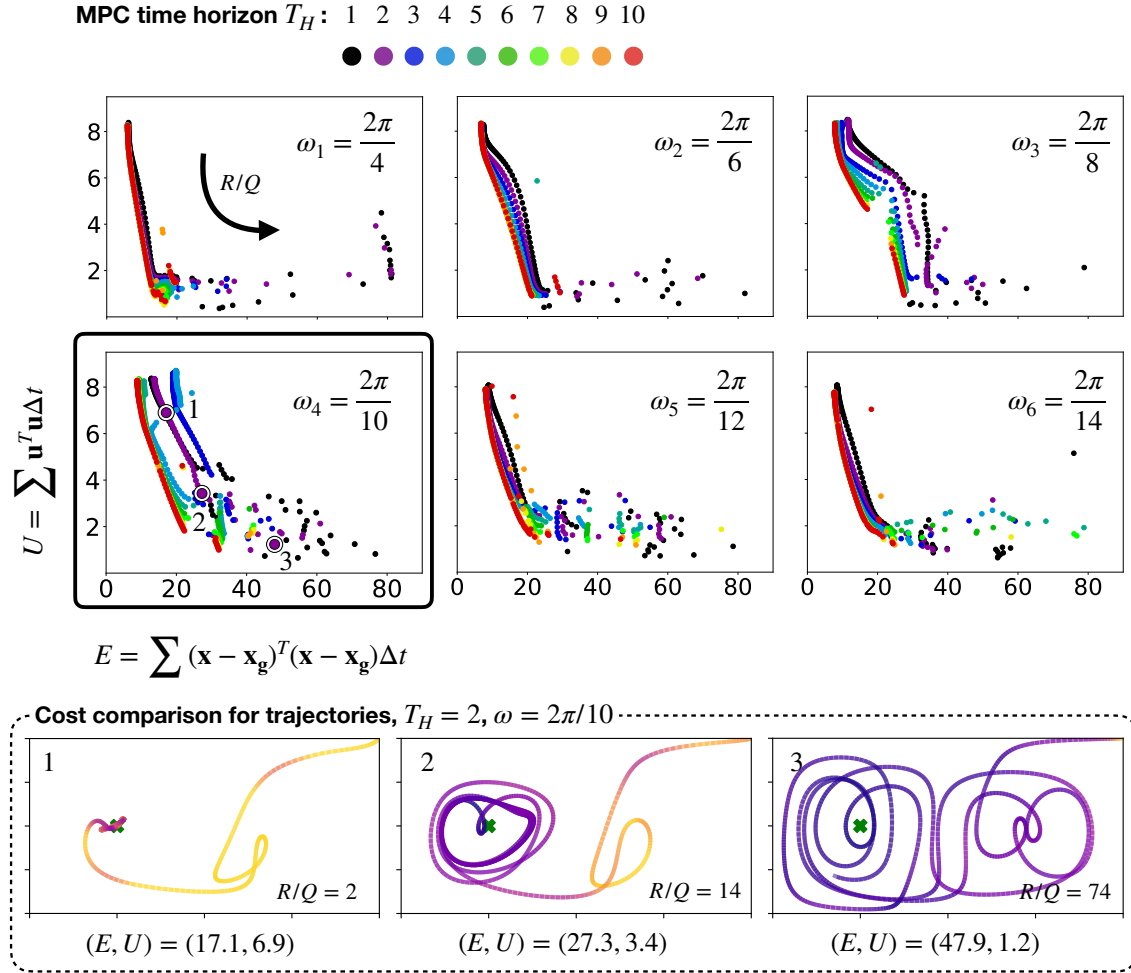


Figure 2.7: Multiple simulations were carried out at each R/Q ratio spaced logarithmically, from 0 to 100, time horizon ranging from 1 to 10, and gyre frequency ranging from $2\pi/4$ to $2\pi/14$. The data presented here is in the form of scatter plots for each gyre frequency with each color representing the Pareto optimal tradeoff curve between the total energy spent along each trajectory and the sum of deviations from target along the trajectory. The trajectories shown in the bottom row correspond to the highlighted purple circles (1,2,3) in the Pareto optimal corresponding to $\omega_4 = 2\pi/10$.

chaotic nature of the flow field dominates the state and energy errors. This phenomenon is more evident for smaller time horizons.

It is observed that longer prediction horizons produce trajectories that are more energy efficient with smaller state errors. This is expected, as longer time horizons include more

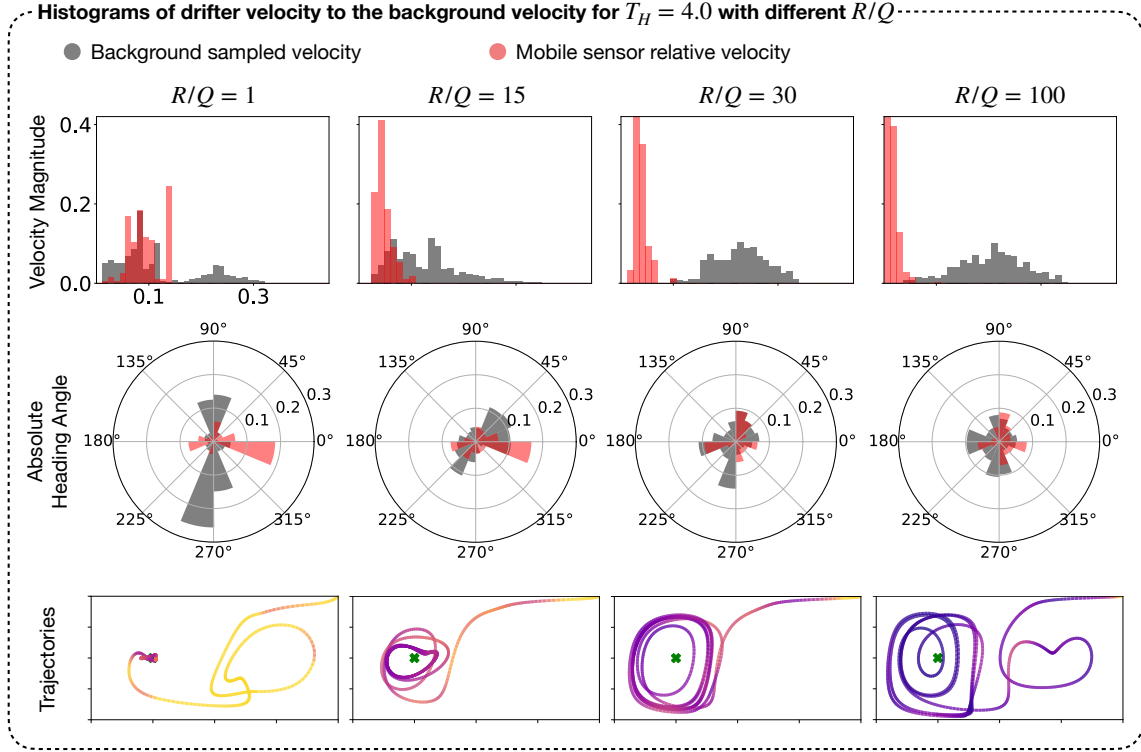


Figure 2.8: As the sensor moves in the double gyre flow field, it is constantly taking control actions $u = [u_x, u_y]$, where, u_x, u_y are the x and y-components of its actuation respectively. At each instant, the sensor is also moving over a background double gyre flow velocity vector whose components, v_x and v_y , are given by (1.15). The top row of histograms are of the magnitude of control actions $\|u\|$ taken (in red), against the magnitude of the background current velocity $\|v(x_s, y_s, t)\|$ (in grey), where, x_s, y_s are the sensor coordinates at time t . The second row shows the heading angle of the sensor (the orientation of dx/dt), plotted in red, against the orientation of the background flow field velocity vector, plotted in grey. The corresponding trajectories are shown on the bottom.

information about the flow field in the optimization. This trend is weaker for small-to-moderate R/Q and is more pronounced for larger R/Q . The shape of the Pareto curve also changes with the double gyre frequency. This shape change is particularly evident for moderate frequencies, suggesting a “resonance” in the interaction of trajectories with the background flow. Resonance with changing gyre frequency has been explored in the context of inertial particles in the double gyre flow [171].

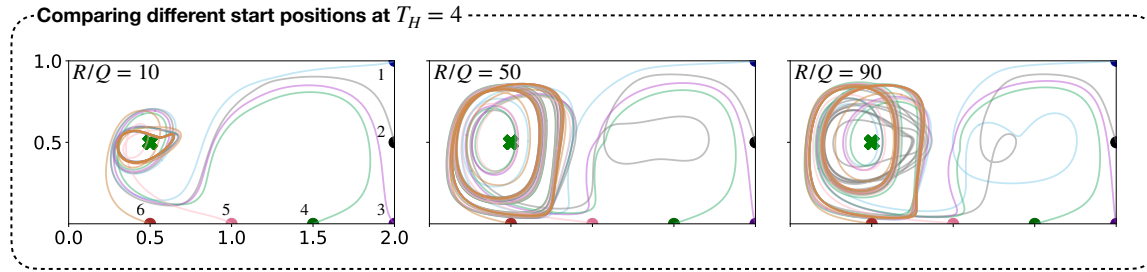


Figure 2.9: MPC trajectories with varying R/Q ratio, starting from different initial conditions with the same goal location. We find that eventually, the trajectories converge to similar periodic orbits thereby using the LCS in similar ways to orbit around the goal location despite having different transients.

2.2.5 Sensor Velocity

To gain further insight into the dependency of sensor actuation velocity on the background flow velocity, we compare their distributions along the resulting trajectory at different R/Q values. Figure 2.8 shows histograms of the magnitude and orientation of the sensor actuation velocity versus the background flow velocity, for a range of R/Q values. It can be observed that more aggressive agents with smaller R/Q have larger actuation velocity magnitudes and tend to move perpendicular to the background flow. Agents with larger R/Q , corresponding to more conservative actuation policies, tend to have smaller actuation velocity and align their actuation in the direction of the flow field to take advantage of the background flow. Except in the most aggressive $R/Q = 1$ case, the mobile sensor rarely uses the maximum control velocity. Additional plots with the x - and y -components of the agent velocity are presented in Figure A.11 in the Appendix.

2.2.6 Different Start and End Locations

Figure 2.9 shows the MPC optimized trajectories for six different starting locations along the right and lower boundaries of the domain. Although the paths have different initial transients, the trajectories evolve onto the same periodic loitering orbits around the goal state, indicating that they are ultimately leveraging similar flow structures. Similarly, Figure 2.10 shows the MPC optimized trajectories for several different goal states. Some goal

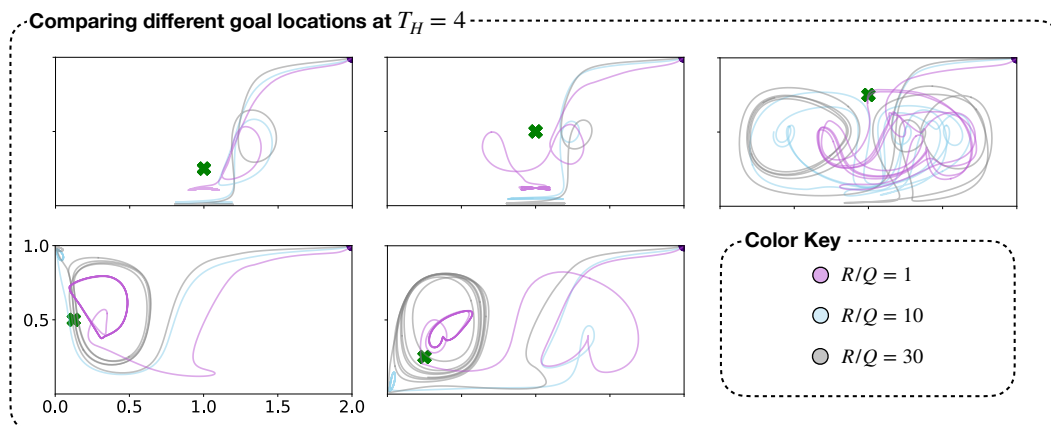


Figure 2.10: MPC trajectories to different goal locations with varying R/Q ratio for the same initial condition. We find that in the case of placing the goal near the middle to bottom region (top left and top middle plots) of double gyre, the sensor is able to form stable orbits near the attracting LCS, where the sensor moves with the base of the attracting LCS. However, placing the goal near the repelling LCS causes more difficulty for MPC in forming small stable orbits (top right). The bottom plots show that it is possible to form small periodic orbits in the corners of the double gyre flow field when the goal is placed close to them.

states are much more difficult to reach than others, because of the strong unsteady background flow field. With more aggressive control ($R/Q = 1$) the trajectories generally form tighter loitering orbits. However, for the case when the goal state is in the upper middle of the domain, it is clear that none of the MPC trajectories are able to find a suitable loitering pattern. This diversity of orbits highlights the importance of choosing a suitable goal location, which would likely involve a higher level of planning.

2.3 Advanced Test Cases

From the analyses in the double gyre flow field, we observed the strong dependency of energy-efficient MPC trajectories on the LCS of unsteady periodic flow fields. In this section, we study the use of MPC on more challenging test cases: an analytical three-dimensional incompressible flow field and a real-world flow field reconstructed from ocean model datasets. The goal is to demonstrate how and when these results generalize to better understand the dependency of energy-efficient MPC trajectories on FTLE-based LCS.

2.3.1 Arnold-Beltrami-Childress

We now demonstrate the use of MPC on the ABC flow field [32, 136], which is an incompressible model for a flow evolving in a three-dimensional periodic domain. It has been studied exhaustively in the past as a stepping stone to understanding turbulent flow fields. The three-dimensional space contains six interwoven vortices. An important feature of this flow field is that even the steady version of the flow field can give rise to chaotic trajectories. However, for our purposes, we investigate the unsteady case. Similar to (1.16) for the double gyre flow field, the equation for a mobile sensor evolving in the ABC flow is given by

$$\frac{d}{dt} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} A(t) \sin(z) + C \cos(y) \\ B \sin(x) + A(t) \cos(z) \\ C \sin(y) + B \cos(x) \end{bmatrix} + \begin{bmatrix} u_x \\ u_y \\ u_z \end{bmatrix} \quad (2.1)$$

with parameters $A, B, C, \epsilon, \omega \in \mathbb{R}$, where $A(t) = A + \epsilon \cos(\omega t)$ is a time-varying component that makes the flow field unsteady. We investigate trajectories in the regime where $A : B : C = \sqrt{3} : \sqrt{2} : 1$, which has been exhaustively studied numerically and analytically (with $C = 0.1$ specifically in our simulations), $u_x, u_y, u_z \leq A + B + C$ and $\omega = 2\pi/10$. All ABC simulations presented were run for 2000 time steps with a step size of $\Delta t = 0.1$. The cost function used was the same as Eq. 1.17. In Figure 2.11, we show the trajectories planned by the MPC for a time horizon $T_H = 5$. Across these simulations, we observed that, similar to the double gyre, that it is possible to form loitering orbits close to the goal point. These loitering orbits become larger for larger R/Q . Results across different time horizons and R/Q are further summarized in Figure 2.12. We found that relatively short time horizon (compared to the period of oscillation $T = 10$) trajectories are able to reach the goal state, and similar to the double gyre, longer time horizons reduce error with lesser energy consumption. The inflection points and breaking off of points from the curve (for example, the black dots near $E = 1500$) correspond to drastic changes in trajectory shape (bifurcations).

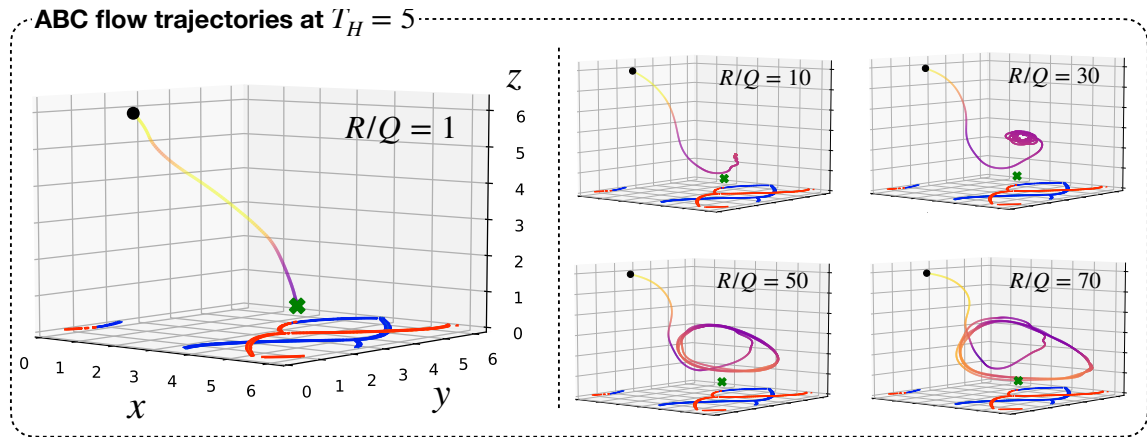


Figure 2.11: MPC trajectories formed by the sensor in an ABC flow field as a function of the R/Q ratio. We observe that for a time horizon of 5, we are able to find several cases of periodic orbits loitering close to the goal location. In these cases, the initial position is $\mathbf{x}_{\text{start}} = [\pi/2, 1, 6]$ and the goal is $\mathbf{x}_{\text{goal}} = [5, 2, 1]$.

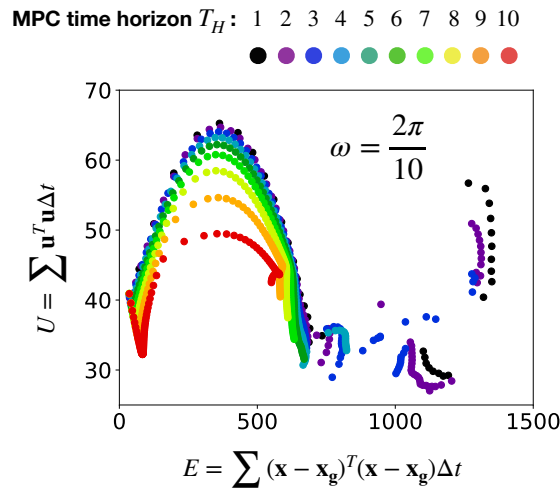


Figure 2.12: This scatter plot shows the performance of MPC in the ABC flow field similar to Figure 2.7. We observe that MPC trajectories with short time horizons are able to reach the goal state and increasing the time horizon has a benefit of decreasing the actuation energy usage. We also observe that the inflection points and breaking off points correspond to bifurcations in the trajectories.

2.3.2 *Gulf of Mexico*

For the final example, we consider the Gulf of Mexico surface velocity estimates from the HYbrid Coordinate Ocean Model (HYCOM). This data-assimilative model synthesizes remotely sensed and in-situ measurements on a hybrid coordinate system. We used daily $1/12.5^\circ$ -resolution data from the HYCOM 1992-1995 experiment 19.0 to generate a vector field. We then used linear interpolation on this vector field in space and time to generate a function that could be used for model predictive control. The parameters were chosen to be $\Delta t = 0.1$ day, $T_H = 0.4$ day, $u_x, u_y \leq 2$ km/hr, and $R/Q = 1$. The step size $\Delta t = 2.4$ hours for the MPC. The full trajectories in Figure 2.14 were computed for 1000 time steps (100 days).

We observe similar spiking behaviour in the energy spent when moving across a repelling LCS, which is seen in panels $t = 6$ to $t = 10$. Here, t is the time in days (specifically, $t = 0$ corresponds to the flow field on Day-1 of the HYCOM dataset). We observe that the sensor synchronizes with the attracting LCS to move towards the goal location. In periodic flow fields such as the double gyre, we observe the formation of periodic orbits. However, in case of the Gulf of Mexico (aperiodic), we observed the formation of an aperiodic loitering trajectory in proximity to the goal location. The time horizon here is relatively short compared to the total trajectory time. We can see that even with short time horizons the MPC trajectories are capable of making it to the goal location. Although the trajectory does evolve towards a loitering orbit near the goal state, it can be seen in Figure 2.14 that for more aggressive control ($R/Q = 0.5$ and $R/Q = 0.1$), the trajectory is able to get much closer to the goal state, with a tighter orbit.

2.4 *Discussion*

In this work, we have investigated the behavior of trajectories optimized over a finite horizon for a controlled mobile sensor in unsteady flow fields, as both the control and flow field parameters were varied. A thorough study was conducted on the double gyre flow field, and selected results were further verified on more advanced flow fields such as the ABC and the Gulf of Mexico. In particular, finite-time model predictive control was used

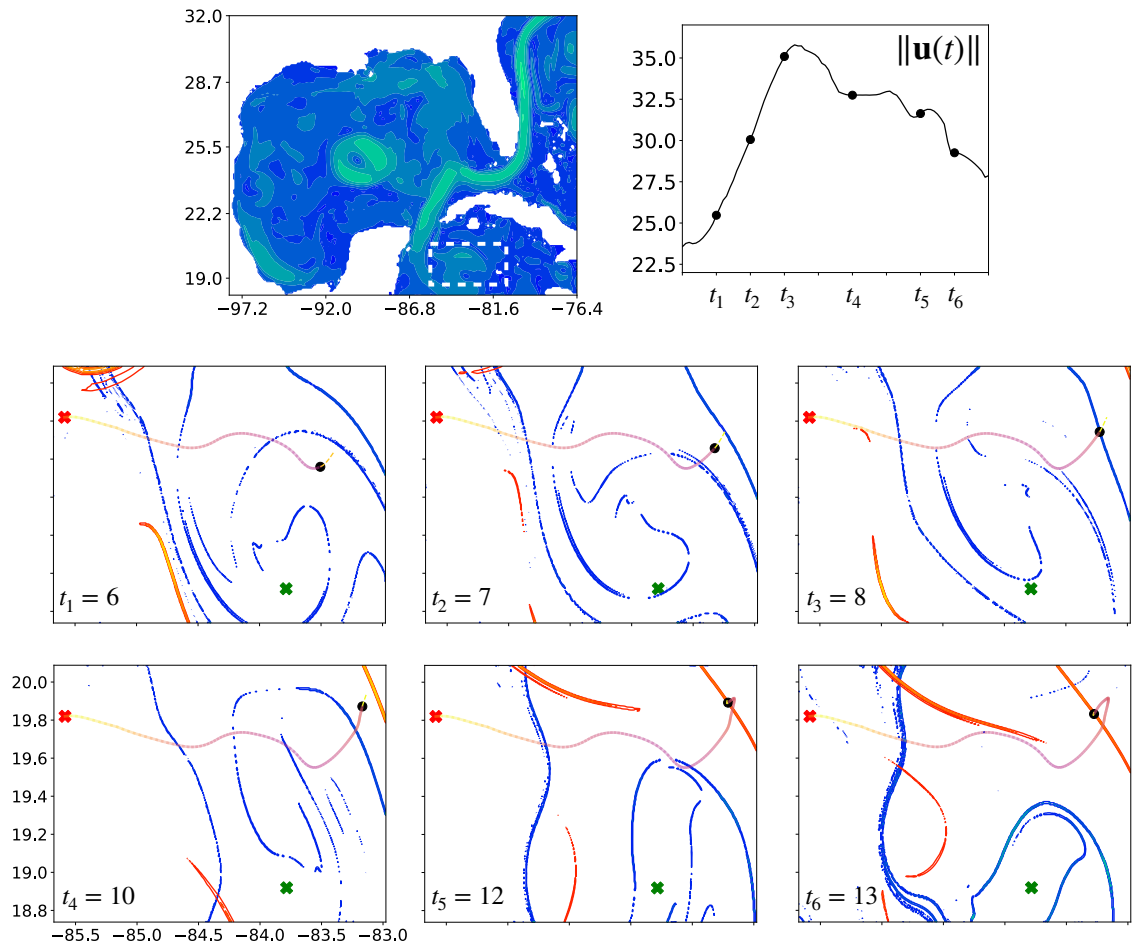


Figure 2.13: This figure shows the use of MPC to plan trajectories in the Gulf of Mexico dataset. We have chosen the region highlighted in dashed dotted lines in the top left plot. The bottom six plots show the trajectory generated for $R/Q = 1$ (full trajectory can be seen in the rightmost plot in Figure 2.14) and $u_x, u_y \leq 2$ km/hr in color shading from yellow to purple to highlight how energy is spent along the path. The units on the x and y axis are longitude and latitude respectively. The red cross in the figures show start location, $\mathbf{x}_{\text{start}} = [-85.5, 19.8]$, the green cross shows the goal location, $\mathbf{x}_{\text{goal}} = [-83.7, 18.9]$. The black dot shows the instantaneous sensor location. Viewing the six figures in sequence we observe that the sensor moves across the blue repelling LCS from $t = 6$ to $t = 8$, where t is the time in days. We observe a spike in the instantaneous energy spent and a slow drop as we move away from the repelling LCS (as seen in the top right plot of $u(t)$). We then observe that the sensor synchronizes with the attracting LCS to move towards the goal location.




Figure 2.14: This figure shows the change in MPC trajectories in the Gulf of Mexico as the R/Q ratio is varied. We see an aperiodic loitering state near the goal location. Given that the time horizon in this case $T_H = 9.6$ hours, which is relatively short compared to the total trajectory time of 100 days, we can see that even with short time horizons the MPC trajectories are capable of making it to the goal location

to generate energy-efficient trajectories for a range of parameters, particularly the prediction horizon and the relative penalty between the state error and control effort. The double gyre oscillation frequency was varied to study its influence on the resulting trajectories. We have constrained the maximum actuation velocity to be less than the largest background flow velocity such that some degree of intelligent planning is required to efficiently traverse the flow.

Through a quantitatively exhaustive study, we have uncovered several interesting trends and established connections between the finite-horizon optimized mobile sensor trajectories and the coherent structures of the underlying flow field. By varying the relative cost of actuation and deviations in the state (i.e., R/Q), the control cost and state error sweep out a Pareto front, and there is often a *sweet spot* where relatively good state tracking performance can be achieved with low actuation costs. These energy-efficient trajectories tend to align with the Lagrangian coherent structures to take advantage of the unsteady background flow. We also found that locations where the sensor spends most energy correlated with the presence of repelling hyperbolic LCS. These findings could be better understood by recalling that repelling LCS are defined as material barriers in the flow field. This explains the expenditure of more energy close to the repelling LCS, as energy needs to be spent to overcome barrier for movement. Furthermore, as the energy spent for movement

goes to zero, the sensor behaves similar to a passive tracer. It is known that the short term dynamics of passive tracers are governed by movement with the FTLE. This suggests that these connections hold for energy optimality. Importantly, we find that it is often possible to generate effective, energy-efficient mobile sensor trajectories with a relatively short prediction horizon, which is promising for the future design of trajectories with limited or partial knowledge of the background flow field.

We observe a rough trend of lower state error when control is less expensive, which agrees with the intuition that the agent is able to more directly pursue the goal state by actuating more aggressively. However, this trend is not monotonic, as there are several cases where slightly decreasing the control cost results in worse state tracking performance. These non-monotonic changes in the cost versus R/Q correspond to *bifurcations* in the agent trajectories, which either correspond to longer trajectories, or to discontinuous jumps in the shape and size of the periodic orbit around the goal state. These bifurcations are more common for smaller prediction horizons, which is also consistent with the intuition that smaller prediction horizons may lead the agent to get trapped by unfavorable flow structures. Similarly, for a fixed relative control cost, there are bifurcations in the optimized trajectory with variations in the time horizon. These bifurcations are relevant in the context of generating mobile sensor trajectories using model predictive control, as small changes in the weights can lead the drastically different trajectories. Upon closer inspection, these bifurcations correspond to the agent trajectory passing through a Lagrangian coherent structure, after which the two trajectory behaviors diverge.

It is also important to note that the energy-efficient trajectories typically result in periodic orbits around the target position, since the unsteady double gyre is periodically oscillating. Previous studies in trajectory generation have mainly focused on solving boundary value optimizations for trajectories keeping the start and end points fixed. Our results show that these assumptions can be relaxed, and moreover, it is possible to reach periodic steady states with little actuation even when the uncontrolled drifter dynamics are chaotic. These periodic orbits correspond to, often desirable, station keeping or hovering behavior.

This work has several implications for the control of individual mobile robots and swarms of robots in geophysical flows. The ability to generate energy-efficient trajectories

that take advantage of the background flow with a short prediction horizon is promising for practical applications. The ability to maintain close periodic orbits around the goal state may also enable efficient long-time monitoring. For example, fix-wing unmanned aerial vehicles must often loiter over an area for sensing and monitoring. Variations in the shape of periodic orbits and the Pareto optimal curves over different R/Q ratios with the gyre oscillation frequency have implications for ocean applications, which exhibit a wide range of spatiotemporal scales with varying oscillating frequencies. We also observed an increase in the expenditure of the sensor's actuation energy as it approached background LCS. This result is beneficial in the context of identifying background coherent structures by observing the energy expenditure patterns of controlled agents. This is an important problem with ongoing work [104, 115, 116]. These results are also potentially useful in the design of scalable navigation algorithms for mobile sensor swarms where the objective is to maintain cohesion or connectivity between agents.

This work motivates a number of interesting future directions. Our results indicate that it is possible to design nearly optimal, energy-efficient trajectories, even with short prediction horizons for the model predictive control; however, it was assumed that the background flow was known perfectly for this short horizon. It will be important to further explore the robustness of these trajectory optimizations to more realistic scenarios with partial, noisy, and uncertain information about the background flow. This analysis may benefit from recent works that have investigated the sensitivity of FTLE calculations to uncertain flow field data [9, 23] as well as how FTLE can be used to propagate uncertainties through chaotic flow maps [102]. Because the optimization result depends strongly on how the MPC trajectories interact with LCS of the background flow field, it may also be possible to incorporate knowledge about the LCS more directly to the optimization. Even with uncertain or partially-observed flow field information, often the LCS are quite persistent, and it may be possible to develop time-varying *maps* of the coherent structures in different geographical regions, for example off the Horn of Africa or in the Gulf of Mexico. In addition, it will be interesting to explore the use of other coherent structure and modal decomposition identification techniques [150, 174]. Further study is also required to characterize the dynamics and coherent structures of the *controlled* vector field of the agent

given a specific control policy. In addition, all the results in this chapter were developed through the study of the double gyre flow field. It will be interesting to perform similar investigations for a variety of flow fields. For example, it will be important to explore how these results change when the flow exhibits a wider range of multiscale behavior in space and time. Particularly, multiscale turbulence will affect the prediction horizon, as uncertainties will be magnified making it challenging to forecast flow structures. These multiscale structures will also impact energy efficiency, both through the forecast uncertainty, but also through making optimal paths more circuitous. Finally, extending the analysis to additional *three-dimensional* turbulent flows will also be critical.

Code Availability

The code used in this study is openly available at <https://github.com/karkris41295/single-agent-MPC-FTLE>.

Chapter 3

FINITE-TIME LYAPUNOV EXPONENT ANALYSIS OF CONTROLLED AGENTS IN FLOW FIELDS

3.1 Introduction

A control policy, for example from MPC or RL, maps each agent state in the domain of fluid flow to an action, which can ultimately be visualized as a vector field. When agents follow a policy within an unsteady flow field, their behavior can be understood as passive drifters operating within an entirely new active flow field – one that is a combination of the original background flow field and the control policy. Given the applicability of FTLE for understanding *passive*, uncontrolled transport in unsteady fluid flow fields, it is natural to extend it to understand *active* transport given an agent’s control policy, which is the focus of this chapter. Policies derived from optimal control methods, and by extension the resulting active flow fields, can often be spatially discontinuous or non-smooth depending on the hyperparameters used. Therefore, our work also contributes to understanding the use of FTLE on non-smooth systems.

In this chapter (drawing results from [85]), we compute FTLE on active agents navigating within an unsteady fluid flow field. Specifically, we examine the controlled trajectories using either a finite-horizon model predictive control optimization or a reinforcement learning strategy to learn a policy that drives an agent towards a goal state. We then perform control FTLE (cFTLE) analysis on the combined flow field generated from the background flow and the control policy, illustrated in Fig. 2.1. Similar studies have been performed in the past for biological applications, where FTLE was used to understand the behavior of plankton actively propelling away from feeding jellyfish [126]. We confirm that, much like the traditional FTLE, cFTLE can be used to visualize transport barriers, which allow us to understand the transport of agents following a control policy. These results are useful in the context of finding ideal deployment locations for rendezvous and




Figure 3.1: This figure outlines the methodology of cFTLE which takes Lagrangian controlled trajectories as input and outputs LCS curves which highlight interesting transport features of the control law which are generally not visualizable by simply plotting all the trajectories or by plotting the control vector field. The left most figure shows that in order to evaluate our control algorithm, we generate the simulation from a mesh grid of initial agent positions at $t = 0$. We end the simulations at $t = T_A$ (as shows in the middle figure). Finally we use the FTLE algorithm on this data to generate “cFTLE” ridges which highlight important features in the domain.

coordinated multi-agent goal tracking. We further relate the features of cFTLE, which uncover exponential divergence of initial conditions, to the sensitivity of optimization cost functions used in trajectory planning. We also show that these findings are not exclusive to policies generated through MPC or RL, but to policies generated by any optimal control method with a particular structure in cost function, which we elaborate on. In a previous work by the authors [87], it was found that large energy expenditure of active agents corresponded with the presence of large background FTLE. Here, we further explore this connection and find that the energy spent has a closer correlation with cFTLE, which are deformations of the passive FTLE ridges. Given the connections to optimal control, we perform sweep through a range of parameters that govern the aggressiveness of the control, the set point to track, and the amount of future information about the background flow field available, which creates spatial discontinuities. In the future, this study could potentially aid in methods for faster learning of control policies with partial knowledge of the control FTLE and vice versa, and aid in summarizing the effectiveness of policies used for flow navigation.

3.2 Methodology for Control FTLE (cFTLE)

In our work, we propose the use of FTLE in the analysis of policies generated through methods such as MPC and reinforcement learning. However, generating policies in this manner presents obstacles in the computation of cFTLE ridges, which will be elaborated in the following subsections. In this section, we outline how the computation of cFTLE can be derived from the traditional direct application of FTLE on unsteady flow fields.

3.2.1 Control Finite Time Lyapunov Exponents (cFTLE)

Extending the FTLE from Chapter 1 to incorporate control, the equations governing the dynamics of an actuated agent from Eq. (1.9) are used. In this work, we shall consider the class of kinematic models [73] where

$$\mathbf{g}(\mathbf{x}(t), \mathbf{u}, t) = \mathbf{v}(\mathbf{x}(t), t) + \mathbf{u}. \quad (3.1)$$

These models assume that the sensor can generate its own flow relative velocity $\mathbf{u}(t) = [u_x, u_y] \in \mathbb{R}^2$ in addition to the flow-induced velocity. Alternatively, models that incorporate inertial effects can also be used [126, 171, 193]. The function $\mathbf{v}(\mathbf{x}(t), t)$ is the unsteady flow field within which the agent is moving. For the majority of this study, we consider this function to be given by the non-autonomous double-gyre equations [74, 165, 166]. Because the control \mathbf{u} is typically a full-state feedback control law, $\mathbf{u}(\mathbf{x}, t)$, the resulting flow field \mathbf{g} may be considered to be a function of \mathbf{x} for a given policy \mathbf{u} .

Similar to traditional FTLE, we initialize a grid of agents at time t_0 and numerically integrate through $\mathbf{g}(\mathbf{x}(t), \mathbf{u}, t)$ for a fixed amount of time (i.e., the advection time) $T_A \in \mathbb{R}$, resulting in a flow map $\hat{\Phi}_{t_0}^{t_0+T_A} : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$:

$$\hat{\Phi}_{t_0}^{t_0+T_A} : \mathbf{x}(t_0) \mapsto \mathbf{x}(t_0) + \int_{t_0}^{t_0+T_A} \mathbf{g}(\mathbf{x}(\tau), \mathbf{u}, \tau) d\tau. \quad (3.2)$$

The flow map operator is particularly important in our work as it will later be used to understand the connections between cFTLE and value functions in optimal control. The

operator $\hat{\Phi}_{t_0}^{t_0+T_A}$ maps an initial agent position $\mathbf{x}(t_0)$ to a position advected forward by time to $\mathbf{x}(t_0 + T_A)$ by the flow field and agent propulsion.

FTLE based on agents integrated forward in time, $T_A > 0$, result in coherent structures that repel agents, which can be seen in the right most plot in Figure 3.1 where the blue curves are visible, and later figures in this work. Compared to the backward time FTLE, the forward time cFTLE is more straightforward, assuming that the policy for all states is known beforehand. To calculate the backward time cFTLE, we need both flow field and policy data backwards in time, which is not possible in realistic scenarios. FTLE based on agents integrated backward in time, $T_A < 0$, results in coherent structures which attract agents. These can be computed for periodic flow fields as both policy and flow field can be extrapolated backwards due to periodicity (as demonstrated in the Appendix).

3.2.2 Computation Using Interpolation

For the computation of cFTLE, Eqs. (1.9) and (3.2) often require the use of a dense grid of initial conditions. Generally, with MPC, this would incur a large computational cost as a full *closed-loop* MPC optimization would need to be computed over the prescribed time horizon, T_H , for the advection time, T_A , for each spatial initial condition at each simulation time step. To reduce the total number of optimizations computed and reduce the requirement for several closed-loop MPC simulations, we pre-compute an open-loop policy for the entire space and time domain. We consider agent initial conditions fixed to a grid (as in the left plot in Figure 3.1) and compute the control actions at the grid node locations (as seen later in the plots of Figure 3.9). To keep the computational cost down, we make use of a coarse grid and then use linear interpolation to compute the policy for off-grid initial conditions. This allows us to build a lookup-table function $\hat{\mathbf{u}}(\mathbf{x}, t)$ for every point in space and time. We then integrate a dense grid of passive particles through the combined vector field of fluid flow and estimated *open-loop* control through the following ODE,

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{v}(\mathbf{x}(t), t) + \hat{\mathbf{u}}, \quad (3.3)$$

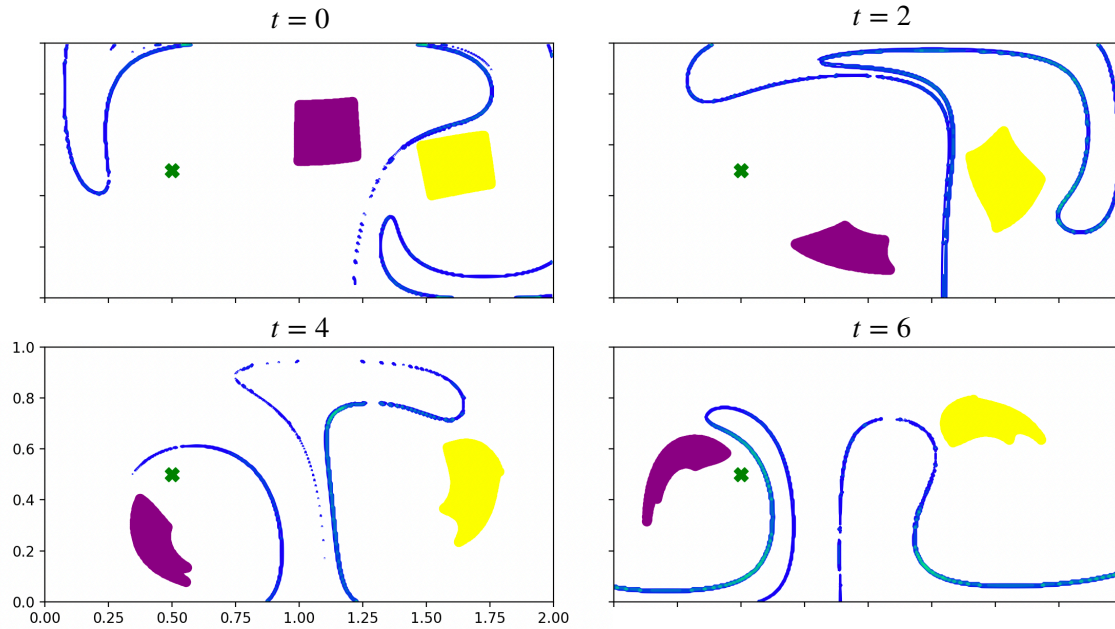


Figure 3.2: Control FTLE ridges separate dynamically different regions of agents moving in the flow field. We show two distinct patches of agents - a purple and a yellow patch. Both patches start initially on different sides of the cFTLE ridge. The time evolution shows that cFTLE ridges separate the patches that undergo different paths and accrue largely different short-term state error cost despite starting close to each other.

where \hat{u} is the estimated open loop control law using linear interpolation. We then follow the computation method outlined in Section 3.2.1 to obtain the Lagrangian trajectories.

3.3 Results

In this section, we first verify that cFTLE behaves much like regular FTLE and demonstrate how this property can be used diagnostically for the analysis of control laws. Next, we highlight a connection between value functions from optimal control theory and cFTLE ridges. Finally, given the aforementioned connection, we study the deformation of cFTLE ridges under the change in parameters of MPC. We show that these changes in cFTLE ridges can be used to understand the change in value functions useful for optimal control. We also show that cFTLE ridges can be a powerful tool for numerically identifying switching surfaces and approximations to time-varying invariant manifolds. This

could potentially aid in using various techniques from non-smooth analysis, and knowledge about global bifurcations to analyze and predict the oncoming changes in system stability as hyperparameters are varied. Note: The discretized time step across all double gyre simulations is $\Delta t = 0.1$. Each simulation was run for 800 steps. In the following sections, we will vary the relative cost of actuation versus state error, given by the ratio R/Q , and analyze how this impacts the mobile sensor trajectories.

3.3.1 Interpretations of cFTLE

We first discuss the interpretations of cFTLE ridges in comparison with the traditional view of FTLE when applied to the transport of agents through a flow field. We will vary the relative cost of actuation versus state error, given by the ratio R/Q , the time horizon T_H , position of the goal and analyze how this impacts the cFTLE. In the limit of an infinite cost of control, the cFTLE should converge to the traditional passive FTLE.

Transport Barriers

Forward time cFTLE ridges form transport barriers similar to traditional FTLE ridges [189]. This can be observed from Figure 3.2, where we have two patches of agents moving towards a common goal. We observe that in the time of advection, the purple patch and yellow patch are separated by the cFTLE ridge as a dividing barrier. Agents that fall on a cFTLE ridge are further stretched apart at later times, which will be especially highlighted later in Section 3.4. Such analysis is useful in the context of analyzing control laws for large-scale transport of under-actuated mobile sensors in fluid flows. For instance, determining the deployment locations for intelligent mobile sensors is a challenging task with significant cost and schedule complications [178]. cFTLE can be potentially useful in finding locations where the control law is able to transport agents effectively.

Rendezvous Through Lobe Transport

Building upon the idea of transport barriers from the previous section, cFTLE barriers can surround agents with barriers that persistently keep the agents trapped. This behavior has

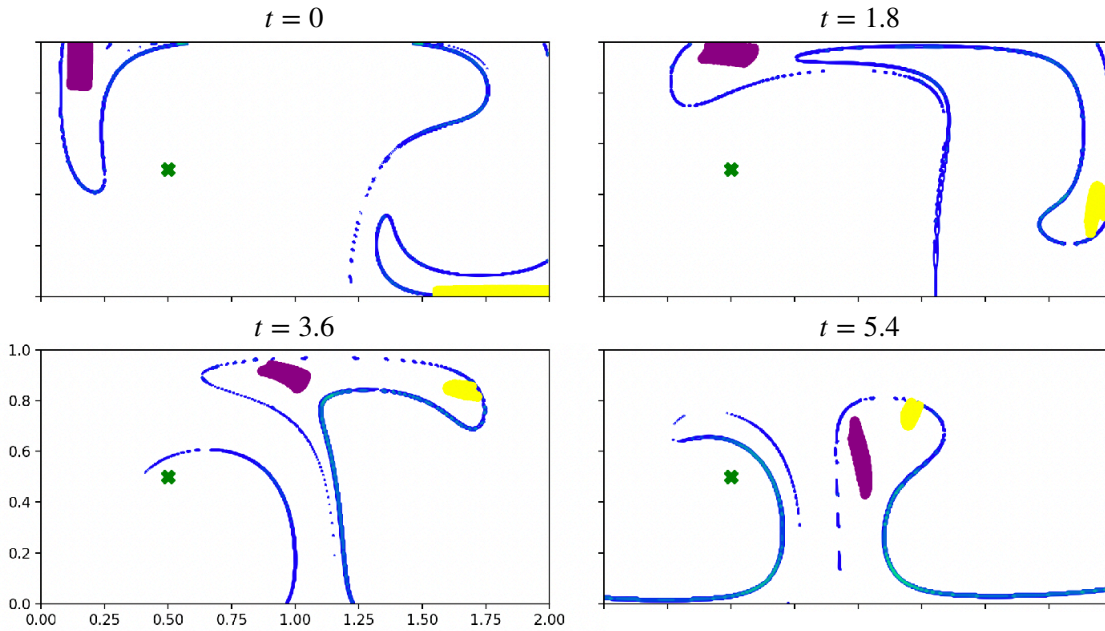


Figure 3.3: cFTLE boundaries enclose regions in the domain which accrue similar cost. Despite the fact that the purple and the yellow patch start far from each other, they meet in short time since they are in the same lobe. This can be viewed particularly at $t = 3.6$ on, where the yellow and purple patch are within the same bounded cFTLE region. These structures are not apparent when viewing the control law or the flow field alone. An application of visualizing such regions in the domain may aid in solving rendezvous problems in the ocean where two distinct groups of agents may need to meet in finite time

been exhaustively studied in dynamical systems theory [145, 188]. Attracting and repelling invariant manifolds intersect and form lobes that trap passive particles. Traditionally, attracting and repelling FTLE ridges have been used to visualize such trapping regions. We observe this in the context of cFTLE in Figure 3.3. Despite not plotting the attracting cFTLE ridges, we can see that the repelling cFTLE barriers, inside of which the purple and yellow patches are placed, form a single trapping structure that brings the two patches to meet in forward time, while they move towards the goal. This idea is potentially useful in the context of multi-agent planning, where drop locations can be found that allow agents to rendezvous in forward time. Such rendezvous problems are important in the context of path planning [163].




Figure 3.4: The plot (a) shows the control FTLE ridges for double gyre setup at an R/Q ratio of 80 on the left, with a time of advection $T_A = 15, T_H = 3$. The agents are tracking the goal set point at $\mathbf{x}_{goal} = [0.5, 0.5]$. On the right (b), we show the accumulated state error J_F for a meshgrid of agents in the domain for the same R/Q and T_H . We see that the cFTLE ridges are located where there is a drastic change in color on the right figure. This indicates that the cFTLE ridges at a particular T_A partition the domain by the cost spent (i.e., large change in J_F).




Figure 3.5: On the left, we plot the instantaneous energy spent by an agent starting from a meshgrid of initial conditions moving to a goal location on the left ($\mathbf{x}_{goal} = [0.5, 0.5]$) at $t = 0$. On the right, we plot the accumulated state error over the time horizon of MPC ($T_H = 4.5, R/Q = 15$). In both plots, the streams are the control law generated by all the MPC agents, and the cFTLE ridges are shown where the time of advection is equal to the time horizon of the MPC, $T_A = T_H$. We observe here that the cFTLE ridge overlaps with regions where actuation energy expenditure peaks locally. Also, similar to Figure 3.4, the cFTLE divides the domain where the cost gradient is the largest.

3.3.2 Connections to Optimal Control

In this section, we discuss cFTLE from the perspective of optimal control. Inspired by the first term in Eq. (1.10), Figure 3.4 (right) displays the value function landscape of

$$J_F = \|\mathbf{x}(T_A) - \mathbf{x}_{goal}\|_2^2, \quad (3.4)$$

which, intuitively, is the square of the agent distance from the goal after a time T_A has elapsed, as a function of where the agent initially started. By definition, J_F only depends on the final state error at time T_A , in contrast to J_{LQR} , which depends on the state error and control energy spent integrated over the trajectory. The general idea of this section is that we will use J_F as a simple model to gain insight into J_{LQR} .

When J_F is viewed in comparison to the cFTLE on the left in Figure 3.4, we observe that the cFTLE ridges overlay with regions where there is a sharpest change in value of J_F . This indicates that cFTLE ridges coincide with regions of high sensitivity in the state error of a cost function, which will be further verified later in this section. This is in agreement with the definition of the FTLE, which is a measure of sensitivity of final conditions to infinitesimally close initial conditions.

Furthermore, the left plot in Figure 3.5 displays the energy spent by the control action taken, $\|\mathbf{u}^*\|$, as a function of initial condition. We see that cFTLE ridges fall on regions where the energy spent by the agent is the largest. We know that Eq. (1.11) implies that $\|\mathbf{u}^*\|_2^2 = \|\mathbf{R}^{-1}\nabla J_{LQR}\|_2^2$, and, we established in the previous paragraph that cFTLE ridges are connected to sharpest *change* in J_F . In image processing, the sharpest change of an image or function value (e.g., J_F) can be found by taking the magnitude of the gradient of the function ($\|\nabla J_F\|_2^2$). These results specifically highlight a connection between cFTLE ridges, sensitivity in value functions $\|\nabla J\|_2^2$, and the instantaneous energy spent by optimal control policies $\|\mathbf{u}^*\|_2^2$.

It is also possible to further explore this connection mathematically. When using the fact that $\mathbf{x}(T_A) = \hat{\Phi}_{t_0}^{t_0+T_A}(\mathbf{x}(t_0))$, the gradient of J_F is given by

$$\nabla J_F = 2[\mathbf{x}(T_A) - \mathbf{x}_{goal}]^T \mathbf{D}\hat{\Phi}_{t_0}^{t_0+T_A}. \quad (3.5)$$

By comparing to the FTLE definition, we can observe that the flow map operator used in the computation of FTLE is deeply connected to the gradient of cost functions with quadratic state error. Specifically, the maximum cost function gradient is related to the maximum singular value of the flow map Jacobian, which is related to the FTLE.

At each instant over the horizon, we compute a complete finite-horizon optimization

and take only the first control action, a standard practice of MPC. On the right plot in Figure 3.5, we plot the integrated state error cost of this finite-horizon optimization at the first step - the first term of J_{LQR} . If we were to make $T_A = T_H$, the J_F plot in Figure 3.4 would look similar to the right plot in Figure 3.5, further highlighting the connection between cFTLE and sensitivity in J_{LQR} . We observed across our simulations that the cost function from J_F and J_{LQR} had very similar qualitative features, and one could potentially use whichever is easier to compute to gain insight into the other.

In model-free control approaches, such as reinforcement learning, one is interested in generating fast *approximations* to the value function as the agent is moving through the flow field. These results show that cFTLE ridges can have a potential impact on estimating these value functions when such model free approaches are used. We will further explore reinforcement learning in Section 3.4.

3.3.3 Deformation of cFTLE with change in MPC parameters

We have now shown that cFTLE ridges contain information about the sensitivity of a control policy, the regions of large magnitude of control, and the boundaries between finite-time invariant sets. In this section, we study how these ridges depend on the aggressiveness of the control law and the location of the goal. Particularly, these two parameters direct the strength and magnitude of fluxes generated by the model predictive control policy over the unsteady flow field. We also study change in cFTLE as the time horizon T_H is varied. This parameter changes the spatial complexity of the control policy. These results can be potentially useful for anticipating changes in the cFTLE for a different parameter value. In the next section (Section 3.4), we finally use the ideas learned in this section to interpret a policy generated through another popular optimization-based control approach – reinforcement learning.

Varying the Cost of Control, R/Q

Figure 3.6 shows cFTLE ridges plotted for incremental changes in R/Q ratio. We find that as the R/Q ratio increases, the aggressiveness of the MPC strategy decreases, and the




Figure 3.6: This figure shows the cFTLE ridges for four different R/Q ratios of 20, 40, 60, 80, and how the barriers shrink as the aggressiveness of the MPC increases. Here, the time of advection for cFTLE $T_A = 15$, and the time horizon of MPC is $T_H = 3$. We observe that as the aggressiveness of the control decreases, the cFTLE approaches the passive FTLE.

average magnitude spent for control decreases, while the direction of control for each state does not change appreciably. Extremely aggressive control (small R/Q) pushes the cFTLE ridges further towards the right when the goal is in the left gyre, and shrinks the size of the left cFTLE ridge completely. For less aggressive control (large R/Q), the cFTLE ridges lengthen and begin to approach passive FTLE ridges. These results are intuitive, given that when the control becomes less aggressive, the agent behaves more like a passive particle due to lack of control authority.

Changing the Goal Location, \mathbf{x}_{goal}

The results of changing the goal location can be found in Figure 3.7. When a control law or policy acts over an unsteady flow field, it generates an added flux to transport particles towards the goal. This is visualized by the grey arrows in Figure 3.7. Here, we observe how the cFTLE changes depending on this newly added flux. For instance, creating a sink in the right gyre moves the cFTLE ridge to the left from the initial position of the corresponding




Figure 3.7: This figure highlights the change in cFTLE ridges as a function of the goal location. Here, $R/Q = 15, T_A = T_H = 4.5$. The streamplot on both plots show the control law. The color in the background shows the regions where most energy is spent (brighter yellow shows larger energy spent and darker shows less). On the left plot, the goal location is set to the right gyre, and on the right plot, the goal location is set to the left gyre. We observe how the fluxes generated by the control law moves the passive FTLE ridge. A left flux moves the cFTLE ridge to the right, and vice versa. This figure continues to highlight the connection between the terms in the cost function and cFTLE ridge as in Figure 3.4.

passive FTLE ridge. Creating a sink to the left moves the cFTLE ridge to right. This shows us that cFTLE ridges move in the direction opposite to the flux of flow generated.

Varying the Time Horizon, T_H

Figure 3.8 shows the change in cFTLE field as the time horizon parameter T_H is varied. Note that we plot the contours of the FTLE field as opposed to the FTLE ridges since agents following the policy cause the ridges to be less sharp, and therefore, more challenging to extract and visualize through thresholding. Unlike the deformation of the cFTLE in the previous section, looking at the deformation of cFTLE in Figure 3.8, we observe that cFTLE undergoes a deflation-type effect and the cFTLE field becomes non-smooth with several branches. In Figure 3.9, the first two figures show the policy changing as a function of T_H . The time horizon parameter captures the amount of future knowledge of the flow field being incorporated in taking the present action. Therefore, the policy changes from a naive sink-like behavior to become spatially more complex as the time horizon is increased. These spatially complex policies make better use of actuation to escape regions or lobes dictated by the *passive* FTLE, where the agent will be advected away from the goal and

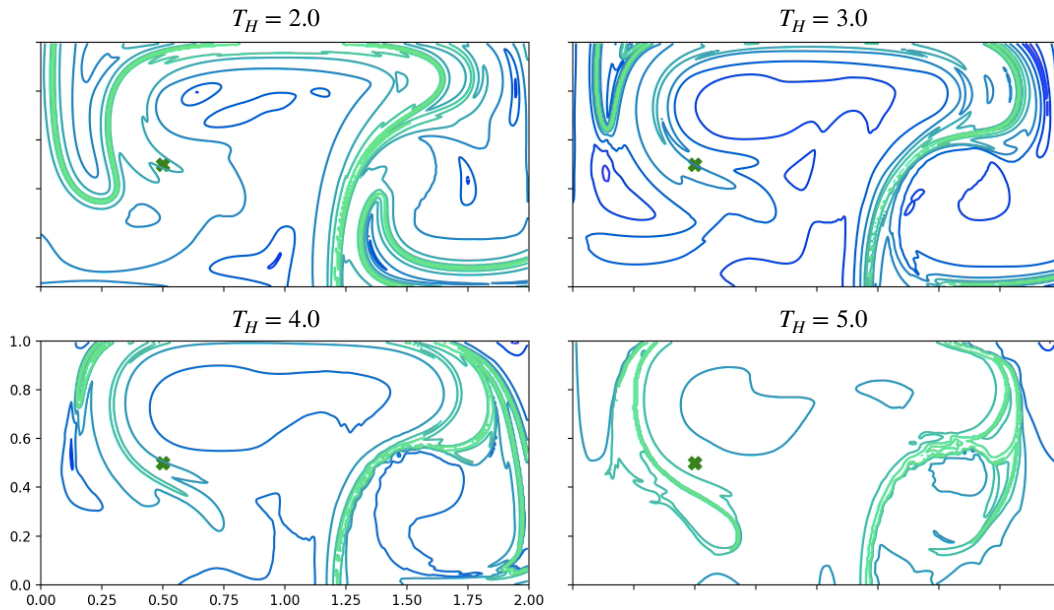


Figure 3.8: This figure shows the change in cFTLE field for time horizon values $T_H = \{2, 3, 4, 5\}$, using an advection time of $T_A = 15$ and $R/Q = 50$ at $t_0 = 0$. Unlike in Figure 3.6, the cFTLE ridge changes due to more intelligent use of control and not because of sheer use of greater effort. The cFTLE ridges exhibits more structures, and the curves have more branches due to several ridges coming close to each other and collapsing on each other.

increase the cost function value.

We also observe in Figures 3.5 and 3.7 that the cFTLE highlights regions in space where the policy becomes spatially discontinuous. The streamlines on either side of the cFTLE ridge are in different directions. This is intuitive, as regions spatial discontinuity in a vector field can potentially also generate large Lyapunov exponent values, in addition to effects from shear and normal hyperbolicity. When using optimal control methods, the policy can often make the controlled system a non-smooth dynamical system [149]. These results shows that the cFTLE can be used for switching manifold detection in controlled systems.

3.4 Example from Reinforcement Learning

Thus far, we have discussed the use of cFTLE exclusively on policies generated from model predictive control. However, cFTLE can be used to analyze policies generated from




Figure 3.9: This figure shows the range of policies that can be generated at time $t = 0$. The policies considered in this chapter vary with time t . The left and middle policies are generated using MPC with a time horizon of 1.0 and 3.0, respectively. The policy on the right is generated by using DDPG and can be considered the “infinite-horizon” case. All policies have an R/Q ratio of 70. From left to right, there is an increase in spatial complexity of the policy as it begins to use more information of the flow field. This spatial complexity in turn influences the cFTLE ridges computed, which can be seen in Figure 3.8 and Figure 3.10.

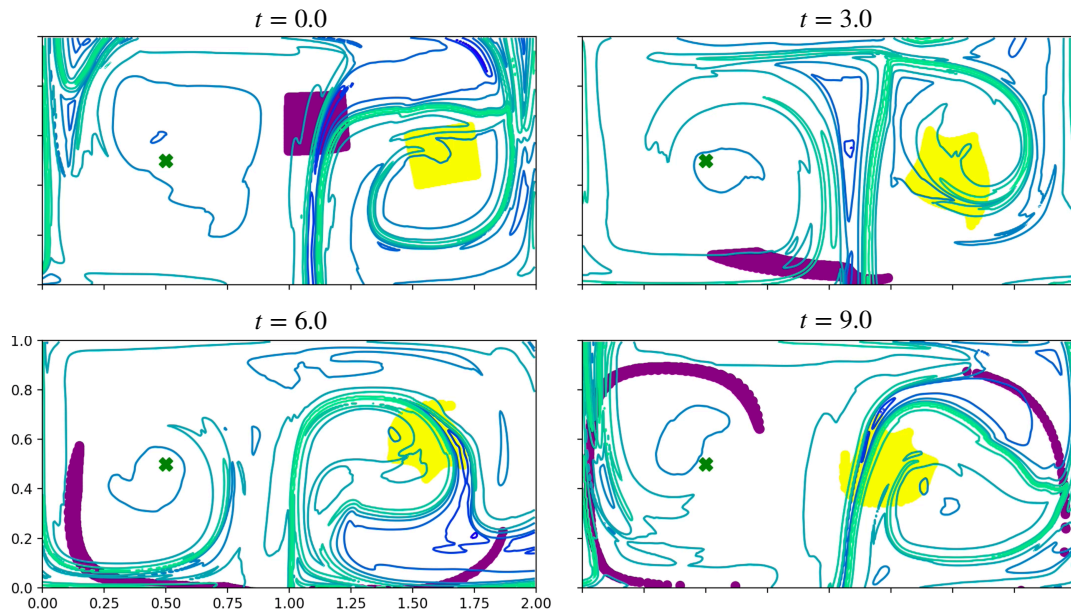


Figure 3.10: This example shows the repelling cFTLE field plotted for a policy generated using DDPG with $R/Q = 70$. Much like Figure 3.2, we use two patches to highlight transport mechanisms. We also use the same initial patch locations for comparison.

any control or planning methods. Reinforcement learning (RL) is a powerful alternative paradigm to model predictive control for policy generation [25, 134]. A key difference be-

tween the two paradigms is that RL is typically model free, i.e., policies can be generated without any knowledge of the flow field, but rather from data alone. In this study, we use the deep deterministic policy gradient (DDPG) method [97] available in the stable baselines package [130] to demonstrate the use of cFTLE analysis on reinforcement learning policies.

Much like model predictive control, RL requires the definition of a reward function. Typically, the cumulative reward over the simulation is computed. For this, we make use of the negative of the discretized LQR cost function, defined in Eq. (1.10) due to the fact that the convention is to *maximize* the reward, as opposed to minimizing the cost in methods such as MPC. We make use of an Euler time-stepping scheme to propagate an agent's state forward in time with a time step of 0.1. Each complete RL simulation (or episode) is run for 800 time steps, which totals to 80 units as in the MPC case.

At a high level view, RL runs multiple episodes with random policies and random initial states to generate data of cumulative rewards over each episode. This data is then used to iterate towards a policy that maximizes the reward over each episode, also known as the cumulative reward. The setup of DDPG used in this chapter does not explicitly incorporate a time horizon; however, the resulting policy can be interpreted as a solution of an infinite horizon problem. The right plot in Figure 3.9 shows an RL policy generated over the double gyre flow field for the same objective outlined in the problem setup section. Metrics such as the average cumulative reward are often used to judge the performance of the policy [173]. These metrics do not provide a spatial description or boundaries in space where the reward declines, which is possible through cFTLE analysis. RL generates a lookup table policy function, described by a neural network $u(\mathbf{x}, t)$ for each state in space and time. Therefore, unlike in the MPC case, computing the FTLE of an RL policy does not require the use of interpolation as mentioned earlier. As in the previous section, we plot the cFTLE *field* as opposed to cFTLE ridges to better visualize these invariant structures.

Figure 3.10 shows the integration of two patches of different colors with the same initial configuration of Figure 3.2. We observe in contrast to Figure 3.2, that, since the purple patch starts on the cFTLE ridge as opposed to on one side of it, it undergoes large deformation normally along the cFTLE ridge in the middle of the domain, while the yellow patch

maintains cohesion for longer time. However, at time $t = 9.0$, we see that a part of the yellow patch lies on the the other side of the cFTLE ridge, therefore, we can predict that in the future, the yellow patch will also be stretched out. Since the RL policy can be viewed as an infinite-horizon extension of the MPC case, we can interpret the several streaks of large cFTLE value as passive FTLE structures that have collapsed very close to each other. This can be visualized at $t = 3.0$ where the FTLE ridge takes a “P” shape that appears as a branch off the middle hyperbolic LCS, which occurs due to the shrinking and collapsing of a passive FTLE lobe as seen from the large time horizon case. This is similar to the case of time horizon 4.0 and 5.0 in Figure 3.8.

Value functions for RL methods can find interesting pathways through passive dynamics that traditional methods miss. This is due to the highly non-convex nature of RL algorithms, which makes it challenging to bound and analyze the policies generated through these methods. cFTLE provides a principled method of analyzing these policies. Moreover, given that value functions are connected to the cFTLE as demonstrated earlier, we can potentially use cFTLE to minimize the steps required to compute policies, leading to faster convergence.

3.5 Discussion

We studied the use of the FTLE method on active agents that use actuation to move towards a fixed goal in an unsteady flow field. In particular, we focused on the FTLE analysis of agents using model predictive control and reinforcement learning independently to generate policies in a double gyre flow field. Broadly speaking, this work can be interpreted through the lens of its potential applications to the following problems:

1) *Invariant sets*: We first compute the cFTLE for agents using policies generated through MPC. We confirm that, much like passive FTLE, cFTLE can be used to find barriers that separate different regions of agents following a particular policy. This suggests that cFTLE can be useful to identify finite-time invariant sets and invariant manifolds for active agents, which is particularly useful in the context of navigating unsteady aperiodic flow fields as coherent structures only persist for finite time. Invariant sets are useful in optimal

control from the perspective of analyzing robustness and designing model predictive control policies [131]. The areas with smaller cFTLE values correspond to regions where the control is robust to small changes in the parameters of the controller.

2) *Multi-agent path planning*: Next, we find that cFTLE can be used to identify trapping regions or regions closed off by barriers or lobes. When viewing these trapping regions at large scales in an unsteady flow field, we can find regions in which, when following the optimal policy, agents can move towards a goal cohesively without getting separated at large distances. When viewing these barriers at small scales, they can be used to find deployment locations where the agents would not collide with physical objects in the flow field or other agents following the same policy.

The latter part of the chapter focuses on the deformation of cFTLE ridges, or rather, the intuition behind how the cFTLE ridges move when new fluxes from the control policy are added to the passive unsteady flow field. We find that the cFTLE ridge moves *opposite* to the direction of the added flux from the flow field. We also find that control can cause a trapping barriers to shrink and collapse into multiple branches emerging from a larger ridge. This can highlight regions (or lobes) in space that are integral in escaping or entering to move towards the goal effectively. An interesting observation is that cFTLE ridges can highlight the presence of switching manifolds in controlled systems, where the optimal policy renders the controlled dynamical system non-smooth. This is possible since regions of discontinuity in a vector field can generate large Lyapunov exponents. Finally, the computation of cFTLE ridges itself is challenging. Control policies generate sources and sinks, where innovations in the computation of compressible FTLE can be used [58]. Potentially, many of these ideas can be tested in realistic experimental systems with three-dimensional flows and multi-scale turbulence.

Code Availability

The code for this work has been made available on GitHub at <https://github.com/karkris41295/cFTLE>.

Chapter 4

CONTROL OF POINT VORTEX DYNAMICS USING KOOPMAN EIGENFUNCTIONS

4.1 Introduction

The control of vortices in fluid flows is of prime importance in many engineering applications [40, 45, 83, 119]. Point vortex models, governed by the Biot-Savart equations, offer a powerful framework to study vortex dynamics of ideal two-dimensional flows [5, 120]. For example, point vortex models have been used to describe various vortex phenomena in wakes [168, 169]. In this chapter, we develop a new set of strategies for controlling point vortices based on a recently developed connection between conserved quantities and Koopman control theory [76, 77]. These results are from my paper [86].

In the past, ideas from optimal control theory, such as the direct application of the Pontryagin maximum principle [121, 183] have been used to find control strategies for simple cases involving a small number of point vortices in the domain [180, 181]. Other methods arising from chaos theory, such as OGY (Ott, Grebogi and Yorke) control [128] have also been used to find control strategies for periodic flows. However, little work has been done in recasting the point vortex equations in terms of conserved quantities for control. This idea is inspired by recent work in Koopman operator theory [27, 114] that shows that invariant quantities are Koopman eigenfunctions corresponding to zero eigenvalue [77], and that it is possible to control these eigenfunctions using a Koopman with control framework [76]. Direct application of model-predictive control on the governing equations *has* been used to control vortex dynamics [148], although, in this chapter, we study the use of model-predictive control on the model formulated in terms of the conserved quantities. This gives a global description of the dynamics written in terms of energy, angular impulse, and linear impulse, as opposed to the local description of the positions of each vortex.

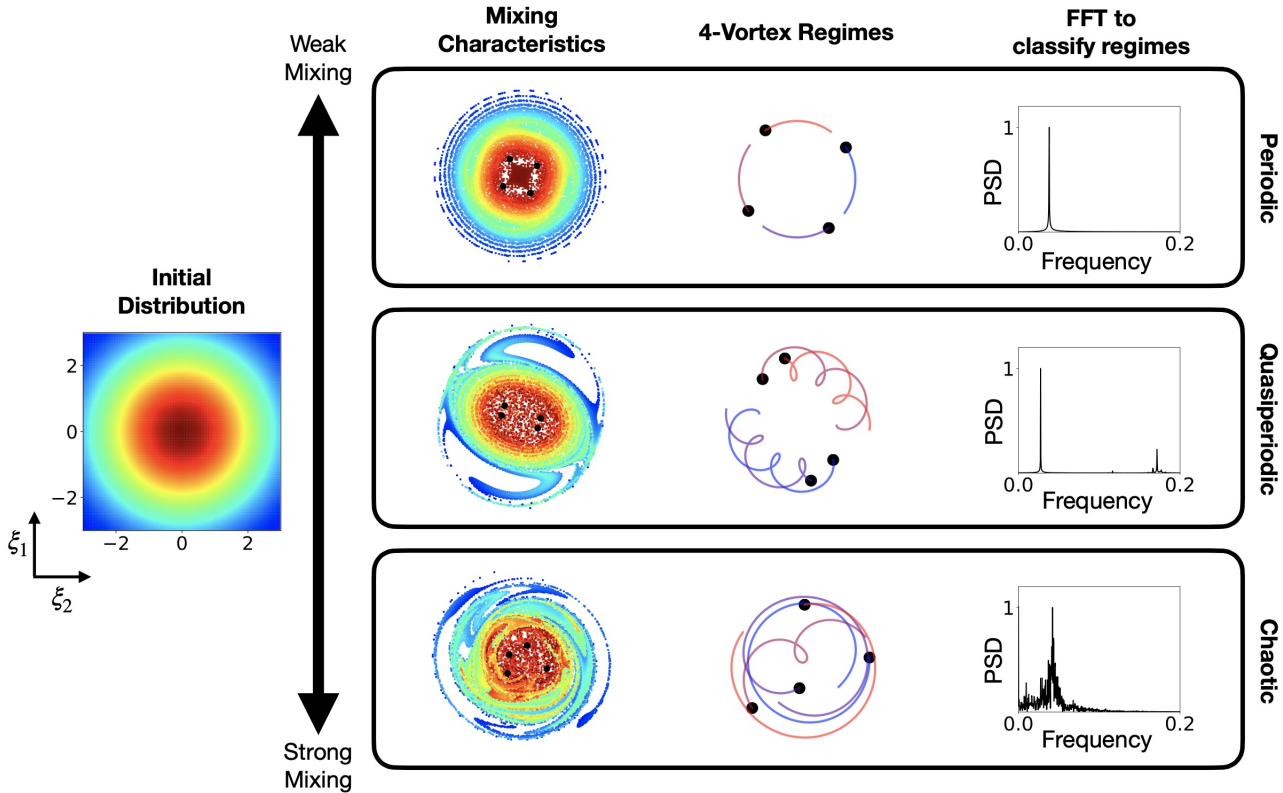


Figure 4.1: **Uncontrolled vortex dynamics:** The key dynamic states of a 4-vortex system are highlighted: periodic (or single-frequency periodic), quasi-periodic (or multi-frequency periodic), and chaotic (broadband frequency). These dynamic states differ by the angle the vortices initially make with the origin [18]. The control demonstrated in this work aims at switching between these different states. From left to right, we show the mixing of tracer particles (starting from the initial distribution to the left), vortex trajectories and the Fourier transform (FFT) applied to the x-component of a single vortex to characterize occurring frequencies. A broad spectrum implies chaos, whereas sharp peaks at discrete frequencies can imply quasiperiodicity or periodicity.

We primarily focus on the four-vortex system. The dynamics of point vortices in a domain are well understood for up to four vortices [120]. It is known that for up to three vortices, the conserved quantities of the system can be used to reduce the dynamics to an integrable form that can exhibit mostly periodic solutions. However, when $n = 4$, the system of vortices loses integrability and behaves like a system of coupled oscillators [6, 21]. When the number of vortices is $n \geq 4$, the vortex system begins to exhibit chaotic behavior.

The 4-vortex system in particular, has been exhaustively studied with respect to mapping out the phase space and the period-doubling bifurcations the system undergoes as a reduced Hamiltonian energy is varied [21]. Four vortex systems have also been studied from the perspective of mixing, where periodic, quasi-periodic, and chaotic vortex configurations have been parameterized by the angle between a single vortex and the origin [18]. These configurations are shown in Figure 4.1. Therefore, the four-vortex system is an ideal test case for studying invariant control due to the range of phenomena exhibited by the configuration and its simplicity.

In this work, we alter the vortex dynamics from a periodic configuration characterized by the four vortices arranged in a square (Thomson solution) as in the top panel of Figure 4.1. Control is applied through a virtual cylinder which is modeled as a stationary point vortex, whose circulation can be changed through model-predictive control (MPC). Importantly, the control is applied in a new coordinate system, defined by the vortex invariants, as this establishes a Koopman invariant subspace in which control becomes more tractable [76, 77]. We first investigate changing the conserved quantities independently, with no constraints on the other conserved quantities, with relatively intuitive results:

- increasing/decreasing the Hamiltonian has the effect of merging/separating vortices;
- increasing the linear momentum produces translation; and,
- changing the angular momentum affects the variance of vorticity in the distribution.

We find that changing a single invariant typically causes other invariant values to change; therefore, we first investigate the change in a single invariant while keeping other invariants fixed through penalization in the MPC cost function. We then show that by using multi-invariant control, we can restrict the transition to states not reachable through single invariant control alone. We also show that conserved quantities may be used to enforce or break symmetry, which enables both quasi-periodic and chaotic states with the same Hamiltonian/kinetic energy. In other words, it is possible to limit the configuration to

reach a quasi-periodic state as opposed to chaotic states using control. From [21], it is also known that Thomson solutions are surrounded by regions of quasi-periodic configurations in phase space characterized by the vortices arranged in a rectangle. When using invariant control, we observe transitions to these nearby states, particularly, when changing the linear impulse. Finally, we highlight the difference between single-invariant and multi-invariant control in a simple mixing example. We compare the KL divergence of these states and how much actuator energy it requires to move to them.

4.2 Methodology

4.2.1 Point vortex dynamics

Consider the governing dynamics of an inviscid, incompressible fluid given by the Euler equation as

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p, \quad (4.1)$$

where $\mathbf{u}(\mathbf{x}, t)$ is the flow velocity, and $p(\mathbf{x}, t)$ is the scalar pressure field at each spatial coordinate \mathbf{x} and time t . With a systematic reduction of the symmetries of the system [106], the vorticity field ($\omega = \nabla \times \mathbf{u}$) from the Euler equations in the two-dimensional plane can be distilled to a discrete point vortex solution of the form

$$\omega(\mathbf{x}) = \sum_{i=1}^n \frac{\kappa_i}{2\pi} \delta(\mathbf{x}) \quad (4.2)$$

where $\delta(\mathbf{x})$ is the Dirac delta function, and κ_i is the strength of point vortex $i = 1, 2, \dots, n$. The position of each point vortex $\mathbf{r}_i = (x_i, y_i)$ evolves according to the Biot-Savart law as

$$\frac{d\mathbf{r}_i}{dt} = \sum_{j=1, i \neq j}^n \frac{\kappa_j}{2\pi} \frac{\hat{k} \times (\mathbf{r}_i - \mathbf{r}_j)}{\|\mathbf{r}_i - \mathbf{r}_j\|^2}, \quad (4.3)$$

where \hat{k} is a unit vector normal to the 2D plane. Each vortex is influenced by the induced velocity of every other point vortex in the system and the self-induced components are zero.

An important property of the point vortex system is that it exhibits a Hamiltonian struc-

ture [120, 146]. The Hamiltonian which represents the interaction kinetic energy between the point vortices is given by

$$H = -\frac{1}{4\pi} \sum_i^n \sum_{j, i \neq j}^n \kappa_i \kappa_j \log \|\mathbf{r}_i - \mathbf{r}_j\|, \quad (4.4)$$

such that $dH/dt = 0$ for the unforced point vortex system. In addition to time-invariance, the Hamiltonian also exhibits invariance with respect to the translation and rotation groups [36]. As a consequence of Noether's theorem, the point vortex system conserves the linear and angular impulse due to translation and rotational invariance, respectively. The impulses are

$$\text{Linear impulse : } (X, Y) = \sum_i^n \kappa_i \mathbf{r}_i, \quad \text{Angular impulse : } A = \sum_i^n \kappa_i \|\mathbf{r}_i\|^2. \quad (4.5)$$

By observing the form of these equations, we note that the linear impulse is directly proportional to the center of vorticity of the point vortex system, and the angular impulse is directly proportional to the variance of vorticity of the vortex configuration. The two vortex system ($n = 2$) is trivially integrable. Through a series of canonical transformations [3, 6], the three-vortex identical point vortex system ($n = 3$) can be completely described by the specification of the Hamiltonian and impulses. This configuration has been studied exhaustively on different geometries [19, 20], and in itself has been useful for understanding higher dimensional vortex phenomena [168].

In general, the motion of a system with four or more vortices ($n \geq 4$) is non-integrable, exhibiting a range of dynamical behaviors including chaos. The invariant dynamics of the point vortex system reduces to

$$\frac{d}{dt} \boldsymbol{\varphi} = 0 \quad (4.6)$$

with $\boldsymbol{\varphi} = [H, A, X, Y]^T$. The invariant dynamics represents the global behavior of the point vortex system as opposed to the local dynamics for each point vortex of the Biot-Savart law. Next, we introduce forcing inputs which alter the global dynamics of the system of point vortices.

4.2.2 Multi-invariant control of point vortices

The (forced) Biot-Savart equation can be written as

$$\frac{d}{dt} \hat{\mathbf{R}} = \mathbf{N}(\hat{\mathbf{R}}) + \mathbf{f}, \quad (4.7)$$

where, $\hat{\mathbf{R}} = [\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n]^T \in \mathbb{R}^{2n}$, $\mathbf{N}(\hat{\mathbf{R}})$ is the nonlinear term on the right-hand side of Eq. (4.3), and \mathbf{f} is the external forcing added to the flow. In this work, we introduce this external forcing through “virtual vortices”. A physical example of these virtual vortices could be the use of rotating cylinders in the domain. Rotating stirring rods are common in mixing applications. Rotating cylinders can also be used to model fans or propellers. Assuming that the vortices are sufficiently far from the stirring rods/cylinders, the viscous effects of their interaction with the vortices can be ignored.

For a system with a large number of point vortices, multiple virtual vortices can be considered. For a system of m actuator cylinders, the forcing input $\mathbf{f} = \mathbf{B}\mathbf{u}$ is defined such that

$$\mathbf{B} = \begin{bmatrix} D(\mathbf{r}_1, \mathbf{a}_1) & D(\mathbf{r}_1, \mathbf{a}_2) & \dots & D(\mathbf{r}_1, \mathbf{a}_m) \\ D(\mathbf{r}_2, \mathbf{a}_1) & D(\mathbf{r}_2, \mathbf{a}_2) & \dots & D(\mathbf{r}_2, \mathbf{a}_m) \\ \vdots & \vdots & \vdots & \vdots \\ D(\mathbf{r}_n, \mathbf{a}_1) & D(\mathbf{r}_n, \mathbf{a}_2) & \dots & D(\mathbf{r}_n, \mathbf{a}_m) \end{bmatrix} \in \mathbb{R}^{2n \times m}, \quad (4.8)$$

where \mathbf{a}_i is the position of actuator cylinders in the domain. The interaction function D computes the induced velocity of the control cylinders (of unit strength) on the vortices in the domain, given by the same law as the interaction between vortices as

$$D(\mathbf{r}_i, \mathbf{a}_j) = \frac{1}{2\pi} \frac{\hat{k} \times (\mathbf{r}_i - \mathbf{a}_j)}{\|\mathbf{r}_j - \mathbf{a}_j\|^2}. \quad (4.9)$$

The influence of the actuator cylinder on the vortices is modulated by the circulation strength of each cylinder $\mathbf{u} = [u_1, u_2, \dots, u_m]^T \in \mathbb{R}^m$.

We describe the global invariant dynamics of the forced system as

$$\frac{\partial \varphi}{\partial \hat{\mathbf{R}}} \cdot \frac{d}{dt} \hat{\mathbf{R}} = \frac{\partial \varphi}{\partial \hat{\mathbf{R}}} \cdot \mathbf{N}(\hat{\mathbf{R}}) + \frac{\partial \varphi}{\partial \hat{\mathbf{R}}} \cdot \mathbf{B}\mathbf{u}. \quad (4.10)$$

The first term on the right hand side of the equation cancels to zero from Eq. (4.6) yielding

$$\frac{d}{dt}\varphi = \nabla_{\hat{\mathbf{R}}}\varphi \cdot \mathbf{B}\mathbf{u}. \quad (4.11)$$

The right-hand side of this equation is the directional derivative of the invariants with respect to the state $\hat{\mathbf{R}}$ of the vortices. We can explicitly calculate the gradient of the invariants with respect to the vortex position from Eq. (4.4) and Eq. (4.5). Due to the symplectic structure of vortex dynamics and Hamiltonian systems, the gradient of the Hamiltonian is orthogonal and equal in magnitude to $d\mathbf{R}/dt$. Thus, it can be calculated by swapping the x and y components of Eq. (4.3). To compute the circulation strengths of the actuator cylinders \mathbf{u} , we use model predictive control using Eq. (4.11) as our model.

4.2.3 Model predictive control

Model predictive control (MPC) [1, 34, 46, 56, 91, 109, 117] has gained widespread adoption due to its success in a range of applications, its ability to incorporate customized cost functions and constraints, and extensions to nonlinear systems. In particular, MPC has become the *de-facto* standard advanced control method in process industries [129] and has gained considerable traction in the aerospace industry due to its versatility [46].

MPC represents an optimal control problem over a receding horizon, subject to system dynamics and constraints, to determine the next control action. The optimization problem aims to solve for a sequence of control inputs $\{\mathbf{u}(t_0), \mathbf{u}(t_1), \dots, \mathbf{u}(t_0 + T_H)\}$ over the time horizon T_H that minimizes a pre-defined objective function J . The time horizon $T_H = N\Delta t$, where N is the number of time steps over which the optimization is calculated, and Δt is the discrete size of the time step of $d\tau$, given in the integral in Eq (4.12). Typically, only the first control input $\mathbf{u}(t_0)^{\text{opt}}$ is applied. The optimization problem is re-initialized each time a new measurement is collected and, thus, adaptively determines optimal control actions adjusting to model discrepancies and changing conditions in the environment and disturbances. The most critical part of MPC is the identification of a dynamical model that accurately and efficiently represents the system behavior in the presence of actuation. If the model is linear, minimization of a quadratic cost functional subject to linear con-

straints results in a tractable convex problem. In this work, we combine nonlinear MPC with control-affine models describing the dynamical evolution of the conserved quantities under the influence of an external control variable.

The receding-horizon optimization problem can be stated as follows. Nonlinear MPC [1] aims to minimize the following quadratic objective function,

$$J = \int_{t_0}^{t_0+T_H} [\mathbf{e}(\tau)^T \mathbf{Q} \mathbf{e}(\tau) + \mathbf{u}(\tau)^T \mathbf{R} \mathbf{u}(\tau)] d\tau, \quad (4.12)$$

subject to, nonlinear system dynamics,

$$\frac{d}{dt} \mathbf{x}(t) = \mathbf{N}(\mathbf{x}(t)) + \mathbf{B} \mathbf{u}(t), \quad (4.13)$$

and input constraints,

$$\mathbf{u}_{\min} \leq \mathbf{u}(t) \leq \mathbf{u}_{\max}. \quad (4.14)$$

where $\mathbf{u}(t)$ is the control input, $\mathbf{e}(t) \triangleq \mathbf{x}(t) - \mathbf{x}_{\text{goal}}$ is the difference in the desired state and predicted state along a finite-horizon trajectory. In our work specifically, $\mathbf{e}(t) = \boldsymbol{\varphi}(t) - \boldsymbol{\varphi}_{\text{ref}}$. The weight matrices $\mathbf{R} \in \mathbb{R}^{q \times q}$, and $\mathbf{Q} \in \mathbb{R}^{n \times n}$ are positive semi-definite and penalize the inputs, and deviations of the predicted output along a trajectory respectively, and set their relative importance. We define the control sequence to be solved over the receding horizon as $\{\mathbf{u}(t_0), \mathbf{u}(t_1), \dots, \mathbf{u}(t_0 + T_H)\}$ given the measurement $\mathbf{x}(t_0)$. The measurement $\mathbf{x}(t_0)$ is the current output of the system, whose dynamics are governed by the invariants of the Biot-Savart equations (4.11), and is used to provide the initial condition for the optimization problem.

If the model is linear, minimization of a quadratic cost functional subject to linear constraints results in a tractable convex problem. Nonlinear models may yield significant improvements; however, they render MPC a nonlinear program, which can be expensive to solve, making it particularly challenging for real-time control. Fortunately, improvements in computing power and advanced algorithms are increasingly enabling nonlinear MPC for real-time applications. Generally, MPC applied to the Navier-Stokes equations would result in a high-dimensional, nonlinear optimal control problem. By reformulating the dy-

namics in terms of conserved quantities, we aim to reformulate the problem as a bilinear, significantly lower-dimensional problem which is significantly more feasible.

4.2.4 Specific control setup

We set the weight matrix \mathbf{Q} as

$$\mathbf{Q} = \begin{bmatrix} Q_H & 0 & 0 & 0 \\ 0 & Q_A & 0 & 0 \\ 0 & 0 & Q_X & 0 \\ 0 & 0 & 0 & Q_Y \end{bmatrix}. \quad (4.15)$$

The limits on actuation are such that

$$|\mathbf{u}(t)| \leq 1. \quad (4.16)$$

The actuator limits are set so that the maximum circulation the actuator produces does not exceed the circulation of the vortices. Thus, the controller will make use of the natural dynamics of the system. The model we use for model-predictive control is given by Eq (4.11). We optimize over this model to generate an action $\mathbf{u}(t)$, which is then fed into the \mathbf{f} term in Eq (4.7) for simulation. The time horizon is given by $N\Delta t = 0.03$ time units with time step of $\Delta t = 0.01$ and $N = 3$. The $\nabla_{\hat{\mathbf{R}}}\varphi \cdot \mathbf{B}$ matrix changes with change in state. However, for the purposes of optimization, we assume that over our short time horizon T_H , the state $\hat{\mathbf{R}}$ is constant. The MPCTools [138] and CasADi [2] packages were used for the computation of control in this chapter.

4.3 Single invariant control results

In this section, we show how changes in invariants manifest in changes in the dynamics of vortices. We use the invariants to influence transitions between chaotic, quasiperiodic, and periodic states, which can in turn impact the mixing properties of the system. For this study, we investigate the dynamics of a 4-vortex system, as 4-vortices can exhibit chaotic

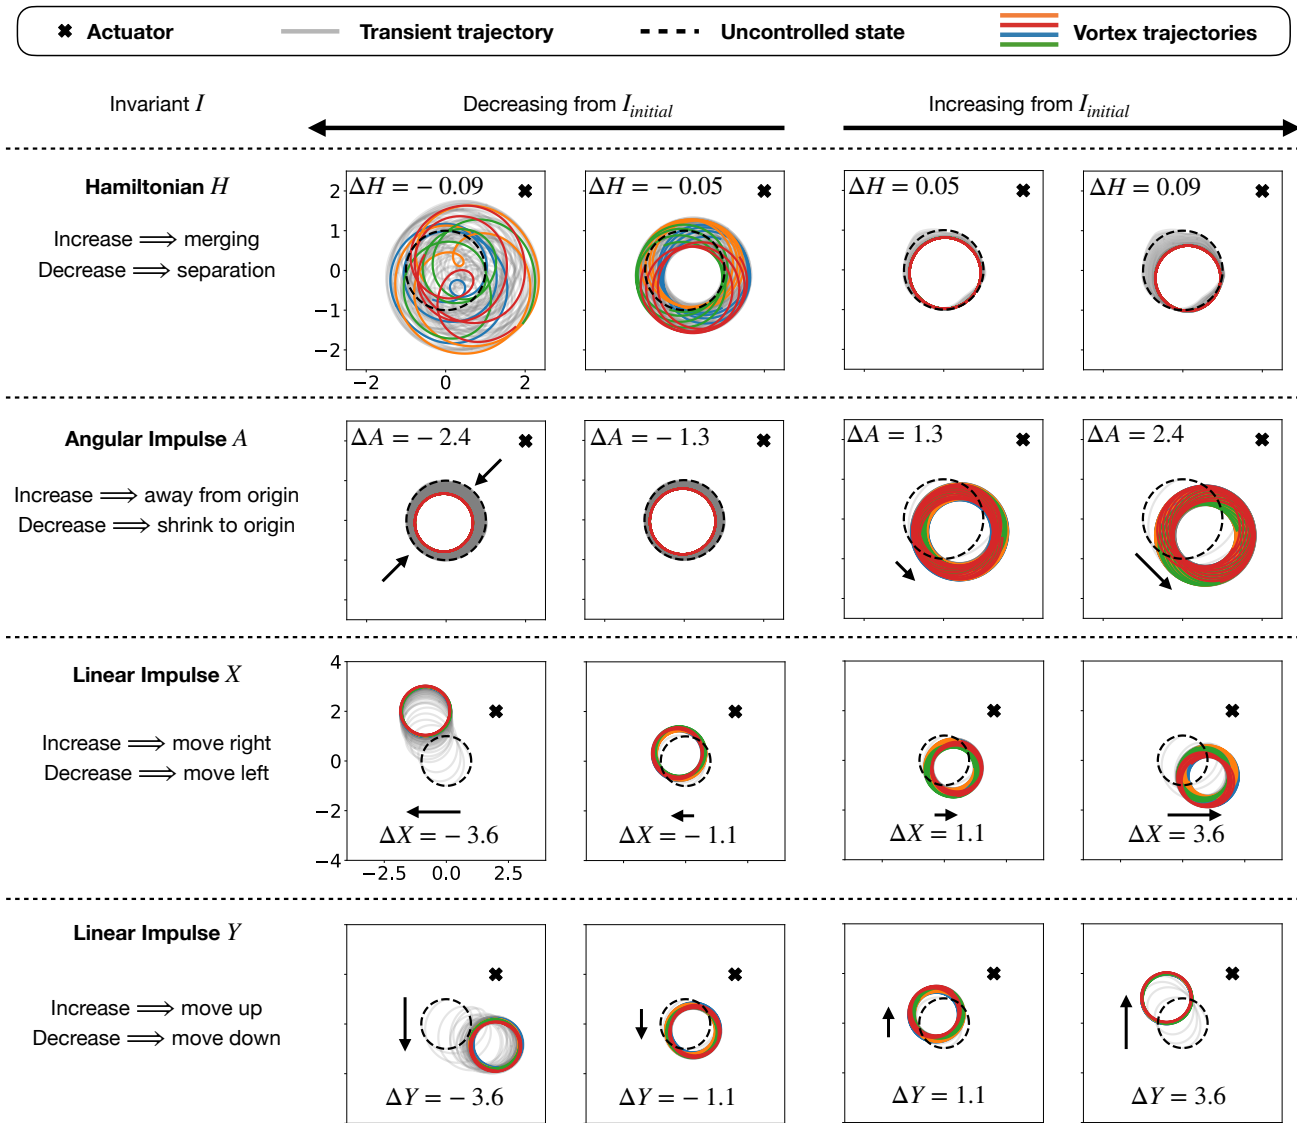


Figure 4.2: Impact of reference tracking of the four invariants of the vortex system on the vortices' trajectories using an actuator in a fixed position. From top row to bottom: (1) Different dynamic regimes appear depending on a higher or lower reference value for the Hamiltonian. (2) Increasing angular impulse translates the vortices away from the origin. Negative A shows that vortices merge closer in a similar strategy to increasing the Hamiltonian. (3-4) Change in linear impulse manifests as translation of the vortex configuration.

behaviour due to non-integrability of the governing equations. These configurations have also been studied extensively in the past [6, 18, 21]. The initial vortex cores are arranged in a square. The circulation κ of all our vortices are taken to be 1.0 for this configuration, so that the uncontrolled behaviour is for the vortices to move in a periodic circle. This is known as the Thomson equilibrium. The invariant values for this configuration are $(H, A, X, Y) = (-0.22, 4, 0, 0)$ when the vortices are all equally spaced at 1 unit from the origin. Over the next few subsections, we explore the changes in vortex dynamics with changing invariant values, and the form of the control signal used by MPC used to change the vortex dynamics.

4.3.1 *Controlled vortex dynamics*

The plots corresponding to positive ΔH show that increasing the Hamiltonian has the effect of bringing vortices closer together. Decreasing the Hamiltonian has the effect of vortex cores moving apart from each other. This observation is in agreement with the form of the equation of the Hamiltonian, which is the sum of the log of the distance between each vortex pair. When there are multiple vortices present (not arranged symmetrically in a square), the MPC control law will pick one pair of vortices and bring them closer together when the Hamiltonian is increased. Using MPC and the conserved quantity framework, it is also possible to make the Hamiltonian oscillate or follow complicated trajectories. This manifests in the system as the vortices moving in closer and farther apart. The panels in the first row of Figure 4.2 show the trajectories of the 4-vortex system for different reference values for the Hamiltonian to track. We observe that decreasing the Hamiltonian makes the vortex system enter a quasi-periodic regime and further decrease makes it transition into a chaotic regime as can be seen by taking the Fourier transform of the x-component of trajectories (after it has converged to the steady state). However, increasing the Hamiltonian moves it into a periodic regime with slight translation. The mean peak in the Fourier transform plots correspond to translation of the vortex system.

The linear impulse is linked to the mean or center of vorticity for the system of vortex cores. Therefore, we observe that changing the linear impulse results in translation of the

vortex configuration. The X or Y impulse corresponds to translations along the x or y axis respectively. This is because moving the vortex distribution changes the center of vorticity distribution.

Unlike the previous two cases, the vortex dynamics for angular impulse changes depending on if we are increasing or decreasing the angular impulse. For increasing angular impulse, the behavior is to translate the vorticity distribution, but decreasing angular impulse occurs by moving inwards symmetrically towards the origin. The angular impulse of a vortex configuration is proportional to the variance from the origin. Thus, one way of increasing the variance from the origin is by translating the distribution away from the origin.

A key observation across all the cases is that changing one invariant also causes a change in other invariants. For example, angular impulse is increased through translation, which, also increases the linear impulse. We can also increase the angular impulse without translation with MPC. We discuss this in a later section.

4.3.2 MPC Control Law

Hamiltonian

We observe in Figure 4.3 that the control commanded by the MPC algorithm to increase the Hamiltonian is an oscillating input, where the frequency of the oscillation increases over time. For decreasing, the switching frequency decreases. In the Hamiltonian control plot, the H was increased/decreased by 0.07 units. Vainchtein et al. [182] showed that merging a vortex pair through control can be performed by perfectly timing a train of Dirac delta pulses in phase with the rotation of the vortex pair. Given that increasing the Hamiltonian can be viewed as bringing vortices closer (or merging), our results are consistent with this observation, as the oscillating control is also in phase with the rotation of the four vortices. Moreover, as the vortices move closer, the frequency of cylinder vorticity switching increases, to match the co-rotating frequency of the 4-vortex system.




Figure 4.3: MPC control signals (circulation) corresponding to the different cases of invariant control displayed in Fig.2. Red curves are associated with a higher reference value, while blue curves represent a lower reference value in the objective function. From left to right: (1) We can observe here that the red curve is oscillatory with increasing frequency in time while the blue curve decreases the frequency in time. (2) While an increase in angular momentum results in a step function, a decrease yields consistent periodic behavior similar to the Hamiltonian but without convergence. (3) In contrast, convergence can be quickly achieved using a step function for both increasing and decreasing linear impulse.

Linear impulse

In Figure 4.3 we observe that the control law used for increasing the linear impulse takes the form of a step function. In the plot, the linear impulse was increased/decreased by 0.2 units. A step function control law causes the vortex configuration to translate along a curved trajectory as seen in Figure 4.2. Changing the X linear impulse causes a translation in the x-direction and the effect is similar in the y-direction. We observe in Figure 4.2 panels corresponding to the linear impulse that the translation also causes a net shearing effect. Vortices closer to the actuator translate more than vortices farther away from it because vortices impact each other proportional to $\frac{1}{r}$. This shearing effect is also the mechanism by which vortex trajectories move from perfectly periodic states to quasi-periodic states. The paper by Aref et al. [6] diagrammatically maps out the phase space of the 4-vortex system, showing the proximity of periodic and quasi-periodic states on a KAM

(Kolmogorov–Arnold–Moser) torus.

Angular impulse

In Figure 4.3, we observe that for increasing the angular impulse, the control law is a step function, similar to the case of linear impulse. The cylinder actuation produces translation by “pivoting” the vortex system around the actuator. This means that the actuator can only translate up to a limit (where the vortex distribution moves diametrically across from the starting point, taking the actuator as the center). For the case of an actuator located at $(2, 2)$, the maximum translation corresponds to an angular impulse $\Delta A \approx 130$. For angular impulse, the direction of translation is the direction where the control expenditure is the least, as given in the linear impulse results. The control law for decreasing the angular impulse is the same control law used in increasing the Hamiltonian. This is due to the fact that increasing the Hamiltonian causes merging, which brings the vortex distribution closer to the origin when starting from a uniform 4-vortex configuration. In the plots, the angular impulse was increased/decreased by 2 units.

4.4 Multi-invariant control results

In the previous section, we discussed single-invariant control. This section is about multi-invariant control. As discussed earlier, the Biot-Savart equations for vortex dynamics has 4 conserved quantities, and we can change multiple conserved quantities at once. The advantage of simultaneously changing multiple conserved quantities is to change the distribution of vortices to states that are not possible through single invariant control. We demonstrate this with two examples.

4.4.1 Restricting translation in angular impulse control

We observed in §4.3.2 that the angular impulse increases through translation in single-invariant control. This also implies that single-invariant control changes the quantity A by changing the linear impulse. We can now use multi-invariant control to change the angular impulse by penalizing the linear impulse. This forces the controller to keep the linear

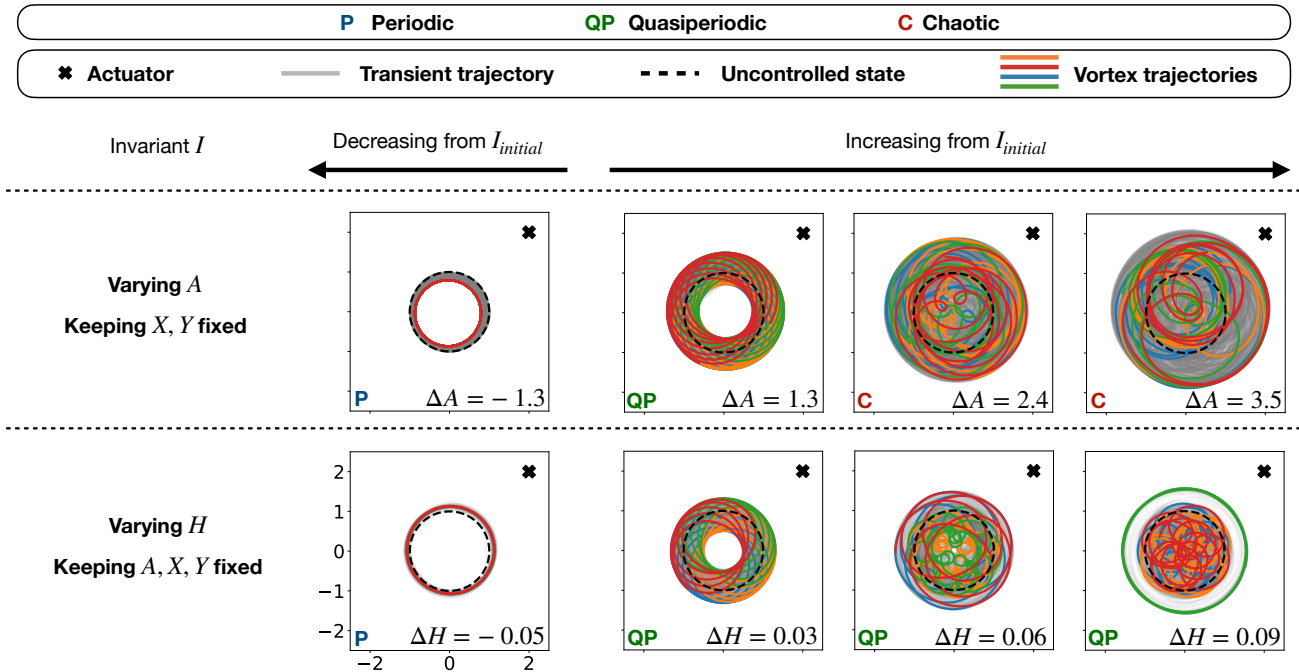


Figure 4.4: This figure summarizes the resulting dynamics from multi-invariant control. The first row shows the vortex trajectories with the reference change in angular impulse. The second row is for the change in Hamiltonian

impulse close to 0. The result can be seen in Figure 4.4, where, as we increase the angular impulse, the trajectories move outward in a manner that keeps the linear impulse close to 0. We observe that for our choice of MPC parameters, the trajectories transition from periodic to quasiperiodic to chaotic as the angular impulse increases. Since the controller is unable to hold the linear impulse exactly at 0, we observe chaos upon larger increases in A , which can be understood using the right plot schematic in Figure 4.5. Moving the center of vorticity causes chaotic transitions.

4.4.2 Enforcing symmetry in Hamiltonian control

We observed earlier that using single-invariant control to change the Hamiltonian changes the other conserved quantities. In this example, we use multi-invariant control to change the Hamiltonian, but use the MPC cost function to penalize changes in the other invari-

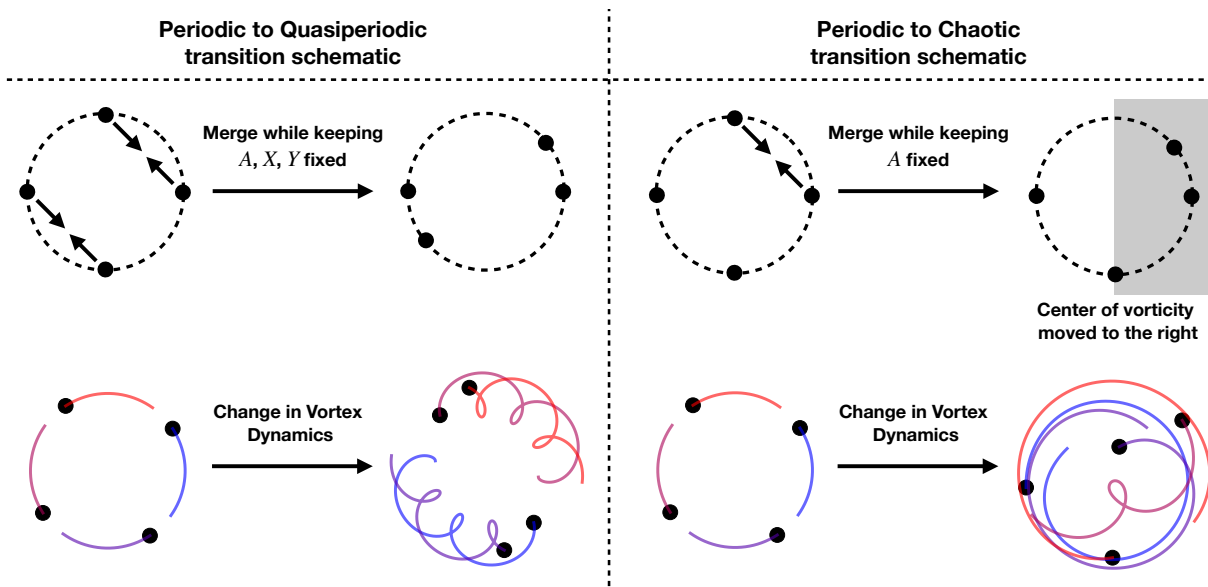


Figure 4.5: **Schematic diagram for transitions:** On the left, we show how multi-invariant control can be used to move between periodic and quasi-periodic regimes as seen in Figure 4.1. Here, the average radius/variance of the vortex distribution and the center of vorticity or the mean can both be controlled. When the Hamiltonian is increased under these constraints, the vortex merging is constrained to states along the dotted circle, which exhibit quasi-periodic behavior. Transitioning between different dynamic regimes is very limited using single-invariant control, whereas multi-invariant control enables one to reach a more diverse set of states. On the right, we show that a similar transition is possible from periodic to chaotic behavior by shifting the centre of vorticity to the right, thereby causing an asymmetry in the vortex configuration. These particular configurations have been exhaustively studied independently in [18]. The key difference between these configurations is the angles the vortices make with the origin.

ants. Particularly, when we increase the Hamiltonian, vortices come closer and move in a smaller periodic circle as seen in the top row of Figure 4.2. However, if the other invariants are not allowed to change, we can move to quasi-periodic states as the vortices are now constrained to merge in a way that preserves the variance or center of vorticity (as shown in Figure 4.5). We note that the final states reached in Figure 4.4 are not exactly the final states described on the left in Figure 4.5, however, Figure 4.5 provides a useful guide for understanding switching between regimes. In Figure 4.2, we observe that in the case of increasing Hamiltonian, the vortices merge such that translating the center of vorticity is

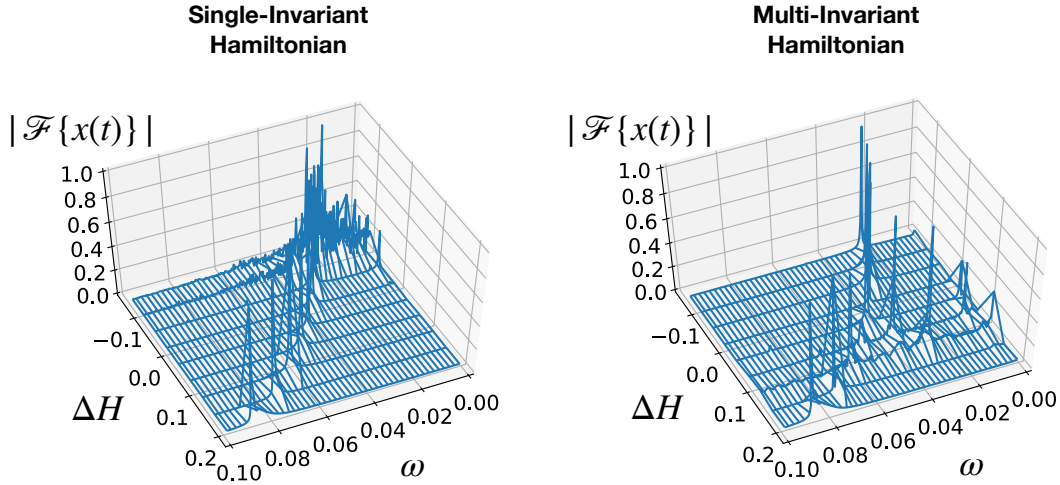


Figure 4.6: This plot shows the power spectral density of the x-component of a single vortex for both single- and multi-invariant Hamiltonian control. This information can be used to classify different dynamic regimes as shown in Figure 1. We see that in single-invariant control, increasing the Hamiltonian yields periodic behavior with a single frequency, whereas decreasing the Hamiltonian modifies the dynamics towards chaotic behavior with broadband frequency characteristics. However, for multi-invariant control, this chaotic transition does not occur. Instead, when increasing the Hamiltonian quasi-periodicity with multiple frequencies can be observed. This is due to enforcing symmetry in the case of multi-invariant control.

prohibited, and maintaining the same average radius/variance from the center. We also make note that the symmetry of the configuration is generally broken when vortices move to chaotic states.

We also perform a sweep through different increases in H with and without keeping the other invariants fixed and compute the frequency spectra of the vortex time series. Similar numerical studies have been performed in [37] to highlight the transition between different regimes. The power spectral density for the time series of an x-component for different baseline regimes can be seen in Figure 4.1. We observe that when the angular impulse and center of vorticity are fixed, transitions to chaotic states are prevented. We observe on the right in Figure 4.6 that the final states are quasi-periodic when performing multi-invariant control. At $\Delta H = 0.09$, the 4-vortex system groups into a cluster of 3 vortices, as though it is a single vortex, and the remaining vortex orbits the 3-vortex system. The

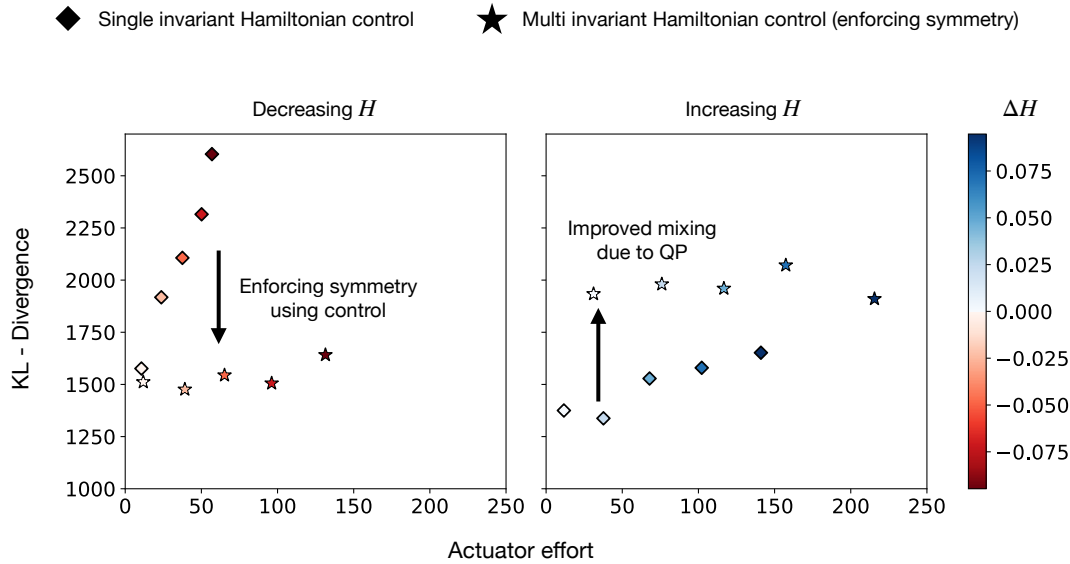


Figure 4.7: This plot shows the comparison between single invariant and multi invariant Hamiltonian control in terms of mixing. The plots are scatter plots of actuator effort $\int u(t)dt$ over time on the x -axis and the KL-Divergence on the y -axis. We observe that enforcing symmetry decreases the KL-divergence in the case of decreasing Hamiltonian. This is due to vortex dynamics being restricted to periodic behavior instead of chaotic. In the case of increasing Hamiltonian, the mixing improves due to a transition to quasiperiodicity. We also observe that multi-invariant control uses more energy than single-invariant control.

states of decreasing Hamiltonian converge to their reference value at an extremely slow rate, and for these cases, the simulation was ended at 80% of the reference value. Over our simulations, we observed that states which retain the periodic nature, but reducing radius, which occur also in the case of decreasing single-invariant A as seen in Figure 4.2, take extremely long to reach the reference state and hence, use a large amount of energy.

4.5 Mixing

To enhance mixing, we are interested in rapidly dispersing some substance within the fluid such that the substance is well distributed within the domain. There has been past work using chaotic advection through fixed stirring rods [4, 22], and methods such as adjoint optimization [44]. In this section, we formulate the mixing problem in terms of control

of conserved quantities in the 4-vortex system. The four vortices act as “moving stirring rods” for the passively advecting distribution.

For the simulations in this section, we advect a grid of particles, using a symplectic Euler scheme, for a short time through the flow field generated by the 4 vortices after the invariants converge to the reference values. These results are plotted in Figure 4.1, where the left plot shows the initial colored tracer distribution. Then, we compute the KL-divergence of the distribution, with respect to the initial distribution of the grid of particles to compare our results. The KL-divergence is a measure of similarity between two probability distributions [28]. We initialize a grid of uniformly distributed points and observe how the histogram of this grid distorts under the action of the flow field induced by the vortex dynamics. The right panels in Figure 4.1 show the final distribution depending on the vortex configuration. The formula of KL-divergence is given by

$$KL = \sum_{x \in \chi} P(T) \log \frac{P(T)}{P(0)}, \quad (4.17)$$

where, $P(0)$ is the initial histogram distribution of tracer particles and $P(T)$ is the final distribution of tracer particles. The single-invariant and multi-invariant Hamiltonian control cases we encountered so far encompass the wide range of dynamical phenomena that the 4-vortex system can exhibit. Hence, for the mixing study, we use this data, summarized in Figure 4.7. On the left plot, we have the results for decreasing the Hamiltonian, and on the right, we have results for increasing the Hamiltonian. We find that in the regime of increasing Hamiltonian, for the same ΔH , multi-invariant control can generate greater KL-divergence than in the case of single-invariant control with a slight increase in actuation effort. It is possible to increase the KL-divergence the most by decreasing the Hamiltonian in single-invariant control. However, this causes the vortices to move apart chaotically. In applications to mixing in confined spaces, it may be beneficial to use multi-invariant control.

Code Availability

The code for this work has been made available on GitHub at <https://github.com/karkris41295/invariant-vrtx-ctrl>

Chapter 5

CONCLUSION

In this work, two parallel threads run that make use of different notions of coherent structures from dynamical systems theory. The first is the use of Lagrangian coherent structures for path planning *within* fluid flow fields, and the second is the use of Koopman eigenfunctions for the control of *of* a fluid flow (vortex dynamics). In the following sections, we review and summarize these threads, along with suggestions for future work.

5.1 Control within a flow field

5.1.1 Summary

Chapters 1-3 investigate the problem of a robot operating within a generic unsteady fluid flow field. Out of the discussion sections of each of these chapters, the main motivating idea was to determine if the control action \mathbf{u} of the robot can be written directly as a function of the (Lagrangian) coherent structures in the background. By this, we are interested in writing $\mathbf{u} = f(LCS)$. This would have massive implications for control algorithms in flow fields that are able to potentially compute control laws without a computationally expensive non-convex optimization algorithm in a chaotic system.

Following this thread of thought, Chapter 1 investigated how individual low-energy trajectories of a robot and the background LCS overlap over a range of planning parameters. Chapter 2 looked at the same problem from the perspective of *policies* or in other words, “What the instantaneous control action from each location in space has to be to achieve the control objective?”. The latter perspective is different from the former in that the former is more interested in individual trajectories. The duality between Chapters 1 and 2 is very similar to the duality between the Lagrangian and Eulerian views in fluid dynamics. Interestingly, a similar duality also exists in optimal control theory between the Pontryagin maximum principle approach to generating trajectories and the Bellman

method of computing instantaneous control in state space.

A highlight contribution of this thread can be found in Equation (3.5), closing the loop on if control can be traced back to coherent structures. This equation can also be used to explain the phenomena we encounter in Chapter 1 - why the control magnitude is largest at the FTLE fields, the direction of control vector is orthogonal to the FTLE ridges, why the trajectories for longer time horizons do not converge to a smooth trajectory or how the non-convexity increases with the time horizon. The following subsection covers selected future work in line with the overarching themes of this thesis, however, more ideas can also be found within the discussion sections of chapters 2 and 3.

5.1.2 Future Work

Optimal control theory

Value functions are scalar functions defined over the domain that map each state to the expected sum of future rewards from that state. From this perspective, control policies (in the context of the kinematic models considered in this work) are gradients of the value function that direct individual agents to “climb up” to states with greater values. Value functions are, generally speaking, solutions to the HJB equation. Modern equation-free control methods such as reinforcement learning ultimately attempt to solve the HJB equation approximately. Recent works along this direction in the context of navigating flow fields include [43, 187]. In Chapter 3, we show that cFTLE ridges can be deeply connected to the boundaries of a value function, where the cost function is most sensitive to perturbations. We also highlight that regions most sensitive to perturbation are also the regions where greatest energy is spent by the agent. These results can potentially aid developing methods for faster computation of value functions or policies. For example, only the value of states where the cFTLE field is large can be updated, since other states will be less sensitive to perturbations when a parameter like the R/Q is varied. This can save on redundant computation at states which do not change value under perturbation. Furthermore, a deeper connection can be potentially highlighted mathematically connecting the singular values/vectors of the flow map Jacobian, which are related to cFTLE, and the

value function.

Bifurcation and stability analysis

In Chapter 3, when the R/Q ratio approaches 0, the policy has the actuation capability to turn the goal state into an attracting fixed point. As we increase the R/Q ratio, the policy loses the ability to create a strongly attracting fixed point, and instead, creates a limit cycle oscillation around the goal state, as was reported previously in [87]. When there is no actuation (i.e., in the limit of R/Q going to ∞), trajectories in the double gyre are chaotic. This points to the possibility that changing the R/Q ratio generates successive bifurcations, ending in chaos. Invariant manifolds, cFTLE ridges in this case, play a crucial role in global bifurcations, which could be useful in understanding the change in stability described above [60, 61]. In the context of transport of autonomous agents, this would allow us to understand and predict the onset of loss of stability around the goal for different R/Q ratios, extrapolating from known positions of the cFTLE ridges. This would ultimately aid in tuning the R/Q parameter.

5.2 Control of a flow field

5.2.1 Summary

In Chapter 4, we investigate the control of vortex dynamics using vortex invariants, inspired by a recent Koopman-based control scheme. Specifically, the Hamiltonian and the linear and angular impulses are conserved quantities that form a set of Koopman eigenfunctions, and we demonstrate that it is possible to control these invariants to manipulate the overall vortex dynamics. We demonstrate this approach on the 4-vortex system, where there is a range of possible phenomena, including periodic, quasi-periodic, and chaotic dynamics. Using model predictive control, we show that it is possible to manipulate single invariants or multiple invariants at once, and that changing these invariants results in broad behavioral changes to the vortex dynamics (e.g., periodic to chaotic, or vice versa). Finally, we investigated the effects of invariant control on mixing by integrating a grid of passive tracers through the flow field generated by the 4-vortices and evaluating the

KL-divergence of the distribution.

Through our study, we found that increasing (or decreasing) the Hamiltonian causes merging (or separation) of vortices. This is achieved by oscillating the vorticity of our control cylinder at a frequency that depends on the vortex configuration. We also observed that changing a single conserved quantity using a vortex actuator causes a change in the other invariants. Multi-invariant control enables us to manipulate the dynamics of a 4-vortex system from periodic behavior to different regimes by specifying different invariant values. Importantly, we were able to do this by using a very short time horizon in model-predictive control. This is promising for the control of larger vortex systems, particularly when used with methods that cluster vortices [111, 118, 172, 175]. Our method could also be potentially combined with recent methods that extract point-vortex models from flow field data [41].

5.2.2 *Future Work*

There are many future directions based on this work. Extending this approach to the control of dissipative, viscous flows would be an interesting study. Such flow fields or vortex systems do not conserve Hamiltonian/kinetic energy, and therefore the approximate short-time conserved quantities can be learned, upon which, the framework used in this chapter can be applied. Investigating the control of higher-dimensional vortex systems would also bring this closer to relevant flows, such as turbulent shear layers [7]. Studying lower-dimensional vortex configurations has been useful for understanding higher-dimensional phenomena, including control of 2-vortex systems for wakes [38], and control of 3-vortex systems for plasma dynamics [127]. This study could also be extended to understand the role of conserved quantities in modeling vortex-body interactions, as in the case of robot/animal locomotion [82, 159, 185, 186], capturing three-dimensional fluid-structure interactions in the context of lift and drag on turbines [142, 143], and mitigating or deflecting large scale vortical structures in wakes behind aircraft [39, 113]. Many of these applications involve complicated boundary conditions that require the use of additional mirror vortices. Future work may involve formulating these objectives within the frame-

work of varying conserved quantities. In addition, the controllability and observability of point vortex systems can be more rigorously validated using differential geometric methods. Finally, the change of coordinates to conserved quantities can be incorporated into Hamiltonian/energy modeling frameworks such as port-Hamiltonian systems [112, 184].

BIBLIOGRAPHY

- [1] Frank Allgöwer, Rolf Findeisen, and Zoltan K Nagy. Nonlinear model predictive control: From theory to application. *J. Chin. Inst. Chem. Engrs*, 35(3):299–315, 2004.
- [2] Joel A E Andersson, Joris Gillis, Greg Horn, James B Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.
- [3] Hassan Aref. Motion of three vortices. *The Physics of Fluids*, 22(3):393–400, 1979.
- [4] Hassan Aref. Stirring by chaotic advection. *Journal of Fluid Mechanics*, 143:1–21, 1984.
- [5] Hassan Aref. Point vortex dynamics: a classical mathematics playground. *Journal of mathematical Physics*, 48(6):065401, 2007.
- [6] Hassan Aref and Neil Pomphrey. Integrable and chaotic motions of four vortices. I. the case of identical vortices. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 380(1779):359–387, 1982.
- [7] Hassan Aref and Eric D. Siggia. Vortex dynamics of the two-dimensional turbulent shear layer. *Journal of Fluid Mechanics*, 100(4):705–737, 1980.
- [8] V.I. Arnold. *Mathematical methods of classical mechanics*, volume 60. Springer, 1989.
- [9] Sanjeeva Balasuriya. Uncertainty in finite-time Lyapunov exponent computations. *Journal of Computational Dynamics*, 7(2):313, 2020.
- [10] Sanjeeva Balasuriya, Nicholas T Ouellette, and Irina I Rypina. Generalized Lagrangian coherent structures. *Physica D: Nonlinear Phenomena*, 372:31–51, 2018.
- [11] James G. Bellingham and Kanna Rajan. Robotics in remote and hostile environments. *Science*, 318(5853):1098–1102, 2007.

- [12] Simon A Berman, John Buggeln, David A Brantley, Kevin A Mitchell, and Thomas H Solomon. Transport barriers to self-propelled particles in fluid flows. *Physical Review Fluids*, 6(1):L012501, 2021.
- [13] Francisco J Beron-Vera, Maria J Olascoaga, and GJ Goni. Oceanic mesoscale eddies as revealed by Lagrangian coherent structures. *Geophysical Research Letters*, 35(12), 2008.
- [14] Francisco J Beron-Vera, María J Olascoaga, George Haller, Mohammad Farazmand, Joaquín Triñanes, and Yan Wang. Dissipative inertial transport patterns near coherent Lagrangian eddies in the ocean. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(8):087412, 2015.
- [15] Francisco J Beron-Vera, María J Olascoaga, George Haller, Mohammad Farazmand, Joaquín Triñanes, and Yan Wang. Dissipative inertial transport patterns near coherent Lagrangian eddies in the ocean. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(8):087412, 2015.
- [16] Pradeep Bhatta, Edward Fiorelli, Francois Lekien, Naomi Ehrich Leonard, Derek Pauley, Fumin Zhang, Ralf Bachmayer, Russ E Davis, David M Fratantoni, and Rodolphe Sepulchre. Coordination of an underwater glider fleet for adaptive ocean sampling. In *Proc. International Workshop on Underwater Robotics, Int. Advanced Robotics Programmed (IARP), Genoa, Italy*, 2005.
- [17] L. Biferale, F. Bonaccorso, M. Buzzicotti, P. Clark Di Leoni, and K. Gustavsson. Zermelo's problem: Optimal point-to-point navigation in 2D turbulent flows using reinforcement learning. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(10):103138, 2019.
- [18] Stefanella Boatto and Raymond T Pierrehumbert. Dynamics of a passive tracer in a velocity field of four identical point vortices. *Journal of Fluid Mechanics*, 394:137–174, 1999.

- [19] A. Borisov and Vladimir Lebedev. Dynamics of three vortices on a plane and a sphere - II. general compact case. *Regular and Chaotic Dynamics*, 3, 04 2005.
- [20] A. Borisov and Alexander Pavlov. Dynamics and statics of vortices on a plane and a sphere - I. *Regular and Chaotic Dynamics*, 3, 04 2005.
- [21] AV Borisov, A. Kilin, and Ivan Mamaev. Transition to chaos in dynamics of four point vortices on a plane. *Doklady Physics*, 51:262–267, 05 2006.
- [22] Philip L. Boyland, Hassan Aref, and Mark A. Stremler. Topological fluid mechanics of stirring. *Journal of Fluid Mechanics*, 403:277–304, 2000.
- [23] Amir E BozorgMagham and Shane D Ross. Atmospheric Lagrangian coherent structures considering unresolved turbulence and forecast uncertainty. *Communications in Nonlinear Science and Numerical Simulation*, 22(1-3):964–979, 2015.
- [24] Amir E BozorgMagham, Shane D Ross, and David G Schmale III. Real-time prediction of atmospheric Lagrangian coherent structures based on forecast data: An application and error analysis. *Physica D: Nonlinear Phenomena*, 258:47–60, 2013.
- [25] S. L. Brunton and J. N. Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2nd edition, 2022.
- [26] Steven L Brunton. Flow map composition to identify coherent structures. In *Proceedings 18th International Symposium on Flow Visualization*. ETH Zurich, 2018.
- [27] Steven L Brunton, Marko Budišić, Eurika Kaiser, and J Nathan Kutz. Modern Koopman theory for dynamical systems. *SIAM Review*, 64(2):229–340, 2022.
- [28] Steven L Brunton and J Nathan Kutz. *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press, 2019.
- [29] Steven L Brunton and Bernd R Noack. Closed-loop turbulence control: progress and challenges. *Applied Mechanics Reviews*, 67(5):050801, 2015.

- [30] Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- [31] Steven L Brunton and Clarence W Rowley. Fast computation of finite-time Lyapunov exponent fields for unsteady flows. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(1):017503, 2010.
- [32] Marko Budišić, Stefan Siegmund, Doan Thai Son, and Igor Mezic. Mesochronic classification of trajectories in incompressible 3d vector fields over finite times. 05 2016.
- [33] Michele Buzzicotti, Luca Biferale, Fabio Bonaccorso, Patricio Clark di Leoni, and Kristian Gustavsson. Optimal control of point-to-point navigation in turbulent time-dependent flows using reinforcement learning, 2021.
- [34] Eduardo F Camacho and Carlos Bordons Alba. *Model Predictive Control*. Springer Science & Business Media, 2013.
- [35] Fei Chai, Kenneth S. Johnson, Hervé Claustre, Xiaogang Xing, Yuntao Wang, Emmanuel Boss, Stephen Riser, Katja Fennel, Oscar Schofield, and Adrienne Sutton. Monitoring ocean biogeochemistry with autonomous platforms. *Nature Reviews Earth & Environment*, 1(6):315–326, 2020.
- [36] David M. F. Chapman. Ideal vortex motion in two dimensions: Symmetries and conservation laws. *Journal of Mathematical Physics*, 19(9):1988–1992, 1978.
- [37] A. T. Conlisk, Y. G. Guezennec, and G. S. Elliott. Chaotic motion of an array of vortices above a flat wall. *Physics of Fluids A: Fluid Dynamics*, 1(4):704–717, 1989.
- [38] Luca Cortelezzi. Nonlinear feedback control of the wake past a plate with a suction point on the downstream wall. *Journal of Fluid Mechanics*, 327:303 – 324, 11 1996.
- [39] Jeffrey Crouch. Airplane trailing vortices and their control. *Comptes Rendus Physique*, 6(4-5):487–499, 2005.

- [40] J. O. Dabiri. Optimal vortex formation as a unifying principle in biological propulsion. *Annual Review of Fluid Mechanics*, 41:17–33, 2009.
- [41] Darwin Darakananda, André Fernando de Castro da Silva, Tim Colonius, and Jeff D Eldredge. Data-assimilated low-order vortex modeling of separated flows. *Physical Review Fluids*, 3(12):124701, 2018.
- [42] Michael Dellnitz, Gary Froyland, and Oliver Junge. The algorithms behind GAIO—Set oriented numerical methods for dynamical systems. In *Ergodic theory, analysis, and efficient simulation of dynamical systems*, pages 145–174. Springer, 2001.
- [43] Manan M Doshi, Manmeet S Bhabra, and Pierre FJ Lermusiaux. Energy–time optimal path planning in dynamic flows: Theory and schemes. *Computer Methods in Applied Mechanics and Engineering*, 405:115865, 2023.
- [44] Maximilian F Ettl and Peter J Schmid. Mixing enhancement in binary fluids using optimised stirring strategies. *arXiv preprint arXiv:1911.10760*, 2019.
- [45] Jeff D Eldredge and Anya R Jones. Leading-edge vortices: mechanics and modeling. *Annual Review of Fluid Mechanics*, 51:75–104, 2019.
- [46] Utku Eren, Anna Prach, Başaran Bahadır Koçer, Saša V Raković, Erdal Kayacan, and Behçet Açıkmeşe. Model predictive control in aerospace systems: Current state and opportunities. *Journal of Guidance, Control, and Dynamics*, 2017.
- [47] Mohammad Farazmand and George Haller. Computing Lagrangian coherent structures from their variational theory. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(1):013128, 2012.
- [48] Edward Fiorelli, Naomi Ehrich Leonard, Pradeep Bhatta, Derek A Paley, Ralf Bachmayer, and David M Fratantoni. Multi-AUV control and adaptive sampling in Monterey Bay. *IEEE Journal of Oceanic Engineering*, 31(4):935–948, 2006.
- [49] Eric Forgoston, Simone Bianco, Leah B Shaw, and Ira B Schwartz. Maximal sensitive

- dependence and the optimal path to epidemic extinction. *Bulletin of mathematical biology*, 73(3):495–514, 2011.
- [50] Trygve O. Fossum, Glaucia M. Fragoso, Emlyn J. Davies, Jenny E. Ullgren, Renato Mendes, Geir Johnsen, Ingrid Ellingsen, Jo Eidsvik, Martin Ludvigsen, and Kanna Rajan. Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics*, 4(27), 2019.
- [51] E. Franco, D. N. Pekarek, J. Peng, and J. O. Dabiri. Geometry of unsteady fluid transport during fluid-structure interactions. *Journal of Fluid Mechanics*, 589:125–145, 2007.
- [52] G. Froyland and K. Padberg. Almost-invariant sets and invariant manifolds – connecting probabilistic and geometric descriptions of coherent structures in flows. *Physica D*, 238:1507–1523, 2009.
- [53] G. Froyland, N. Santitissadeekorn, and A. Monahan. Transport in time-dependent dynamical systems: Finite-time coherent sets. *Chaos*, 20(4):043116–1–043116–16, 2010.
- [54] Gary Froyland. Statistically optimal almost-invariant sets. *Physica D: Nonlinear Phenomena*, 200(3):205–219, 2005.
- [55] Carlos E Garcia, David M Prett, and Manfred Morari. Model predictive control: Theory and practice—A survey. *Automatica*, 25(3):335–348, 1989.
- [56] Jorge L Garriga and Masoud Soroush. Model predictive control tuning methods: A review. *Industrial & Engineering Chemistry Research*, 49(8):3505–3515, 2010.
- [57] Guillaume Gauthier, Matthew T Reeves, Xiaoquan Yu, Ashton S Bradley, Mark Baker, Thomas A Bell, Halina Rubinsztein-Dunlop, Matthew J Davis, and Tyler W Neely. Negative-temperature onsager vortex clusters in a quantum fluid. *arXiv preprint arXiv:1801.06951*, 2018.

- [58] DR González, RL Speth, DV Gaitonde, and MJ Lewis. Finite-time Lyapunov exponent-based analysis for compressible flows. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26(8):083112, 2016.
- [59] M. A. Green, C. W. Rowley, and G. Haller. Detection of Lagrangian coherent structures in 3D turbulence. *Journal of Fluid Mechanics*, 572:111–120, 2007.
- [60] John Guckenheimer, Bernd Krauskopf, Hinke M Osinga, and Björn Sandstede. Invariant manifolds and global bifurcations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(9):097604, 2015.
- [61] John Guckenheimer and Alexander Vladimirsky. A fast method for approximating invariant manifolds. *SIAM Journal on Applied Dynamical Systems*, 3(3):232–260, 2004.
- [62] Peter Gunnarson, Ioannis Mandralis, Guido Novati, Petros Koumoutsakos, and John O Dabiri. Learning efficient navigation in vortical flow fields. *arXiv preprint arXiv:2102.10536*, 2021.
- [63] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D*, 149:248–277, 2001.
- [64] G. Haller. Exact theory of unsteady separation for two-dimensional flows. *Journal of Fluid Mechanics*, 512:257–311, 2004.
- [65] G. Haller. An objective definition of a vortex. *Journal of Fluid Mechanics*, 525:1–26, 2005.
- [66] George Haller. Lagrangian coherent structures from approximate velocity data. *Physics of fluids*, 14(6):1851–1861, 2002.
- [67] George Haller. A variational theory of hyperbolic Lagrangian coherent structures. *Physica D: Nonlinear Phenomena*, 240(7):574–598, 2011.
- [68] George Haller. Lagrangian coherent structures. *Annual Review of Fluid Mechanics*, 47:137–162, 2015.

- [69] George Haller and Themistoklis Sapsis. Where do inertial particles go in fluid flows? *Physica D*, 237(5):573–583, 2008.
- [70] Christoffer R Heckman, M Ani Hsieh, and Ira B Schwartz. Controlling basin break-out for robots operating in uncertain flow environments. In *Experimental Robotics*, pages 561–576. Springer, 2016.
- [71] Philip Holmes and John Guckenheimer. *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, volume 42 of *Applied Mathematical Sciences*. Springer-Verlag, Berlin, Heidelberg, 1983.
- [72] Tobias Gybel Hovgard, Lars FS Larsen, John Bagterp Jørgensen, and Stephen Boyd. Fast nonconvex model predictive control for commercial refrigeration. *IFAC Proceedings Volumes*, 45(17):514–521, 2012.
- [73] Tamer Inanc, Shawn C Shadden, and Jerrold E Marsden. Optimal trajectory generation in ocean flows. In *Proceedings of the 2005, American Control Conference, 2005.*, pages 674–679. IEEE, 2005.
- [74] Shi Jiang, Fei-fei Jin, and Michael Ghil. Multiple equilibria, periodic, and aperiodic solutions in a wind-driven, double-gyre, shallow-water model. *Journal of Physical Oceanography*, 25(5):764–786, 1995.
- [75] Yusheng Jiao, Feng Ling, Sina Heydari, Nicolas Heess, Josh Merel, and Eva Kanso. Learning to swim in potential flow. *Physical Review Fluids*, 6(5):050505, 2021.
- [76] Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of koopman eigenfunctions for control. *arXiv preprint arXiv:1707.01146*, 2017.
- [77] Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Discovering conservation laws from data for control. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 6415–6421. IEEE, 2018.
- [78] Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Sparse identification of non-

- linear dynamics for model predictive control in the low-data limit. *Proceedings of the Royal Society A*, 474(2219):20180335, 2018.
- [79] Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Data-driven approximations of dynamical systems operators for control. *arXiv preprint arXiv:1902.10239*, 2019.
- [80] Douglas H Kelley, Michael R Allshouse, and Nicholas T Ouellette. Lagrangian coherent structures separate dynamically distinct regions in fluid flows. *Physical Review E*, 88(1):013017, 2013.
- [81] Douglas H Kelley and Nicholas T Ouellette. Emergent dynamics of laboratory insect swarms. *Scientific reports*, 3(1):1–7, 2013.
- [82] Scott David Kelly and Hailong Xiong. Self-propulsion of a free hydrofoil with localized discrete vortex shedding: analytical modeling and simulation. *Theoretical and Computational Fluid Dynamics*, 24(1-4):45–50, 2010.
- [83] M. Kinzel, Q. Mulligan, and J. O. Dabiri. Energy exchange in an array of vertical-axis wind turbines. *Journal of Turbulence*, 13(38):1–13, 2012.
- [84] Wang Sang Koon, Martin W Lo, Jerrold E Marsden, and Shane D Ross. Dynamical systems, the three-body problem and space mission design. *California Institute of Technology, Pasadena, CA, USA*, 2006.
- [85] Kartik Krishna, Steven L Brunton, and Zhuoyuan Song. Finite time lyapunov exponent analysis of model predictive control and reinforcement learning. *arXiv preprint arXiv:2304.03326*, 2023.
- [86] Kartik Krishna, Aditya Nair, Anand Krishnan, Steven L Brunton, and Eurika Kaiser. Control of vortex dynamics using invariants. *arXiv preprint arXiv:2308.03920*, 2023.
- [87] Kartik Krishna, Zhuoyuan Song, and Steven L Brunton. Finite-horizon, energy-efficient trajectories in unsteady flows. *Proceedings of the Royal Society A*, 478(2258):20210255, 2022.

- [88] Dhanushka Kularatne, Subhrajit Bhattacharya, and M Ani Hsieh. Time and energy optimal path planning in general flows. In *Robotics: Science and Systems*, 2016.
- [89] Francis D Lagor, Kayo Ide, and Derek A Paley. Touring invariant-set boundaries of a two-vortex system using streamline control. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 2217–2222. IEEE, 2015.
- [90] Francis D Lagor and Derek A Paley. Active singularities for multivehicle motion planning in an N-vortex system. In *International Conference on Dynamic Data-Driven Environmental Systems Science*, pages 334–346. Springer, 2014.
- [91] Jay H Lee. Model predictive control: Review of the three decades of development. *International Journal of Control, Automation and Systems*, 9(3):415–424, 2011.
- [92] F. Lekien, C. Coulliette, A. J. Mariano, E. H. Ryan, L. K. Shay, G. Haller, and J. E. Marsden. Pollution release tied to invariant manifolds: A case study for the coast of Florida. *Physica D*, 210:1–20, 2005.
- [93] Francois Lekien, Shawn C Shadden, and Jerrold E Marsden. Lagrangian coherent structures in n-dimensional systems. *Journal of Mathematical Physics*, 48(6):065404, 2007.
- [94] Naomi Ehrich Leonard and Joshua G Graver. Model-based feedback control of autonomous underwater gliders. *IEEE Journal of Oceanic Engineering*, 26(4):633–645, 2001.
- [95] Naomi Ehrich Leonard, Derek A Paley, Francois Lekien, Rodolphe Sepulchre, David M Fratantoni, and Russ E Davis. Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE*, 95(1):48–74, 2007.
- [96] Pierre FJ Lermusiaux. Adaptive modeling, adaptive data assimilation and adaptive sampling. *Physica D: Nonlinear Phenomena*, 230(1-2):172–196, 2007.
- [97] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez,

- Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [98] D. Lipinski and K. Mohseni. A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures. *Chaos*, 20:017503, 2010.
- [99] Doug Lipinski and Kamran Mohseni. Cooperative control of a team of unmanned vehicles using smoothed particle hydrodynamics. In *AIAA Guidance, Navigation, and Control Conference*, page 8316, 2010.
- [100] Doug Lipinski and Kamran Mohseni. A master-slave fluid cooperative control algorithm for optimal trajectory planning. In *2011 IEEE International Conference on Robotics and Automation*, pages 3347–3351. IEEE, 2011.
- [101] Douglas Lipinski and Kamran Mohseni. Feasible area coverage of a hurricane using micro-aerial vehicles. In *AIAA Atmospheric Flight Mechanics Conference*, page 0894, 2014.
- [102] D. M. Luchtenburg, S. L. Brunton, and C. W. Rowley. Long-time uncertainty propagation using generalized polynomial chaos and flow map composition. *Journal of Computational Physics*, 274:783–802, 2014.
- [103] Theodore MacMillan, Nicholas T Ouellette, and David H Richter. Detection of evolving Lagrangian coherent structures: A multiple object tracking approach. *Physical Review Fluids*, 5(12):124401, 2020.
- [104] Kenneth Mallory, M Ani Hsieh, Eric Forgoston, and Ira B Schwartz. Distributed allocation of mobile sensing swarms in gyre flows. *Nonlinear Processes in Geophysics*, 20(5):657–668, 2013.
- [105] Carlo Marchioro and Mario Pulvirenti. *Mathematical theory of incompressible nonviscous fluids*, volume 96. Springer Science & Business Media, 2012.
- [106] Jerrold Marsden and Alan Weinstein. Coadjoint orbits, vortices, and clebsch variables for incompressible fluids. *Physica D: Nonlinear Phenomena*, 7(1-3):305–323, 1983.

- [107] Jerrold E. Marsden and Tudor S. Ratiu. *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems*. Springer Publishing Company, Incorporated, 2010.
- [108] M. Mathur, G. Haller, T. Peacock, J. E. Ruppert-Felsot, and H. L. Swinney. Uncovering the Lagrangian skeleton of turbulence. *Physical Review Letters*, 98:144502–1–144502–4, 2007.
- [109] David Q Mayne. Nonlinear model predictive control: An assessment. In V. J. C. Kantor, Carlos E Garcia, and B. Carnahan, editors, *Proc. Chemical Process Control*, pages 217–231. AIChE, 1997.
- [110] David Q Mayne, James B Rawlings, Christopher V Rao, and Pierre OM Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- [111] Muralikrishnan Gopalakrishnan Meena, Aditya G Nair, and Kunihiko Taira. Network community-based model reduction for vortical flows. *Physical Review E*, 97(6):063103, 2018.
- [112] Volker Mehrmann and Benjamin Unger. Control of port-hamiltonian differential-algebraic systems and applications. *arXiv preprint arXiv:2201.06590*, 2022.
- [113] Patrice Meunier, Stéphane Le Dizes, and Thomas Leweke. Physics of vortex merging. *Comptes Rendus Physique*, 6(4-5):431–450, 2005.
- [114] Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325, 2005.
- [115] Matthew Michini, M Ani Hsieh, Eric Forgoston, and Ira B Schwartz. Robotic tracking of coherent structures in flows. *IEEE Transactions on Robotics*, 30(3):593–603, 2014.
- [116] Matthew Michini, Kenneth Mallory, Dennis Larkin, M Ani Hsieh, Eric Forgoston, and Philip A Yecko. An experimental testbed for multi-robot tracking of manifolds

- and coherent structures in flows. In *Dynamic Systems and Control Conference*, volume 56130, page V002T32A002. American Society of Mechanical Engineers, 2013.
- [117] Manfred Morari and Jay H Lee. Model predictive control: past, present and future. *Computers & Chemical Engineering*, 23(4):667–682, 1999.
- [118] Aditya G Nair and Kunihiro Taira. Network-theoretic approach to sparsified discrete vortex dynamics. *Journal of Fluid Mechanics*, 768:549–571, 2015.
- [119] J. C. Nawroth, H. Lee, A. W. Feinberg, C. M. Ripplinger, M. L. McCain, A. Grosberg, J. O. Dabiri, and K. K. Parker. A tissue-engineered jellyfish with biomimetic propulsion. *Nature Biotechnology*, 30:792–797, 2012.
- [120] Paul K Newton. *The N-vortex problem: analytical techniques*, volume 145. Springer Science & Business Media, 2013.
- [121] B.R. Noack, I. Mezic, and A. Banaszuk. Controlling vortex motion and chaotic advection. In *Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No.00CH37187)*, volume 2, pages 1716–1723 vol.2, 2000.
- [122] Maria J Olascoaga, II Rypina, Michael G Brown, Francisco J Beron-Vera, Huseyin Koçak, Larry E Brand, GR Halliwell, and Lynn K Shay. Persistent transport barrier on the West Florida Shelf. *Geophysical research letters*, 33(22), 2006.
- [123] Nicholas T Ouellette, Charlie AR Hogg, and Yang Liao. Correlating Lagrangian structures with forcing in two-dimensional flow. *Physics of Fluids*, 28(1):015105, 2016.
- [124] Kathrin Padberg, Thilo Hauff, Frank Jenko, and Oliver Junge. Lagrangian structures and transport in turbulent magnetized plasmas. *New Journal of Physics*, 9:400, 2007.
- [125] J. Peng and J. O. Dabiri. The ‘upstream wake’ of swimming and flying animals and its correlation with propulsive efficiency. *The Journal of Experimental Biology*, 211:2669–2677, 2008.

- [126] J Peng and JO Dabiri. Transport of inertial particles by lagrangian coherent structures: application to predator–prey interaction in jellyfish feeding. *Journal of Fluid Mechanics*, 623:75–84, 2009.
- [127] A Pentek, JB Kadtke, and Z Toroczka. Stabilizing chaotic vortex trajectories: an example of high-dimensional control. *Physics Letters A*, 224(1-2):85–92, 1996.
- [128] Bartosz Protas. Vortex dynamics models in flow control problems. *Nonlinearity*, 21(9):R203, 2008.
- [129] S. J. Qin and T. A. Badgwell. An overview of industrial model predictive control technology. In *AIChE Symposium Series*, volume 93, pages 232–256, 1997.
- [130] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.
- [131] Saša V Raković, Eric C Kerrigan, David Q Mayne, and Konstantinos I Kouramas. Optimized robust control invariance for linear discrete-time systems: Theoretical foundations. *Automatica*, 43(5):831–841, 2007.
- [132] AG Ramos, VJ García-Garrido, AM Mancho, S Wiggins, J Coca, Scott Glenn, Oscar Schofield, Joshua Kohut, D Aragon, J Kerfoot, et al. Lagrangian coherent structure assisted path planning for transoceanic autonomous underwater vehicle missions. *Scientific Reports*, 8(1):1–9, 2018.
- [133] Dushyant Rao and Stefan B Williams. Large-scale path planning for underwater gliders in ocean currents. In *Australasian conference on robotics and automation (ACRA)*, pages 2–4, 2009.
- [134] Benjamin Recht. A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems*, 2:253–279, 2019.
- [135] Blane Rhoads, Igor Mezić, and Andrew C Poje. Minimum time heading control of

- underpowered vehicles in time-varying ocean currents. *Ocean Engineering*, 66:12–31, 2013.
- [136] Renzo Ricca. *An Introduction to the Geometry and Topology of Fluid Flows*. 01 2001.
- [137] Em Rickinson, Nick G Parker, Andrew W Baggaley, and Carlo F Barenghi. Diffusion of quantum vortices. *Physical Review A*, 98(2):023608, 2018.
- [138] Michael J Risbeck and James B Rawlings. MPCTools: Nonlinear model predictive control tools for CasADi (Python interface), 2015. <https://bitbucket.org/rawlings-group/mpc-tools-casadi>.
- [139] M. P. Rockwood, K. Taira, and M. A. Green. Detecting vortex formation and shedding in cylinder wakes using Lagrangian coherent structures. *AIAA Journal*, 55(1):15–23, 2016.
- [140] Matthew P Rockwood and Melissa A Green. Real-time identification of vortex shedding in the wake of a circular cylinder. *AIAA Journal*, 57(1):223–238, 2019.
- [141] Matthew P Rockwood, Thomas Loiseau, and Melissa A Green. Practical concerns of implementing a finite-time Lyapunov exponent analysis with under-resolved data. *Experiments in Fluids*, 60(4):74, 2019.
- [142] Steven N Rodriguez and Justin W Jaworski. Strongly-coupled aeroelastic free-vortex wake framework for floating offshore wind turbine rotors. part 1: Numerical framework. *Renewable Energy*, 141:1127–1145, 2019.
- [143] Steven N Rodriguez and Justin W Jaworski. Strongly-coupled aeroelastic free-vortex wake framework for floating offshore wind turbine rotors. part 2: Application. *Renewable Energy*, 149:1018–1031, 2020.
- [144] Johan Roenby and Hassan Aref. Chaotic dynamics of a body–vortex pair. *Journal of Fluids and Structures*, 27(5-6):768–773, 2011.
- [145] V Rom-Kedar, A Leonard, and S Wiggins. An analytical study of transport, mixing and chaos in an unsteady vortical flow. *Journal of Fluid Mechanics*, 214:347–394, 1990.

- [146] Philip G Saffman. *Vortex dynamics*. Cambridge university press, 1995.
- [147] R Salazar, A Campos, V Fuentes, and A Abdelkefi. A review on the modeling, materials, and actuators of aquatic unmanned vehicles. *Ocean Engineering*, 172:257–285, 2019.
- [148] Yasuo Sasaki and Daisuke Tsubakino. Model predictive control of a separated flow around a circular cylinder at a low reynolds number. *SICE Journal of Control, Measurement, and System Integration*, 11(3):154–159, 2018.
- [149] Shankar Sastry. *Nonlinear systems: analysis, stability, and control*, volume 10. Springer Science & Business Media, 2013.
- [150] Kristy L Schlueter-Kuck and John O Dabiri. Coherent structure colouring: Identification of coherent structures from sparse data using graph theory. *Journal of Fluid Mechanics*, 811:468–486, 2017.
- [151] Carmine Senatore and Shane D Ross. Fuel-efficient navigation in complex flows. In *2008 American Control Conference*, pages 1244–1248. IEEE, 2008.
- [152] Carmine Senatore and Shane D Ross. Detection and characterization of transport barriers in complex flows via ridge extraction of the finite time Lyapunov exponent field. *International Journal for Numerical Methods in Engineering*, 86(9):1163–1174, 2011.
- [153] Mattia Serra and George Haller. Objective Eulerian coherent structures. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 26(5):053110, 2016.
- [154] Shawn C Shadden. Lagrangian coherent structures. *Transport and Mixing in Laminar Flows: From Microfluidics to Oceanic Currents*, pages 59–89, 2011.
- [155] Shawn C Shadden, John O Dabiri, and Jerrold E Marsden. Lagrangian analysis of fluid transport in empirical vortex ring flows. *Physics of Fluids*, 18(4):047105, 2006.
- [156] Shawn C Shadden, Francois Lekien, and Jerrold E Marsden. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D: Nonlinear Phenomena*, 212(3-4):271–304, 2005.

- [157] Shawn C Shadden, Francois Lekien, Jeffrey D Paduan, Francisco P Chavez, and Jerrold E Marsden. The correlation between surface drifters and coherent structures based on high-frequency radar data in Monterey Bay. *Deep Sea Research Part II: Topical Studies in Oceanography*, 56(3-5):161–172, 2009.
- [158] Shawn C Shadden and Charles A Taylor. Characterization of coherent structures in the cardiovascular system. *Annals of Biomedical Engineering*, 36(7):1152–1162, 2008.
- [159] Banavara N Shashikanth, Jerrold E Marsden, Joel W Burdick, and Scott D Kelly. The hamiltonian structure of a two-dimensional rigid circular cylinder interacting dynamically with n point vortices. *Physics of Fluids*, 14(3):1214–1227, 2002.
- [160] Eric D Siggia and Hassan Aref. Point-vortex simulation of the inverse energy cascade in two-dimensional turbulence. *The Physics of Fluids*, 24(1):171–173, 1981.
- [161] Tapio Simula, Matthew J Davis, and Kristian Helmerson. Emergence of order from turbulence in an isolated planar superfluid. *Physical review letters*, 113(16):165302, 2014.
- [162] Audun Skaugen and Luiza Angheluta. Origin of the inverse energy cascade in two-dimensional quantum turbulence. *Physical Review E*, 95(5):052144, 2017.
- [163] Zhuoyuan Song, Doug Lipinski, and Kamran Mohseni. Multi-vehicle cooperation and nearly fuel-optimal flock guidance in strong background flows. *Ocean Engineering*, 141:388–404, 2017.
- [164] Zhuoyuan Song and Kamran Mohseni. Anisotropic active Lagrangian particle swarm control in a meandering jet. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 240–245. IEEE, 2015.
- [165] S. Speich, H. Dijkstra, and M. Ghil. Successive bifurcations in a shallow-water model applied to the wind-driven ocean circulation. *Nonlinear Processes in Geophysics*, 2(3/4):241–268, 1995.

- [166] S. Speich and M. Ghil. Interannual variability of the mid-latitude oceans: A new source of climate variability. *Sistema Terra*, 3(3):459, 1994.
- [167] Robert F Stengel. *Optimal control and estimation*. Courier Corporation, 1994.
- [168] Mark A Stremler and Hassan Aref. Motion of three point vortices in a periodic parallelogram. *Journal of Fluid Mechanics*, 392:101–128, 1999.
- [169] Mark A Stremler and Saikat Basu. On point vortex models of exotic bluff body wakes. *Fluid Dynamics Research*, 46(6):061410, 2014.
- [170] Deepak N Subramani and Pierre FJ Lermusiaux. Energy-optimal path planning by stochastic dynamically orthogonal level-set optimization. *Ocean Modelling*, 100:57–77, 2016.
- [171] M Sudharsan, Steven L Brunton, and James J Riley. Lagrangian coherent structures and inertial particle dynamics. *Physical Review E*, 93(3):033108, 2016.
- [172] Arun Vishnu Suresh Babu, Kiran Kumar Ramesh, and Ashok Gopalarathnam. Model reduction in discrete vortex methods for 2D unsteady aerodynamic flows. In *34th AIAA Applied Aerodynamics Conference*, page 4163, 2016.
- [173] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. 2018.
- [174] Kunihiko Taira, Steven L Brunton, Scott Dawson, Clarence W Rowley, Tim Colonius, Beverley J McKeon, Oliver T Schmidt, Stanislav Gordeyev, Vassilios Theofilis, and Lawrence S Ukeiley. Modal analysis of fluid flows: An overview. *AIAA Journal*, 55(12):4013–4041, 2017.
- [175] Kunihiko Taira, Aditya G Nair, and Steven L Brunton. Network structure of two-dimensional decaying isotropic turbulence. *Journal of Fluid Mechanics*, 795, 2016.
- [176] Phanindra Tallapragada and Shane D Ross. A set oriented definition of finite-time Lyapunov exponents and coherent sets. *Communications in Nonlinear Science and Numerical Simulation*, 18(5):1106–1126, 2013.

- [177] Phanindra Tallapragada, Shane D Ross, and David G Schmale III. Lagrangian coherent structures are associated with fluctuations in airborne microbial populations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 21(3):033122, 2011.
- [178] L. D. Talley, I. Rosso, I. Kamenkovich, M. R. Mazloff, J. Wang, E. Boss, A. R. Gray, K. S. Johnson, R. M. Key, S. C. Riser, N. L. Williams, and J. L. Sarmiento. Southern Ocean biogeochemical float deployment strategy, with example from the Greenwich Meridian line (GO-SHIP A12). *Journal of Geophysical Research: Oceans*, 124(1):403–431, 2019.
- [179] Emanuel Todorov et al. Optimal control theory. *Bayesian brain: probabilistic approaches to neural coding*, pages 268–298, 2006.
- [180] Dmitri Vainchtein and Igor Mezić. Control of vortex elements: Flat coordinates vs. adiabatic control. *IFAC Proceedings Volumes*, 34(6):471–474, 2001.
- [181] Dmitri Vainchtein and Igor Mezic. Control of a vortex pair using a weak external flow*. *Journal of Turbulence*, 3(51):2–7, 2002.
- [182] Dmitri Vainchtein and Igor Mezić. Optimal control of a co-rotating vortex pair: averaging and impulsive control. *Physica D: Nonlinear Phenomena*, 192(1-2):63–82, 2004.
- [183] Dmitri Vainchtein and Igor Mezić. Vortex-based control algorithms. In *Control of Fluid Flow*, pages 189–212. Springer, 2006.
- [184] Arjan Van Der Schaft, Dimitri Jeltsema, et al. Port-hamiltonian systems theory: An introductory overview. *Foundations and Trends® in Systems and Control*, 1(2-3):173–378, 2014.
- [185] Joris Vankerschaver, Eva Kanso, and Jerrold E Marsden. The geometry and dynamics of interacting rigid bodies and point vortices. *arXiv preprint arXiv:0810.1490*, 2008.
- [186] Chengjie Wang. *High-Fidelity Simulation and Low-Order Modeling of Bio-inspired Fluid Dynamics*. PhD thesis, UCLA, 2014.

- [187] Marius Wiggert, Manan Doshi, Pierre FJ Lermusiaux, and Claire J Tomlin. Navigating underactuated agents by hitchhiking forecast flows. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 2417–2424. IEEE, 2022.
- [188] Stephen Wiggins. *Chaotic transport in dynamical systems*, volume 2. Springer Science & Business Media, 2013.
- [189] Megan M Wilson, Jifeng Peng, John O Dabiri, and Jeff D Eldredge. Lagrangian coherent structures in low Reynolds number swimming. *Journal of Physics: Condensed Matter*, 21(20):204105, 2009.
- [190] Russell B. Wynn, Veerle A.I. Huvenne, Timothy P. Le Bas, Bramley J. Murton, Douglas P. Connelly, Brian J. Bett, Henry A. Ruhl, Kirsty J. Morris, Jeffrey Peakall, Daniel R. Parsons, Esther J. Sumner, Stephen E. Darby, Robert M. Dorrell, and James E. Hunt. Autonomous underwater vehicles (AUVs): Their past, present and future contributions to the advancement of marine geoscience. *Marine Geology*, 352:451–468, 2014.
- [191] Ao Xu, Hua-Lin Wu, and Heng-Dong Xi. Long-distance migration with minimal energy consumption in a thermal turbulent environment. *Phys. Rev. Fluids*, 8:023502, Feb 2023.
- [192] Namik Kemal Yilmaz, Constantinos Evangelinos, Pierre FJ Lermusiaux, and Nicholas M Patrikalakis. Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. *IEEE Journal of Oceanic Engineering*, 33(4):522–537, 2008.
- [193] Weizhong Zhang, Tamer Inanc, Sina Ober-Blobaum, and Jerrold E Marsden. Optimal trajectory generation for a glider in time-varying 2D ocean flows B-spline model. In *2008 IEEE International Conference on Robotics and Automation*, pages 1083–1088. IEEE, 2008.
- [194] Yanwu Zhang, John P Ryan, Brett W Hobson, Brian Kieft, Anna Romano, Benedetto Barone, Christina M Preston, Brent Roman, Ben-Yair Raanan, Douglas Pargett, et al.

A system of coordinated autonomous robots for Lagrangian studies of microbes in the oceanic deep chlorophyll maximum. *Science Robotics*, 6(50), 2021.

s

Appendix A

SUPPLEMENTARY DATA - CHAPTER 2

Here we present additional information that provide a more detailed analysis of the performance of MPC trajectories for various parameters.

In Figure A.1, we see the evolution of trajectories with $R/Q = 2$ and $R/Q = 3$ with a time horizon $T_H = 4$ to explain why the $R/Q = 2$ trajectory initially appears to perform worse than the $R/Q = 3$ trajectory in Figure 2.2 from the main text. In particular, it appears that the more aggressive $R/Q = 2$ agent ends up on the wrong side of the blue LCS, which forces it to take a full revolution in the right gyre before making it to the left gyre where the goal state resides. We observe this phenomena in several different parameter regimes, where small changes in the parameters may cause agents to get forced into extra orbits in the right gyre.

Figures A.2–A.10 provide similar information to Figure 2.2 in the main text, but with different time horizons. Even for a short time horizon of $T_H = 1$, the most aggressive controllers achieve relatively good state tracking performance. However, the cost versus R/Q curves for $T_H \leq 3$ are considerably less monotonic than those for $T_H \geq 4$, indicating several more bifurcations in the trajectory shape. For $T_H \in [4, 7]$, the behavior is fairly regular, exhibiting the same qualitative bifurcation behavior. Interestingly, there is a trend of bifurcations occurring later for larger T_H in this range, as the longer time-horizon controllers are able to achieve slightly better trajectories for larger R/Q values.

Finally, Figure A.11 provides the histograms of the x - and y -components of the agent velocity, complementing the data in Figure 2.8 from the main text.

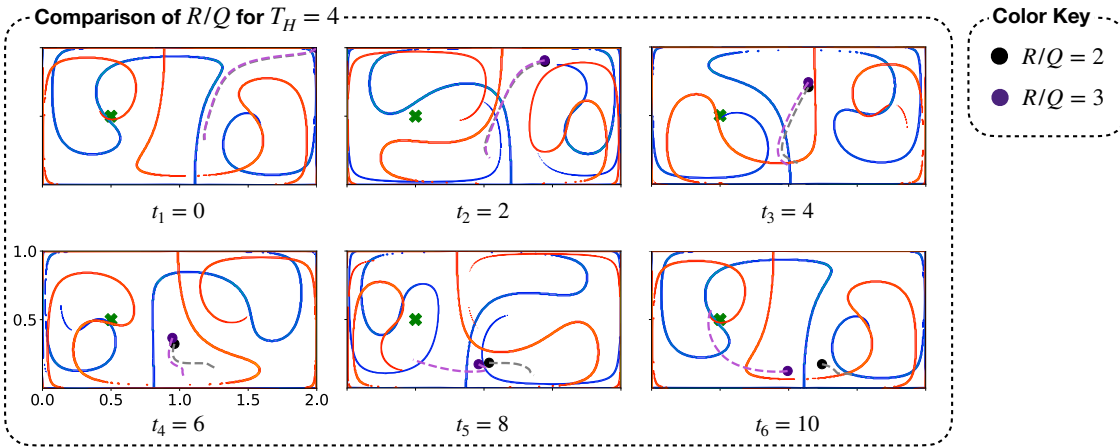


Figure A.1: In Figure 2.2, we observe that for $T_H = 4$, $R/Q = 2$, the trajectory makes a loop in the right gyre before ultimately reaching the goal. This is not the case for $R/Q = 3$ which is more greedy in spending energy. This is somewhat counter-intuitive, since we expect that spending more energy should drive the mobile sensor more rapidly to the goal. In this plot, we show the difference in how the control energy is spent comparing $R/Q = 2$ and $R/Q = 3$. We observe that $R/Q = 2$ takes longer to reach the goal because good performance is very sensitive to timing in the gyre. We see in the above figure that the aggressive control (grey trajectory) moves too far ahead the non aggressive control (purple trajectory) to make use of the gyre dynamics to move directly to the goal state.

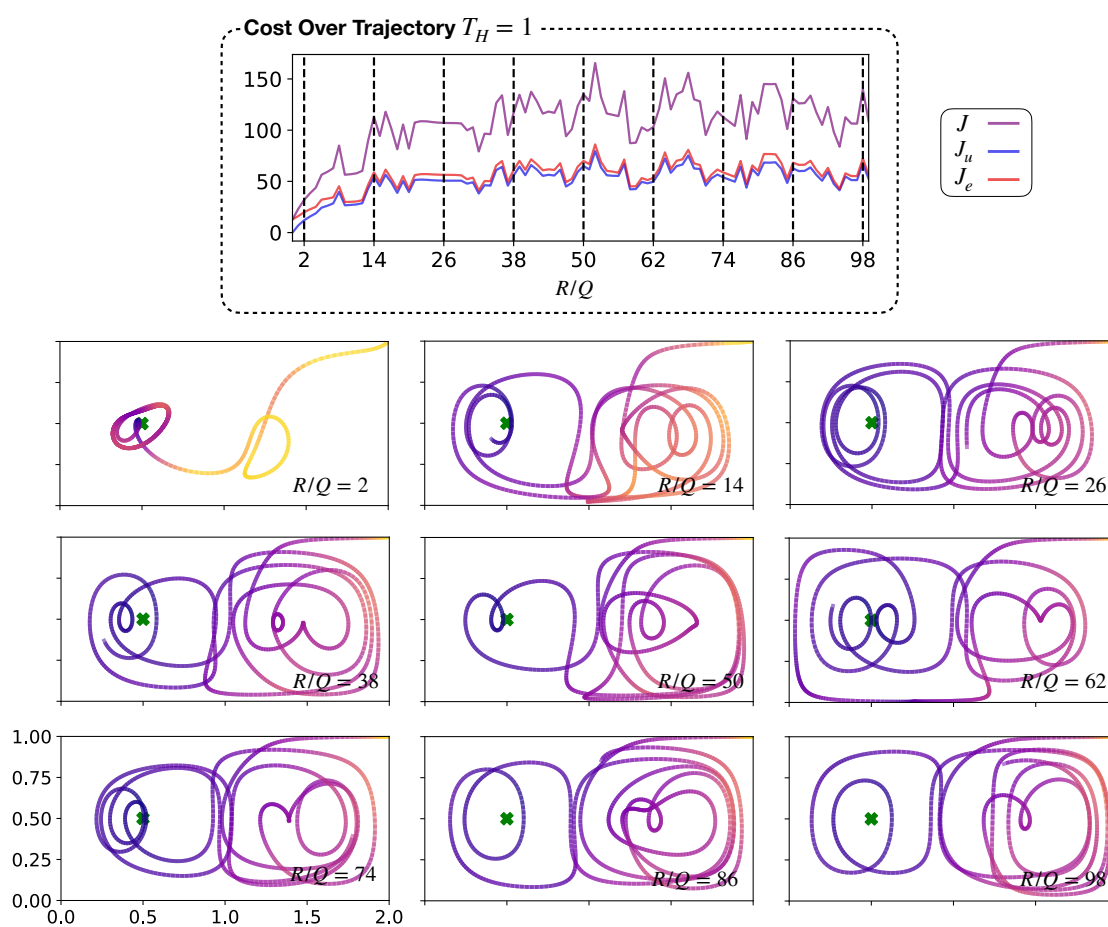


Figure A.2: Trajectories for various R/Q at a time horizon of $T_H = 1$.

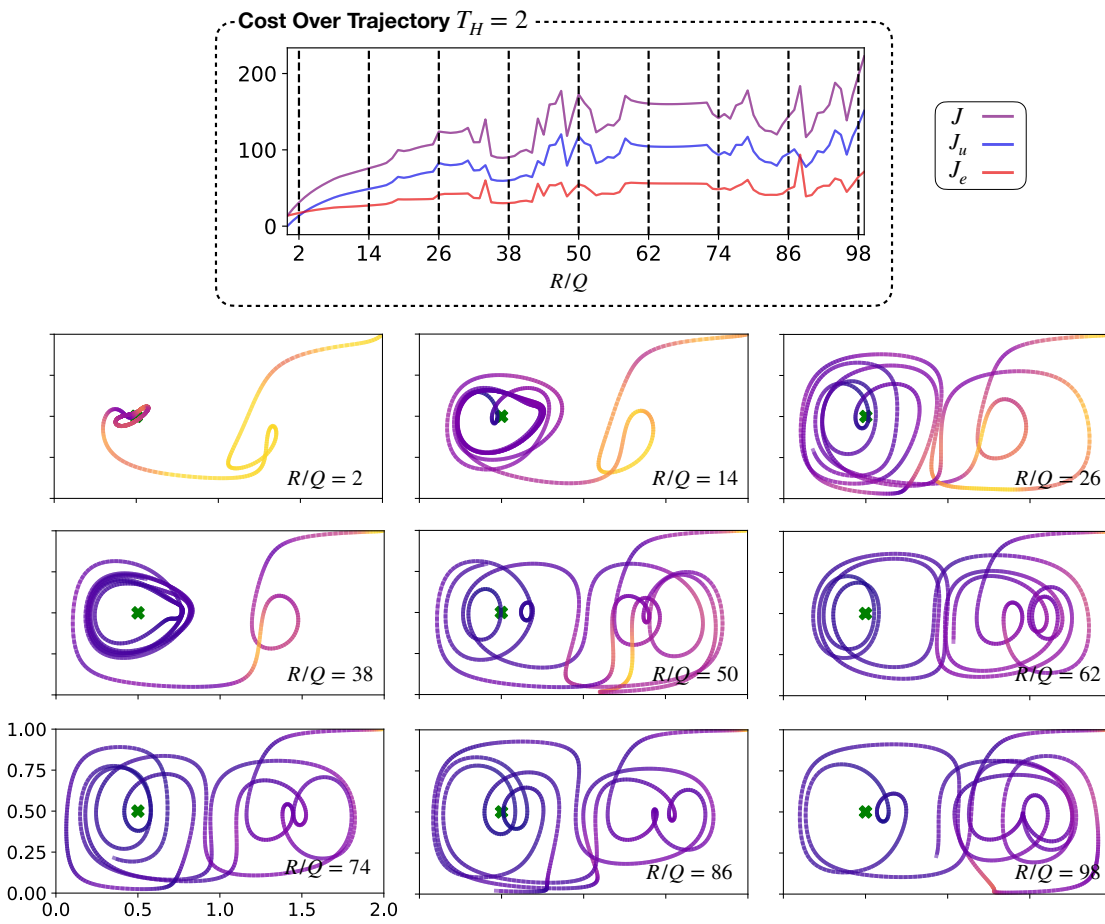


Figure A.3: Trajectories for various R/Q at a time horizon of $T_H = 2$.

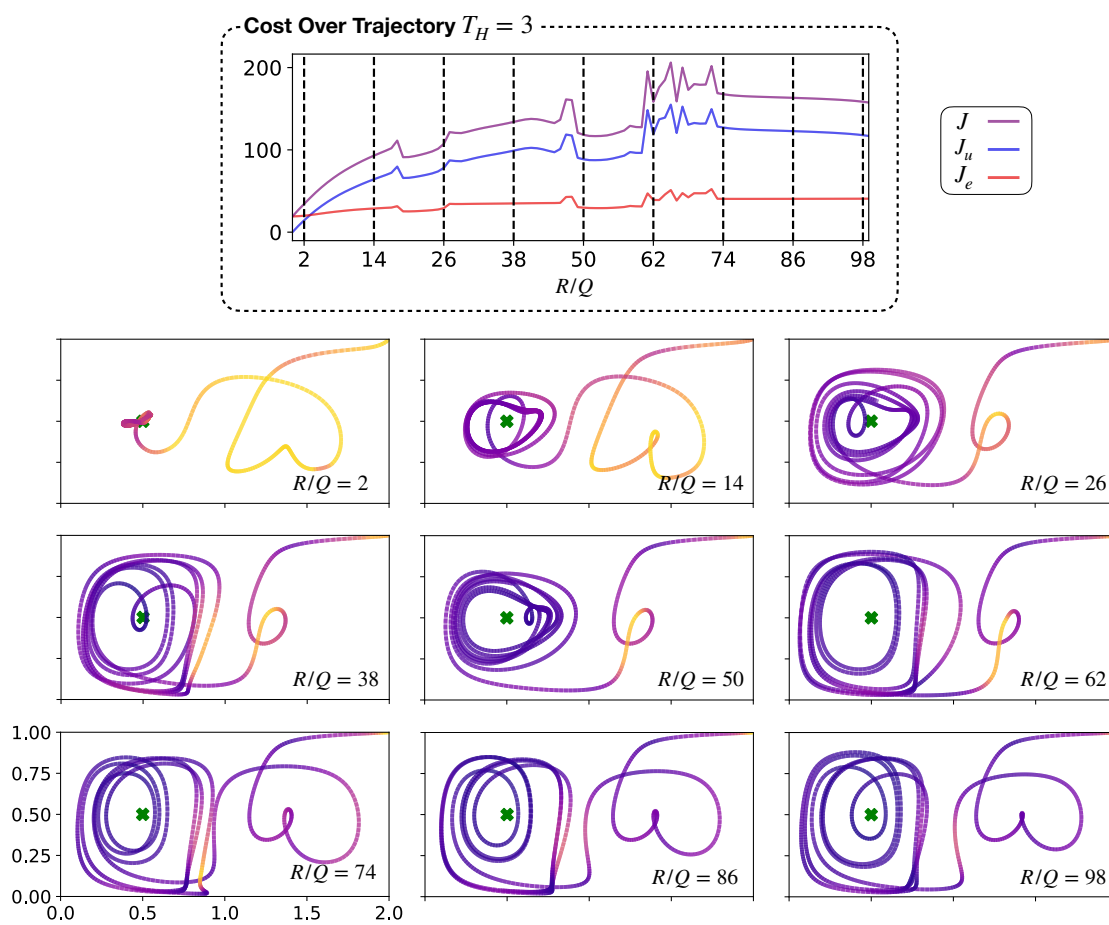
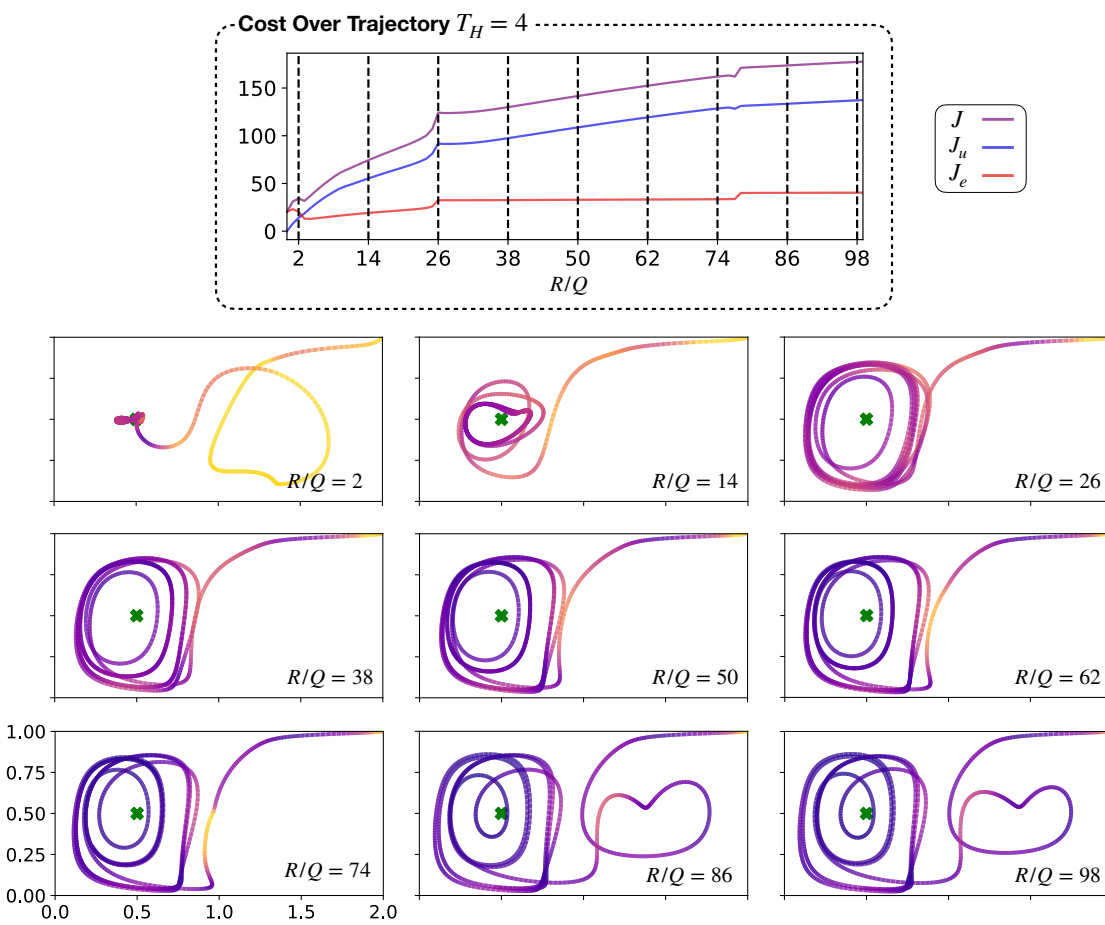
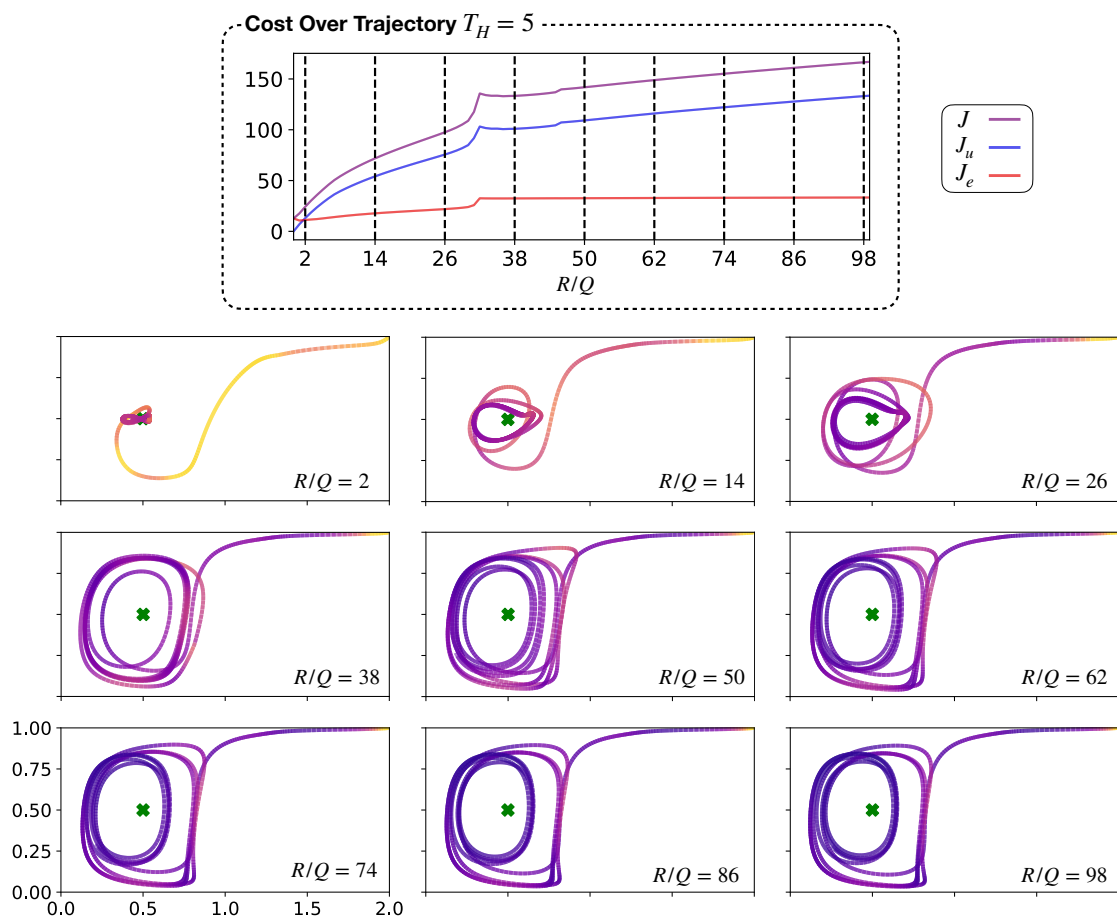
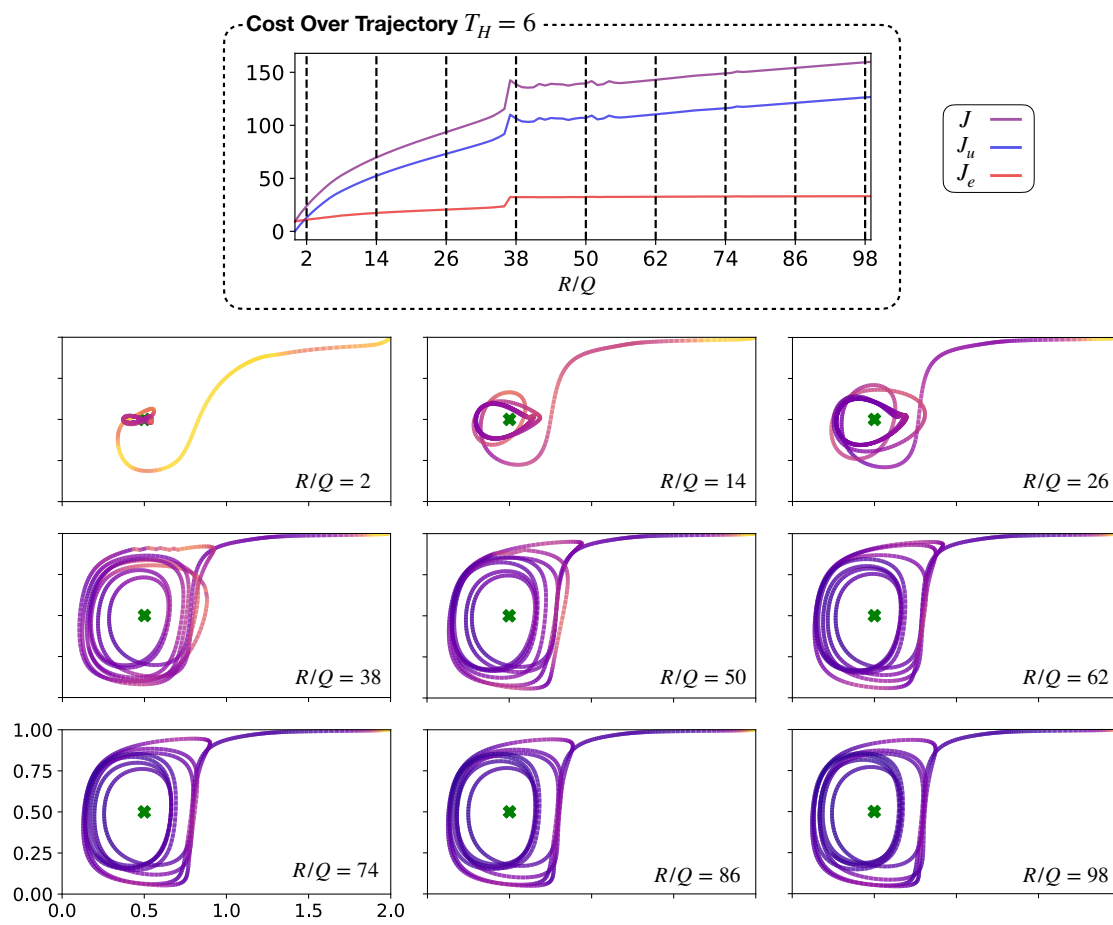
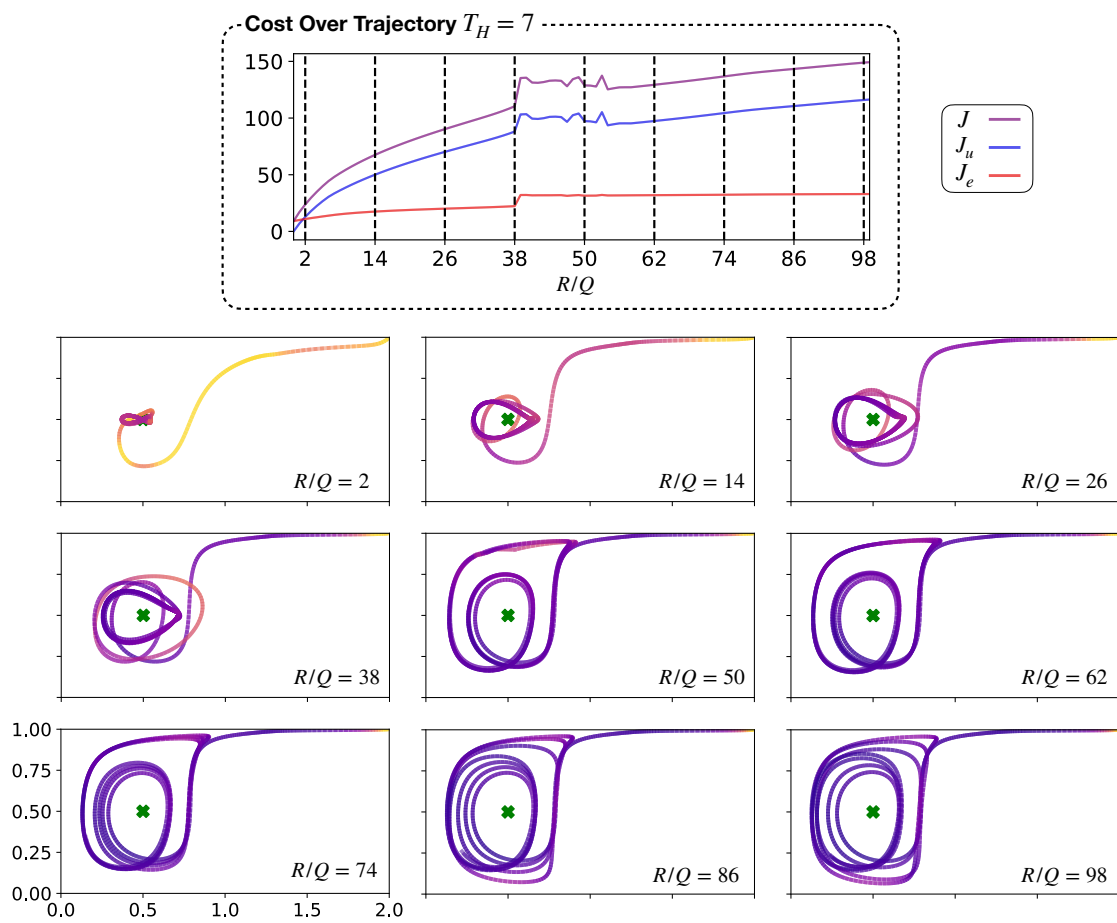


Figure A.4: Trajectories for various R/Q at a time horizon of $T_H = 3$.









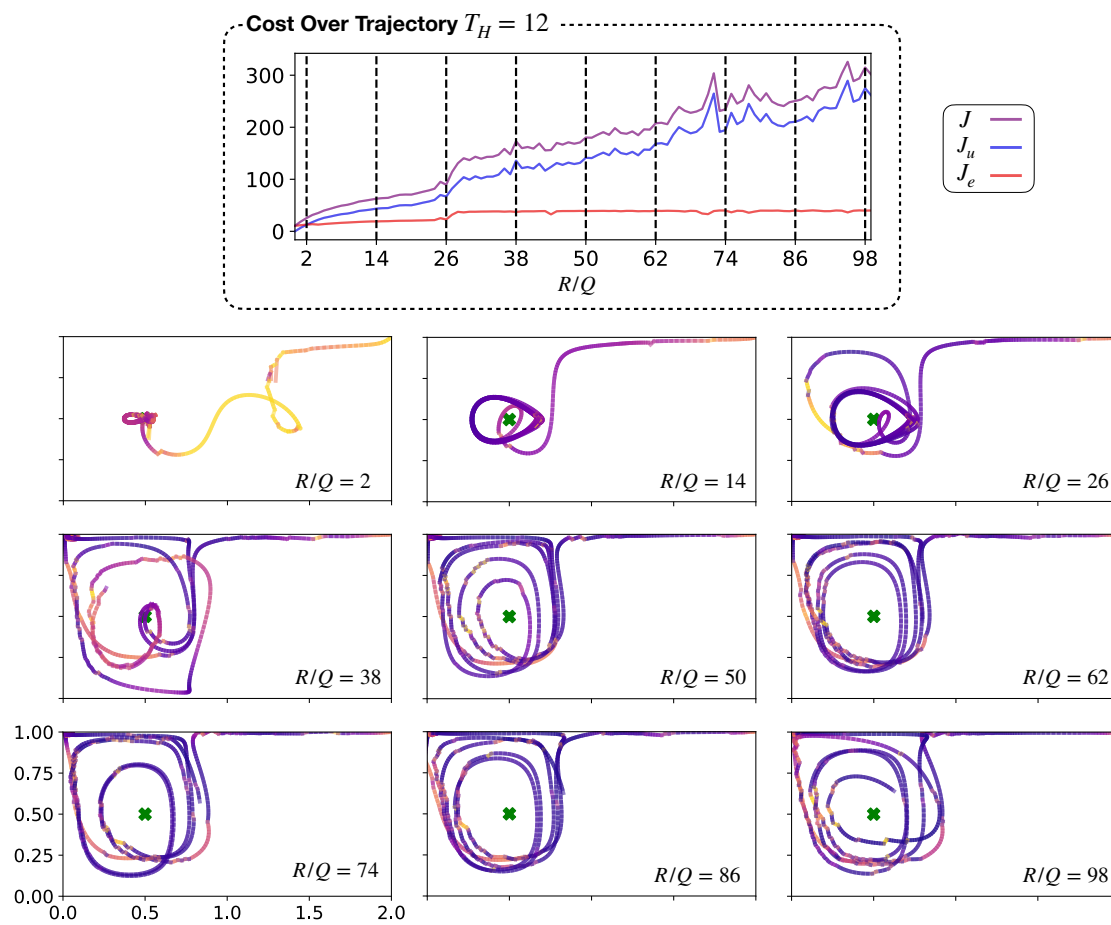


Figure A.9: Trajectories for various R/Q at a time horizon of $T_H = 12$.

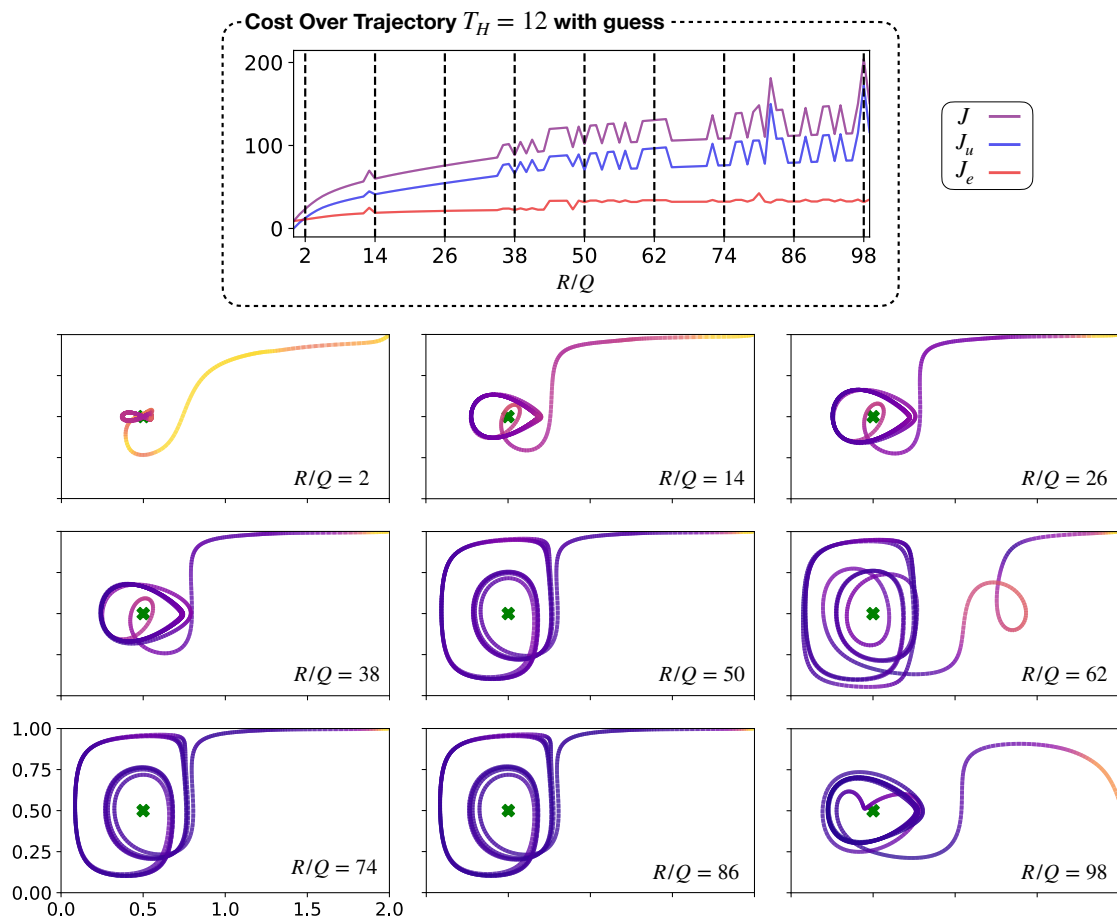


Figure A.10: Trajectories for various R/Q at a time horizon of $T_H = 12$ with warm start.

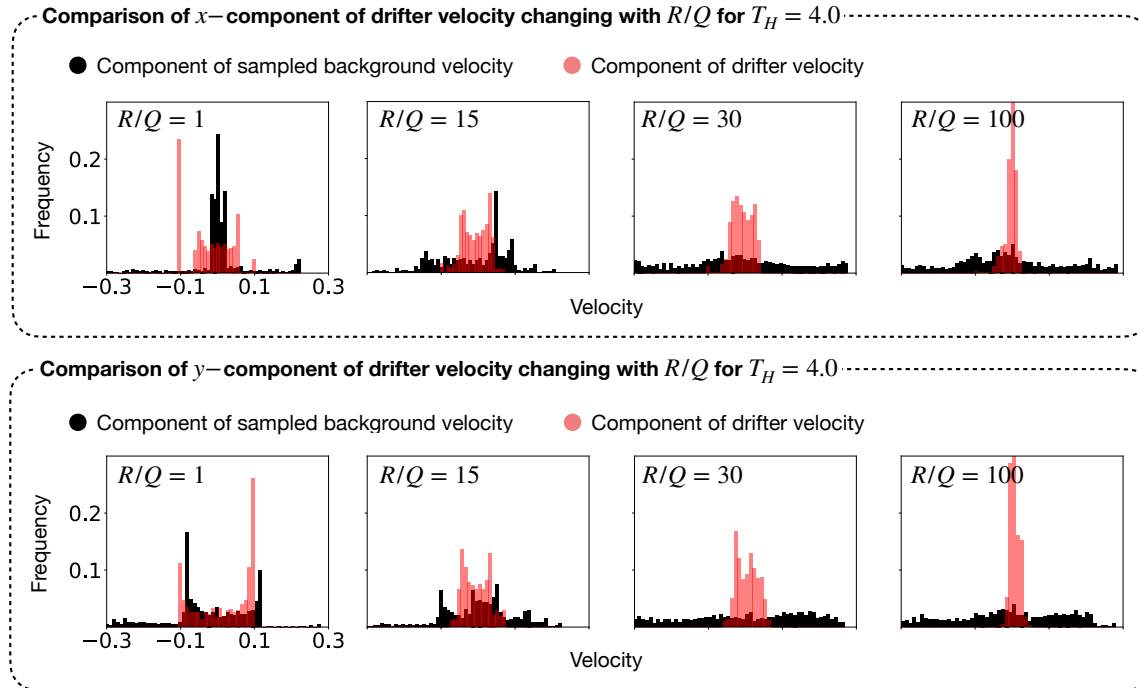


Figure A.11: As the sensor moves in the double gyre flow field, it is constantly taking control actions $u = [u_x, u_y]$, where, u_x, u_y are the x and y -components of its actuation respectively. At each instant, the sensor is also moving over a background double gyre current vector whose components v_x, v_y are given by (1.15). The top row is a histogram of the x -component of control actions u_x taken (in red), against the x -component of the background current velocity $v_x(x_s, y_s, t)$ (in black), where, x_s, y_s are the sensor coordinates at time t . The second row is a similar plot for the y -component. We observe that the gyre takes values beyond the actuation capacity of the sensor, which highlights the under-actuated nature of the problem. Also, at low R/Q ratios, the distribution of control actions follows a distribution with two peaks at ± 0.1 , which corresponds to a situation similar to bang-bang control. As we increase the R/Q ratio, the distribution of control actions move to a single peak centered around zero corresponding to the use of very little control effort when compared to the background velocity.

Appendix B

SUPPLEMENTARY DATA - CHAPTER 3

In Chapter 3, much of the discussion revolved around the repelling cFTLE. This is due to the fact that the policy for moving towards a goal is well known in forward time. To compute an attracting cFTLE, it is important to know the policy in backward time. This is easier if the flow field is periodic since the policy can simply be reversed in time sequence. However, in more realistic situations, it is challenging to have the flow field data and a control policy at times *before* the initial condition, which is generally not possible. Therefore, the attracting cFTLE can only be visualized at times $(t - t_0) > T_A$, where t_0 is the time point at which the flow field data is available.

In Figure B.1, the data from Figure 3.2 is plotted, but this time with the attracting cFTLE also visualized. The attracting cFTLE looks markedly different from the passive FTLE in Figure B.2. The double gyre system contains six fixed points counting the corners of the domain and the two on the edges. When control is applied, the unstable manifolds from each of these fixed points connect to an attractor formed around the goal by the control policy. This causes the attracting cFTLE to have multiple curves leading into the domain, taking a spiral structure. In the passive case, these unstable manifolds align with the boundaries of the domain. Depending on the aggressiveness of the control policy due to the hyperparameters, the attracting set can either be a fixed point, a limit cycle, or a chaotic set. In Figure B.1, T_A is much shorter in timescale than the time it takes for all the trajectories to fall on the attractor. This can lead to the attractor not being fully resolved when the cFTLE ridges are plotted. Being able to visualize attracting sets is useful in the context of global bifurcation analysis [60]. We observe that the attracting cFTLE forms curves in space, to which, patches latch onto in forward time.

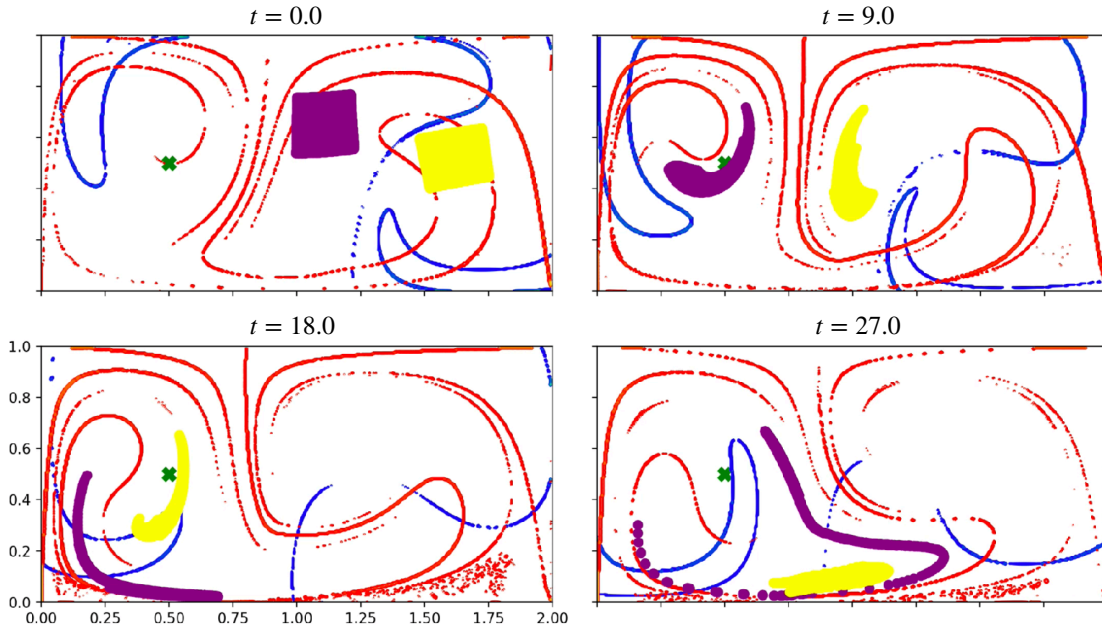


Figure B.1: This figure shows two patches of agents - one purple, and one yellow starting on opposite sides of a repelling cFTLE ridge computed with $T_A = 15$ and evolving through the unsteady flow field with control. We also plot the attracting cFTLE ridges to show the full picture of high the cFTLE ridges govern the dynamics. We see that the attracting cFTLE ridges highlight curves of long term attraction

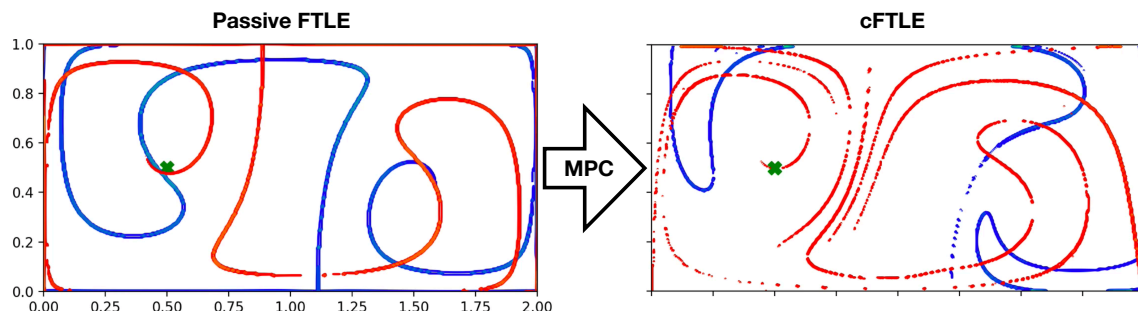


Figure B.2: We plot the passive FTLE ridges in the uncontrolled double gyre on the left, and the deformed ridges under the action of MPC at $T_H = 3.0$, $R/Q = 80$ on the right. The blue curves are repelling structures, and the red curves are attracting structures.