

THE UNIVERSITY OF CHICAGO

RARE EVENTS IN WEATHER AND CLIMATE

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF THE GEOPHYSICAL SCIENCES

BY

DAVID ARIEL PLOTKIN

CHICAGO, ILLINOIS

DECEMBER 2018

“Not fooling around, not bothering nobody, just sitting here mending the Primus and, moreover, I consider it my duty to warn you that the cat is an ancient, inviolable animal.”

– Mikhail Bulgakov, *The Master and Margarita*

TABLE OF CONTENTS

LIST OF FIGURES	v
LIST OF TABLES	vii
ACKNOWLEDGMENTS	viii
ABSTRACT	ix
1 INTRODUCTION	1
2 DISTINGUISHING MEANDERS OF THE KUROSHIO USING MACHINE LEARNING	5
2.1 Introduction	5
2.2 Methods	8
2.2.1 Data	8
2.2.2 Singular Value Decomposition	10
2.2.3 Diffusion Maps and Spectral Clustering	11
2.2.4 A Comparison of SVD and DMSC	11
2.3 Results	12
2.3.1 DSMC: States Characterized by Path Variability and Gyre Location	12
2.3.2 SVD: Gyre Location is Key	15
2.3.3 Comparison with Other Indices	17
2.3.4 Sensitivity Analysis	19
2.3.5 Comparison with AVISO SSH Data	21
2.4 Discussion	22
2.5 Conclusions	23
3 MAXIMIZING SIMULATED TROPICAL CYCLONE INTENSITY WITH ACTION MINIMIZATION	25
3.1 Introduction	25
3.2 Methods	29
3.2.1 Action minimization	29
3.2.2 WRFPLUS	35
3.2.3 Experimental design	35
3.2.4 Action minimization for the study of rapid intensification	38
3.3 Results	41
3.3.1 Comparison with maximum potential intensity	41
3.3.2 Comparison of perturbed and unperturbed storm structure	44
3.3.3 Temperature perturbations	44
3.3.4 Velocity response to heating	52
3.3.5 Water vapor perturbations	56
3.3.6 A note on energetics	59
3.3.7 Parameter sensitivity	61

3.4	Discussion	63
3.5	Conclusions	69
4	CONCLUSIONS	71
	APPENDIX DIFFUSION MAPS AND SPECTRAL CLUSTERING	74
	REFERENCES	77

LIST OF FIGURES

2.1	Kuroshio paths from the years 1956-1959 (panel [a]) and 1959-1962 (panel [b]). The paths are generated by integrating velocity fields over the course of each year. Panel [a] shows small meander years while panel [b] shows large meander years. Panel [b] suggests that the large meander may be characterized by higher variance in path location.	6
2.2	Panel [a] shows the first empirical orthogonal function of the singular value decomposition (SVD EOF 1); as expected, the dominant El Niño variability is in the warming and cooling of the eastern Pacific. Panel [b] shows that the first principal component (PC 1) found by SVD matches the ONI index. Panel [c] shows that DMSC does not identify El Niño as bimodal because there is no sharp drop-off after the first eigenvalue (normalized to $\lambda_1 = 1$).	12
2.3	Small (panel [a]) and large (panel [b]) meander states as separated by diffusion maps and spectral clustering (DMSC). Green lines are paths from every month that fall into the corresponding meander; the background velocity field is the average field in each meander. The mean paths in each state are quite similar, but the large meander includes some paths with very large deviations from the coast. The main difference in the mean velocity fields is the location of the gyre, which shifts about 1°S and 2°W during the large meander. The bottom left path in the large meander results from recirculation around the gyre.	14
2.4	The first empirical orthogonal function (EOF 1) of the yearly velocity time series. The dominant variable feature is the location of the recirculation gyre, which moves between 136.25°E , 31.25°N and 133.75°E , 30.25°N	16
2.5	Diffusion maps and spectral clustering (cyan), gyre (blue), singular value decomposition (green), path (red), and inflow (green) indices. All indices are highly correlated; the DMSC and gyre indices are nearly identical. All indices are normalized to have zero mean and equal variance.	18
2.6	Histograms of the diffusion maps and spectral clustering, gyre, and path indices. The DMSC and gyre indices are clearly bimodal; the path index is qualitatively similar to the other indices.	18
2.7	Velocity fields and paths seeded in several years of interest. 1960 (panel [a]) and 1978 (panel [b]) are large-meander years according to both the diffusion maps and spectral clustering (DMSC) and the path indices; 1967 (panel [c]) is small-meander according to both; 1970 (panel [d]) is large-meander according to the DMSC index but not the path index.	20
3.1	NOAA NHC track Hurricane Danny (Stewart, 2016).	36
3.2	NOAA NHC track Hurricane Fred (Brennan, 2009).	37
3.3	Histogram of minimum pressures for Hurricane Danny obtained by ensemble simulation (blue crosses) and the minimum pressure obtained by action minimization (red cross).	39

3.4	Minimum pressure as a function of time for forward integration (blue triangles) and action minimization output (orange circles) of Hurricanes Danny (panel a) and Fred (panel b). The green crosses and cyan stars are maximum potential intensities according to Emanuel (1991) and Holland (1997), respectively.	42
3.5	Maximum surface winds as a function of time for forward integration (blue triangles) and action minimization output (orange circles) of Hurricanes Danny (panel a) and Fred (panel b). The shaded regions mark wind speeds that must be reached for RI to occur.	43
3.6	Surface wind magnitude and OLR snapshots at final integration time step for Hurricane Danny: optimization output (panels a and c) and forward integration (panels b and d).	45
3.7	Temperature perturbations $\hat{\eta}_{\mathbf{T}}$ at 850 hPa height for action minimization of Hurricanes Danny (panels a-d) and Fred (panels e-h).	46
3.8	As in figure 3.7, but at 500 hPa instead of 850 hPa.	48
3.9	Hovmuller diagrams of 900 hPa temperature perturbations with radius and time for Hurricanes Danny and Fred. Orange triangles indicate the position of an air parcel advected by mean radial flow.	49
3.10	Hovmuller diagrams of 900 hPa temperature perturbations, averaged over $r \in [50, 150]$ km, with angle and time for Hurricanes Danny and Fred.	51
3.11	Average deviation from symmetry of the 900 hPa temperature field for Hurricane Danny at $t = 26$ hours (panel a) and $t = 36$ hours (panel b).	52
3.12	As in figure 3.4a, but with the purely symmetric of portion of action-minimizing perturbations (black triangles) and with purely asymmetric portion of the perturbations (red crosses) for Hurricane Danny.	53
3.13	Average deviation from symmetry of the azimuthal surface wind field as a function of radius for Hurricane Danny (panels a-c) and Hurricane Fred (panels d-f) at $t = 20, 28,$ and 36 hours.	54
3.14	Magnitude of 900 - 500 hPa and 850 - 250 hPa wind shear as a function of time for Hurricanes Danny (panels a and b) and Fred (panels c and d) for forward integration (blue curve) and optimization output (orange curve). Markers denote times at which perturbations are added by action minimization.	56
3.15	Initial time water vapor perturbations (panel a) and initial time relative humidity (panel b) at the surface for Hurricane Danny, as well as time series of near-storm relative humidity differences (optimized minus unperturbed) for the full near-storm domain and only the downstream portion of the domain for Hurricane Danny (panel c) and Hurricane Fred (panel d).	58
3.16	Hovmuller diagrams of 900 hPa temperature perturbations with radius and time for Hurricanes Danny at various perturbation frequencies. Orange triangles indicate the position of an air parcel advected by mean radial flow.	62
3.17	Sensitivity of intensification to the ratio of weights R_i/R_f (panel a), and sensitivity of intensification to time horizon (panel b).	64

LIST OF TABLES

2.1	Data Sources for SODA database (Carton and Giese, 2008)	9
3.1	Energy budget for simulations of Hurricane Danny ($J \times 10^{-16}$).	60

ACKNOWLEDGMENTS

Thank you to my advisors, Dorian Abbot and Jon Weare, as well as to Morgan O’Neill and Robert Webber, for their endless help in formulating the questions, developing the methods, interpreting the results, and editing the text presented in this thesis. Thank you to the members of my committee: Malte Jansen, Fred Ciesla, Noboru Nakamura, and Mihai Anitescu.

We acknowledge support from the National Science Foundation under NSF award number 1623064 and award number NSF DMS-0940261, which was part of the Mathematics and Climate Research Network. This work was supported by the Department of Energy Computational Science Graduate Fellowship Program of the Office of Science and National Nuclear Security Administration in the Department of Energy under contract DE-FG02-97ER25308. This work was completed with resources provided by the University of Chicago Research Computing Center.

ABSTRACT

Rare and/or extreme events in weather and climate often have particularly important implications for human welfare. Despite this, these events are often poorly understood since they are difficult to simulate and are, by definition, rarely observed. In this thesis, we present and apply rare event algorithms to both better characterize and better simulate specific rare events in the ocean and atmosphere systems.

First, we use diffusion maps and spectral clustering, a machine learning technique, to better characterize the large and small meanders of the Kuroshio current near Japan. This current has long been considered to be bimodal; however, there have been few data-driven efforts to confirm this bimodality or to characterize the predominant states. By applying the diffusion maps and spectral clustering algorithm in an oceanographic context for the first time, we show that the Kuroshio is indeed bimodal but that the two most common states are characterized by high and low variability rather than by current location (as was previously thought to be the case). We also show that the meanders correlate with the location of a nearby recirculation gyre, thereby providing evidence for a meander transition mechanism that depends on the movement of this gyre.

Second, we apply action minimization to the study of tropical cyclone rapid intensification by perturbing the Weather Research and Forecasting model into forming more intense storms than it otherwise would. We show that, compared to ensemble methods commonly used in the study of intensification, this method yields significant computational savings in accessing the tail of the intensification distribution. Action minimization generates maximum likelihood pathways of intensification, thereby allowing us to characterize the preferred intensification mechanisms in the model. We find that action minimization chooses physically realistic intensification mechanisms including low-level heating and the reduction of vertical wind shear. Further, we show that asymmetric heating can cause significantly more intensification than purely symmetric heating and discover a regime of non-linear storm response to asymmetric heating that has not been previously observed.

CHAPTER 1

INTRODUCTION

The study of rare and/or extreme events in climate and weather systems has long been of interest (Gretener, 1967). Such events often have disproportionately large human impacts. Rare phenomena such as flash floods, extreme heat, extreme cold, intense tropical cyclones, and droughts are some of the deadliest natural hazards in the United States (Frei and Schar, 2001; Borden and Cutter, 2008; Easterling et al., 2000). Further, climate change is postulated to increase the frequency and/or severity of rare events including droughts, extreme heat, and category 4 and 5 tropical cyclones (Katz and Brown, 1992; Emanuel, 2005; Beniston et al., 2007). This has further heightened interest in the study of extreme climate and weather events in recent years.

Rare events are difficult to study observationally due to both their infrequency and, in many cases, to the danger of obtaining observations. These events are also often difficult to simulate without incurring huge computational costs. In a modeling context, one may operationally define a rare event as one that happens infrequently by comparison to the model time step. This means one must generate very long simulations of the systems in question in order to witness particularly rare events. The problem is often compounded in ocean and atmosphere systems, which require complex models running at high enough resolutions to resolve the many disparate scales involved in these events. As a result, it is often the case that the dynamics and even the characteristics in state space of rare events are poorly understood.

Over the past thirty years, the development of the theory of large deviations has allowed for targeted study of the dynamics of rare events (Freidlin and Wentzell, 1984). Algorithms such as action minimization (E et al., 2004) allow for dynamical models to be biased into visiting rare event states more often than they otherwise would, thereby yielding computational savings in simulations of such events. Concurrently with improvements in modeling capabilities, the development of machine learning techniques including diffusion maps and

spectral clustering (DMSC) (Coifman et al., 2008) has allowed for better characterization of rare events through the use of observations. While the aforementioned algorithms are commonly used in fields such as molecular dynamics (Dellago et al., 2002) and materials science (Laio and Gervasio, 2008), they have yet to receive widespread use in weather and climate applications.

The goal of this thesis is to apply machine learning and rare event simulation methods to improve our understanding of specific rare events in weather and climate systems. We will use DMSC to show that the Kuroshio current, which undergoes rare transitions that significantly affect weather, industry, and biology along the south coast of Japan (Zhang et al., 2012), has had its dominant form of variability mischaracterized for at least half a century. We will then use action minimization to help elucidate the mechanisms responsible for causing tropical cyclones (TCs) to undergo extremely rapid intensification, and in doing so further our understanding of the role of asymmetric heating in the intensification process. Perhaps most importantly, we demonstrate the viability of rare event methods for the study of weather and climate systems, and set the stage for further applications of these and related methods in the geophysical sciences.

In Chapter 2, we apply DMSC to the Kuroshio, which is a western boundary current off the southeastern coast of Japan. The Kuroshio has long been considered to be bimodal in that it spends the vast majority of the time in one of two states, with relatively short transitions between states (Taft, 1972). These states are the small and large meanders, which have often been described as flowing near and far from the coast, respectively. Despite the significance of the Kuroshio and the wealth of literature describing both its bimodality and the predominant states, there have been no studies definitively characterizing these states or using data to show bimodality (Hasumi et al., 2010). Unlike standard tools used in the geophysical sciences such as principal component analysis (PCA) (Bretherton et al., 1992a), which seek to find dominant modes of variability, DMSC specifically seeks bimodal features of dynamical systems. By applying DMSC to reanalysis output from the Simple

Ocean Data Assimilation project (Carton and Giese, 2008), we show that the Kuroshio is indeed bimodal. However, we find that the two states are characterized by variability rather than path location: what was previously thought of as the small meander is really a state of low current variability while the large meander is a state of high current variability. DMSC provides us with an index of the state of the Kuroshio current at each time; we note that this index correlates closely with the location of a recirculation gyre south of the current. This allows for a dynamical interpretation of the results in that they support previously proposed transition mechanisms that rely on the movement of the gyre to cause switching between the small and large meanders. Finally, our index also matches the time series of coefficients produced by PCA, thereby showing that the movement of the gyre and accompanying meander transition is not only the dominant form of bimodality, but also the dominant form of variability in the system.

In Chapter 3, we apply action minimization to the study of TC rapid intensification. Using the Weather Research and Forecasting (WRF) in conjunction with its adjoint model, WRFPLUS, we have developed a code that calculates optimal small perturbations to the model in order to nudge it into forming more intense TCs. Specifically, we solve the action functional, which balances the size and realism of the perturbations against the minimum pressure at the final time of a given simulation. In doing so, we enforce physically relevant perturbations while causing two TCs, Hurricanes Danny and Fred, to clearly undergo rapid intensification. Since the intensification pathway produced by action minimization is necessarily the (locally) likeliest pathway to forming such intense storms, this allows us to learn about the preferred intensification mechanisms within the WRF model. In part, our results validate action minimization as a feasible method for this application by confirming that it chooses reasonable intensification pathways; namely it reduces shear and adds heat to the lower levels of both storms. Further, we show that asymmetric heating at low levels can cause significant intensification beyond purely symmetric heating, as had previously been argued (Nolan and Grasso, 2003; Nolan, 2007). In doing so, we find a regime of non-linear

tropical cyclone response to asymmetric heating that, to our knowledge, has not been studied in previous work. Finally, we provide evidence for the hypotheses that 1) low-level shear is more detrimental to TC intensification than high-level shear, 2) reduction of shear is significantly more impactful when shear is above a threshold of approximately 5 m s^{-1} , and 3) that there is not necessarily a minimum threshold below which reducing shear actually inhibits intensification.

Finally, in Chapter 4, we summarize the results of the previous two chapters and provide an outlook for the future use of rare event methods in the study of weather and climate.

CHAPTER 2

DISTINGUISHING MEANDERS OF THE KUROSHIO USING MACHINE LEARNING

2.1 Introduction

The Kuroshio is a western boundary current that runs along the southeastern coast of Japan, transporting heat and salinity poleward and thereby influencing large-scale weather, industry, and biology (Zhang et al., 2012). The Kuroshio south of Japan is widely claimed to exist in one of two persistent states: a small-meander state in which it does not separate from the coast and a large-meander state in which the current axis is about 2° south of the coast at about 137°E (Taft, 1972). The residence time in each state is reported as 5-10 years (Masuda, 1982; Nitani, 1975), with transitions between the two states taking several months.

The two states of the Kuroshio were noted by Taft (1972) in a study where the current axis was determined as the maximum surface velocity along North-South ship transects spaced 1° latitude apart. Representative paths were obtained by interpolating the locations of the current axis; transects where the maximum surface velocity was below 1 knot were not used. Figure 2.1 shows paths calculated from the Simple Ocean Data Assimilation (SODA) surface velocity reanalysis data set (Carton and Giese, 2008) using a similar algorithm to that used in Taft (1972). The paths show that the Kuroshio was in the small-meander state from March, 1956 to March, 1959, and that it was in the large-meander state from July, 1959 to November, 1962. While these paths provide evidence of two distinct states, there is no consensus on either a rigorous characterization of each state or on the mechanism responsible for transitions between states (Hasumi et al., 2010).

Several studies have proposed that the Kuroshio is bimodal, with the current residing almost exclusively in either the large- or small-meander state. Schmeits and Dijkstra (2001) use a high-resolution ocean general circulation model to support the hypothesis that multiple

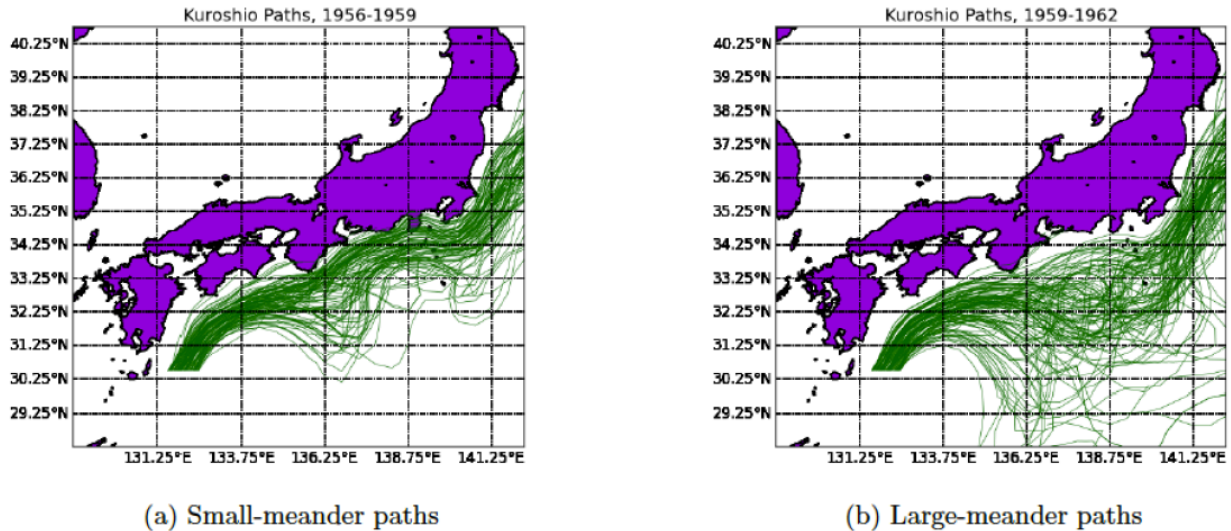


Figure 2.1: Kuroshio paths from the years 1956-1959 (panel [a]) and 1959-1962 (panel [b]). The paths are generated by integrating velocity fields over the course of each year. Panel [a] shows small meander years while panel [b] shows large meander years. Panel [b] suggests that the large meander may be characterized by higher variance in path location.

mean paths of the Kuroshio are possible. Qiu and Miao (2000) use a two-layer primitive-equation model to propose that Kuroshio bimodality is the result of a self-sustained internal oscillation of the recirculation gyre south of Japan. Pierini (2006); Pierini et al. (2009) find that the downstream Kuroshio Extension is an important bimodal feature of the system.

Further, there are many proposed mechanisms for the transition between the small and large meanders. Several modeling studies conclude that the large meander state can exist only when the Kuroshio volume transport is high (White and McCreary, 1976; Chao, 1984; Akitomo et al., 1991; Sekine, 1990) and that the transition occurs at a critical level of inflow. Other models implicate the formation of eddies as driving the transition between states (Ebuchi and Hanawa, 2003; Maltrud and McClean, 2005; Nakano and Hasumi, 2005; Hurlburt et al., 1996). A third set of models relies on wind forcing to trigger the transition between meanders (Kurogi and Akitomo, 2006; Seager et al., 2001). Qiu and Miao (2000) find an internal oscillation based on the strengthening and subsequent destabilization of the recirculation gyre south of the current.

There are likewise a number of potential indices for the state of the Kuroshio. Based

on the volume transport criterion, an index can be constructed from the inflow into the current. Kawabe (1980) used the difference in sea surface height between two points as a proxy for the state of the Kuroshio. An index of the offshore distance of the Kuroshio was presented by Qiu and Miao (2000), who used the location of the 16°C isotherm at a depth of 200 m as the Kuroshio current axis. Further, Kawabe (1995) used sea surface height data to show that there are at least three possible paths of the Kuroshio (nearshore non-large meander, offshore non-large meander, and large meander). As we show in Section 2.3, the existence of many possible paths is consistent with a fundamentally bimodal system because path shape does not define statistical modes; instead, multiple paths can fall into the same mode.

Previous data analyses argue for the existence of multiple states for the Kuroshio, and thus support the claim that the Kuroshio is highly variable. However, these studies do not address bimodality in the sense of two states of relatively high probability separated by states of relatively low probability. According to this definition of bimodality, a system can have two distinct states without being bimodal if it usually resides in between the states. An example of this is the El Niño Southern Oscillation (ENSO), which has two distinct end-member states (El Niño and La Niña) but reaches these extremes only rarely. Previous studies of the Kuroshio have produced indices that show high variability between two accepted states, but these indices do not exhibit bimodal probability distributions. In this study, we show rigorous bimodality by testing separately for bimodality and variability.

Further, the progress we make in characterizing the states of the Kuroshio is crucial for better evaluation of the proposed models and mechanisms. Due to the difficulty in separating the two meanders using simple physical measures, we use two distinct machine learning algorithms to determine the two states and confirm the bimodal behavior of the Kuroshio. The first is singular value decomposition (SVD), which we use to characterize the dominant aseasonal mode of variability present in the Kuroshio system.

The second algorithm is diffusion maps and spectral clustering (DMSC), a method that

reduces the dimensionality of the velocity field time series by treating the data as a random walk. The advantages of DMSC are two-fold: first, it allows us to specifically check for bimodality of the Kuroshio. Second, if bimodality is found, DMSC automatically identifies the modes and returns an index of the state of the system. Using SVD and DMSC in tandem allows us to establish both the dominant mode of variability of the Kuroshio, as well as its dominant bimodal feature. Agreement between SVD and DMSC results indicates that bimodality is not only present, but that the variability associated with this bimodality is the dominant form of variability in the system dynamics.

The remainder of this paper is organized as follows: in Section 2.2, we describe DMSC and SVD and discuss the advantages of each algorithm. In Section 2.3, we apply the algorithms to the Kuroshio system and show that they successfully find two distinct states. In Section 2.4, we give a physical interpretation of the results. In Section 2.5 we conclude and in the appendix we give a detailed description of the DMSC algorithm.

2.2 Methods

2.2.1 Data

We use surface velocity fields from January, 1950 to December, 2008 from the SODA database (Carton and Giese, 2008) as the input for both SVD and DMSC. The database contains monthly reanalysis velocity fields on a $.5^\circ$ (longitude and latitude) mesh. The reanalysis uses raw data sourced from various satellite, drifter, and bathythermograph observations. These are summarized in Table 2.1.

The domain we use is $131.25^\circ\text{E} - 140.25^\circ\text{E}$ and $29.25^\circ\text{N} - 35.25^\circ\text{N}$, although results presented in Section 2.3.4 show that the algorithm is robust to the choice of domain. The velocity field at time t is a point $y_t \in R^{234}$. To construct the input used for both algorithms, we divide the reanalysis time series into individual years and concatenate all of the fields within each year to create 58 points $x_t \in R^{234 \times 12}$. This approach has two advantages over

Table 2.1: Data Sources for SODA database (Carton and Giese, 2008)

Variable	Name	Dates	Description
Ocean Temp	Mechanical bathythermograph	< 1968	Temperature sensors lowered by winches to a depth of 285 m. The number of observations is under 5000 yr^{-1} .
Ocean Temp	Expendable bathythermograph	> 1968	Probes dropped from “ships of opportunity,” measure the thermal profile of the top kilometer of the ocean.
Ocean Temp, Salinity	Argo floats	> 2000	Floats that drift at a depth of 1000 m for 9 days, then descend to 2000 m and get temperature and salinity profiles as they rise to the surface.
Sea Surface Temp (SST)	AVHRR/2 AVHRR/3	> 1981	Satellites, respectively using 5 and 6-channel radiometers to report SST at daily $1^\circ \times 1^\circ$ resolution. Channels used for SST are 3B, 5, and 6, with bands at 3.55 - 3.93, 10.30 - 11.30, and 11.50-12.50 μm , respectively.
Sea Surface Height (SSH)	Topex/Poseidon Jason 1	> 1991	Satellites that get SSH at 10-day resolution, using C- and Ku-band radio altimeters.
Surface Velocity, SST	Global Drifter Program	> 1979	Drifters spaced approximately 500 km from each other that get SST and velocity data.

using individual monthly velocity fields: 1) it prevents the algorithms from simply identifying seasonal modes and 2) it encodes dynamical information in the data. The qualitative features of the results presented in Section 2.3 still hold when the algorithms are used with monthly input, but the seasonal cycle does introduce some noise that is eliminated by the use of yearly data.

To ensure that the results in Sections 2.3.1 and 2.3.2 are not artifacts of the model used in the SODA reanalysis, we check them in Section 2.3.5 using AVISO satellite sea surface height (SSH) data. The AVISO data set contains dynamic topography (SSH with respect to the geoid) in ten-day intervals and a $1^\circ \times 1^\circ$ grid for the years 1993-2009 (NCAR, 2013).

2.2.2 Singular Value Decomposition

Singular value decomposition is a standard method used in oceanography and other fields that seeks to explain the variance of a high-dimensional time series $\{\mathbf{x}\}_{t=1}^N$ by projecting each point onto a set of spatially orthonormal patterns (Bretherton et al., 1992a). These patterns are called empirical orthogonal functions (EOFs); the EOFs are in one-to-one correspondence with Principal Components (PCs) (Hannachi, 2004). Each principal component is a one-dimensional time series of projection coefficients for the corresponding EOF. The EOFs and PCs are ordered in descending order of the magnitude of variance each mode explains.

Importantly, even though SVD projects the original time series onto spatial modes, none of the modes need be independently observed. Each EOF represents a variable feature of the system; the state of the system at any time t is a linear combination of all of the variable features. The original time series can therefore be reconstructed via

$$\mathbf{x}_i = \sum_{i=1}^M \mathbf{EOF}_i \times PC_i. \quad (2.1)$$

2.2.3 Diffusion Maps and Spectral Clustering

The diffusion maps and spectral clustering algorithm (DMSC), introduced by Belkin and Niyogi (2003) and further developed by Coifman et al. (2008), seeks to reduce the dimension of a high-dimensional time series by finding multimodal features of the underlying process. A more detailed description of DMSC can be found in the appendix. DMSC constructs a random walk in phase space among the points of the original time series $\{\mathbf{x}\}_{t=1}^N$. The transition probability between two points \mathbf{x}_i and \mathbf{x}_j falls off exponentially with the Euclidean distance between them as:

$$p_{ij} \propto e^{-\frac{|\mathbf{x}_i - \mathbf{x}_j|^2}{2\epsilon_i \epsilon_j}}. \quad (2.2)$$

The exponential decay of transition probability with distance ensures that the random walk is likely to persist on sets of points with small pairwise separations. If the process that produced the time series is multimodal, these sets correspond to individual modes and determine the number of large eigenvalues of the random walk transition matrix P . Thus, a large drop-off (spectral gap) in the magnitude of eigenvalues from λ_i to λ_{i+1} indicates that the underlying process is likely multimodal, while the lack of a spectral gap suggests a lack of multimodality.

Specifically, we expect a bimodal system to have one large eigenvalue so that $\lambda_1 \gg \lambda_2$ (we ignore the first eigenvalue λ_0 since it is always true that $\lambda_0 = 1$). In this case, the first eigenvector can be used as a one-dimensional index for the state of the system. For an application of DMSC in the geophysical sciences, see Giannakis and Majda (2012).

2.2.4 A Comparison of SVD and DMSC

To illustrate the differences between SVD and DMSC, we apply both algorithms to El Niño Southern Oscillation (ENSO). The sea surface temperature (SST) in the eastern equatorial Pacific is highly variable: large positive anomalies in this region occur during El Niño and large negative anomalies occur during La Niña. The Oceanic Niño Index (ONI: a 3-month




Figure 2.2: Panel [a] shows the first empirical orthogonal function of the singular value decomposition (SVD EOF 1); as expected, the dominant El Niño variability is in the warming and cooling of the eastern Pacific. Panel [b] shows that the first principal component (PC 1) found by SVD matches the ONI index. Panel [c] shows that DMSC does not identify El Niño as bimodal because there is no sharp drop-off after the first eigenvalue (normalized to $\lambda_1 = 1$).

running mean SST anomaly in the region $5^\circ\text{S} - 5^\circ\text{N}$, $120^\circ - 170^\circ\text{W}$) describes the oscillation between these states (Trenberth, 1997). Importantly, strong El Niño or La Niña states are end members of the smooth ONI time series. Because this time series is not bimodal, we expect SVD to find ENSO variability and DMSC not to identify bimodal behavior.

Figure 2.2 summarizes the results of both algorithms when tested on a time series of equatorial Pacific temperature fields from the SODA database. EOF1 (Fig. 2.2a) shows a positive temperature anomaly in the eastern Pacific; we therefore expect PC_1 to correlate strongly with ONI. Fig. 2.2b shows that this is indeed the case, so we conclude that SVD successfully finds ENSO variability. In contrast, Fig. 2.2c shows that the DMSC eigenvalues do not exhibit a spectral gap, reflecting the fact that ENSO is not bimodal.

2.3 Results

2.3.1 DSMC: States Characterized by Path Variability and Gyre Location

Applying DMSC with constant $\epsilon_i = 3 \text{ m s}^{-1} \forall i$ to the time series of yearly velocity fields yields a spectral gap between λ_1 and λ_2 with $\lambda_2 = .42\lambda_1$. This suggests bimodal behavior. We thus consider a new, one-dimensional time series comprising the entries $\mathbf{v}_{1,i}$ of the first

non-stationary eigenvector, \mathbf{v}_1 , of the transition matrix P . To find the persistent features that correspond to this eigenvector, we then run k-means 2 clustering (Xu and Wunsch, 2005) on the entries of \mathbf{v}_1 , which yields a threshold entry value H . The association rule is then


$$\mathbf{x}_i \in \begin{cases} C_0 & : v_i \leq H \\ C_1 & : v_i > H \end{cases} \quad (2.3)$$

where C_0 and C_1 are the two clusters.


To visualize the difference between the two clusters, we seed paths in the southwestern region of the Kuroshio; this is the portion of the domain where the current is most localized. Each path is a streamline emanating from 31°N and at a random longitude between 131.4°E and 132.4°E ; this choice of longitude corresponds to the strongest region of the Kuroshio near the southwestern coast of Japan and excludes most paths that “escape” into the recirculation gyre centered at 136°E , 30.75°N . We first seed a random number of paths, typically $n \in \{0, \dots, 8\}$ paths in each monthly velocity field in each year that falls into C_0 ; we then repeat this in a separate plot for all years that fall into C_1 . The two clusters, with paths seeded in each and the average velocity fields in the background, are given in Figure 2.3. The seeded paths are used only for visualization and are not part of the DMSC algorithm.

Figure 2.3 suggests that C_1 corresponds to the small meander state and that C_0 corresponds to the large meander state. The time series of entries of the first eigenvector \mathbf{v}_1 is used to assign years to C_0 and C_1 ; each entry can thus be viewed as an index of the state of the system during the corresponding year. We call this the diffusion maps and spectral clustering (DMSC) index. The analysis in Section 2.3.3 suggests that the DMSC index is bimodal; further, the switching between clusters occurs on an approximately decadal timescale.

Figure 2.3 also shows that the small and large meanders, respectively, do not correspond to states where paths either always remain close to the coast or always meander southward around 136°E ; this is in contrast with previous work (e.g. Yoshida, K (1961); Maximenko (2002)). Instead, the small meander state contains only paths that remain near the coast,



(a) Average field in the small meander



(b) Average field in the large meander

Figure 2.3: Small (panel [a]) and large (panel [b]) meander states as separated by diffusion maps and spectral clustering (DMSC). Green lines are paths from every month that fall into the corresponding meander; the background velocity field is the average field in each meander. The mean paths in each state are quite similar, but the large meander includes some paths with very large deviations from the coast. The main difference in the mean velocity fields is the location of the gyre, which shifts about 1°S and 2°W during the large meander. The bottom left path in the large meander results from recirculation around the gyre.

while the large meander state contains both paths that remain near the coast and those that exhibit traditional large-meander behavior. In fact, most paths seeded in the mean of C_0 - that is, the average velocity field in the large-meander cluster - remain near the coast and do not look like large-meander paths.

A striking difference between the average velocity fields in the two clusters is the location of the recirculation gyre, which is located at $\sim (136.25^\circ\text{E}, 31.25^\circ\text{N})$ in the small meander and shifts $\sim 1^\circ\text{S}$ and $\sim 2^\circ\text{W}$ during the large meander. The importance of the location of the recirculation gyre is confirmed by the sensitivity analysis in Section 2.3.4: in order to identify the states shown in Figure 2.3, the algorithm domain must include either the gyre or the region immediately downstream.

2.3.2 SVD: Gyre Location is Key

Applying SVD to the yearly velocity data shows that the location of the recirculation gyre is the dominant form of variability in the system. The first EOF (Figure 2.4) shows an anticyclonic gyre centered at $136.25^\circ\text{E}, 31.25^\circ\text{N}$ and a cyclonic gyre centered at $133.75^\circ\text{E}, 30.25^\circ\text{N}$. These actually correspond to different locations of the recirculation gyre in different years: when the first principal component is positive, the gyre is nearer to $136.25^\circ\text{E}, 31.25^\circ\text{N}$. When the first principal component is negative, the gyre is nearer to $133.75^\circ\text{E}, 30.25^\circ\text{N}$. We propose that these states correspond to the small and large meanders, respectively, as described by DMSC.

To test the notion that DMSC and SVD find the same phenomenon, we compare the DMSC index with the first principal component PC_1 . The two indices are very highly correlated, with a Pearson product-moment correlation coefficient of $r = .98$ (Rodgers and Nicewander, 1988a). These indices, along with others discussed in Section 2.3.3, are shown in Figure 2.5. The near-perfect alignment of the DMSC index and PC_1 shows that SVD and DMSC find the same feature of the Kuroshio system; we therefore conclude that the changing location of the gyre is both the dominant bimodal feature and the dominant variable

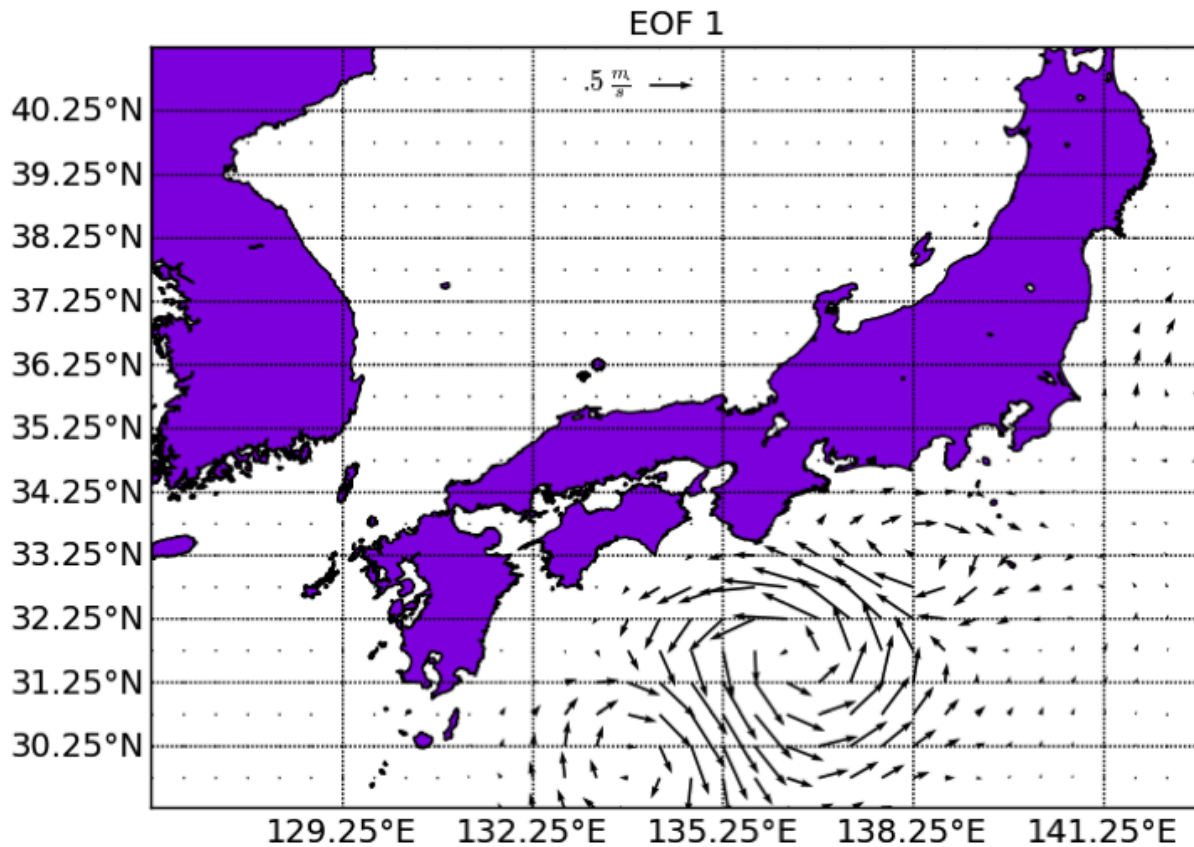


Figure 2.4: The first empirical orthogonal function (EOF 1) of the yearly velocity time series. The dominant variable feature is the location of the recirculation gyre, which moves between 136.25°E , 31.25°N and 133.75°E , 30.25°N .

feature in the system. This suggests that the bimodality of the gyre is a dominant dynamical feature of the system; the implications of this result for the meander transition mechanism are discussed in Section 2.4.

2.3.3 *Comparison with Other Indices*

To further test the validity of the diffusion maps results, we compare the DMSC index to several physically intuitive indices. We use a path index, which is constructed from the average distance from the coast of the farthest 10% of paths. We also construct an index based on the gyre location (gyre index). This is done by calculating the displacement of the gyre center from its small-meander location along a line parallel to the coast of Japan. All indices are then normalized to have zero mean and equal variance.

Figure 2.5 shows a comparison of the DMSC (blue), gyre (blue), SVD (purple), path (red), and inflow (green) indices. There is close agreement among the indices; the gyre and DMSC indices are most closely matched with a correlation coefficient of $r = .93$. The paths in Figure 2.3, as well as the correlation of the path index and the DMSC index suggest the importance of the gyre in defining the state of the Kuroshio. The DMSC index is also well-correlated with both the inflow ($r = .65$) and path ($r = .68$) indices, which suggests that the current and the recirculation gyre are strong during the small meander and weak during the large meander. This agrees with (Qiu and Miao, 2000), who present a model of the Kuroshio system in which the strengthening, migration, and weakening of the gyre is responsible for the transition between meanders.

Figure 2.6 shows histograms of the DMSC index (blue), the gyre index (green), and the path index (red). As expected, the DMSC and gyre indices agree well and show strong bimodality, with the majority of each time series being spent in either the small or large meander; the path index is qualitatively similar and show slightly weaker bimodality. The histograms show a slight preference for the small meander, with 54% of time spent nearer to this state than to the large-meander state.

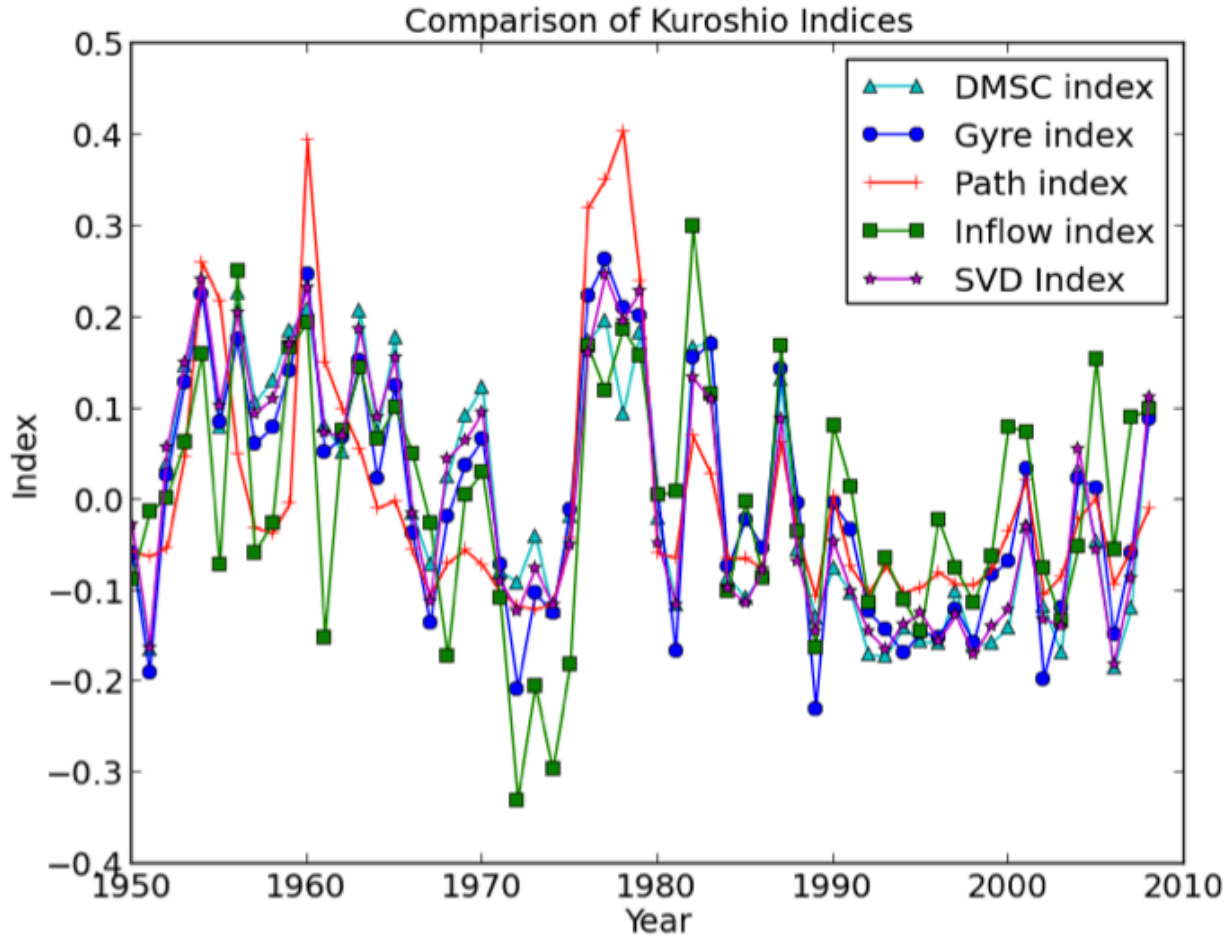


Figure 2.5: Diffusion maps and spectral clustering (cyan), gyre (blue), singular value decomposition (green), path (red), and inflow (green) indices. All indices are highly correlated; the DMSC and gyre indices are nearly identical. All indices are normalized to have zero mean and equal variance.




Figure 2.6: Histograms of the diffusion maps and spectral clustering, gyre, and path indices. The DMSC and gyre indices are clearly bimodal; the path index is qualitatively similar to the other indices.

The major discrepancy between the DMSC and the path index is that DMSC shows excursions into the large meander in 1958, 1970, and 1983. The top panels in Figure 2.7 show the velocity fields with seeded paths in January of 1960 and 1978. The bottom left panel shows paths in January of 1967, which is a strongly small-meander year according to both indices. The bottom right panel shows paths seeded in the velocity field of June, 1970, a year classified as large-meander by the DMSC index but small-meander by the path index. As with other such years, the center of the recirculation gyre ($135^\circ\text{E}, 30.25^\circ\text{N}$) is about 2° to the west and 1° to the south of where it is during small meander years.

The proposed importance of the recirculation gyre also explains the discrepancy between the DMSC and path indices in 1970. As the bottom right panel of Figure 2.7 shows, the recirculation gyre is approximately 2° to the west and 1° to the south of where it is during small meander years; however, there is an anomalous second gyre centered at $139^\circ\text{E}, 31^\circ\text{N}$. This eastern gyre pushes paths back towards the coast and prevents the vast majority of paths from forming a large-meander pattern.

2.3.4 Sensitivity Analysis

We test the sensitivity of the diffusion maps and spectral clustering algorithm to domain size by simultaneously varying the minimum and maximum latitude, lat_{\min} and lat_{\max} , as well as the minimum and maximum longitude, long_{\min} and long_{\max} of the domain in which we perform the analysis. For each domain size, we compute the correlation coefficient of the DMSC index with the new domain and the DMSC index with the original domain.

For any fixed values of $\text{long}_{\min} \leq 133.5^\circ\text{E}$, $\text{lat}_{\min} \leq 30^\circ\text{N}$, the correlation coefficient between the DMSC indices for the old and new domains is $r > .9$ as long as ($\text{long}_{\max} \geq 137^\circ\text{E}$ and $\text{lat}_{\max} \geq 32^\circ\text{N}$). This is consistent with our previous discussion in Section 2.3.1 of the importance of both the recirculation gyre and the region immediately downstream of the gyre (where the current is localized in the small meander and highly variable in the large meander). The algorithm breaks down if the boundaries are chosen such this region is excluded in many

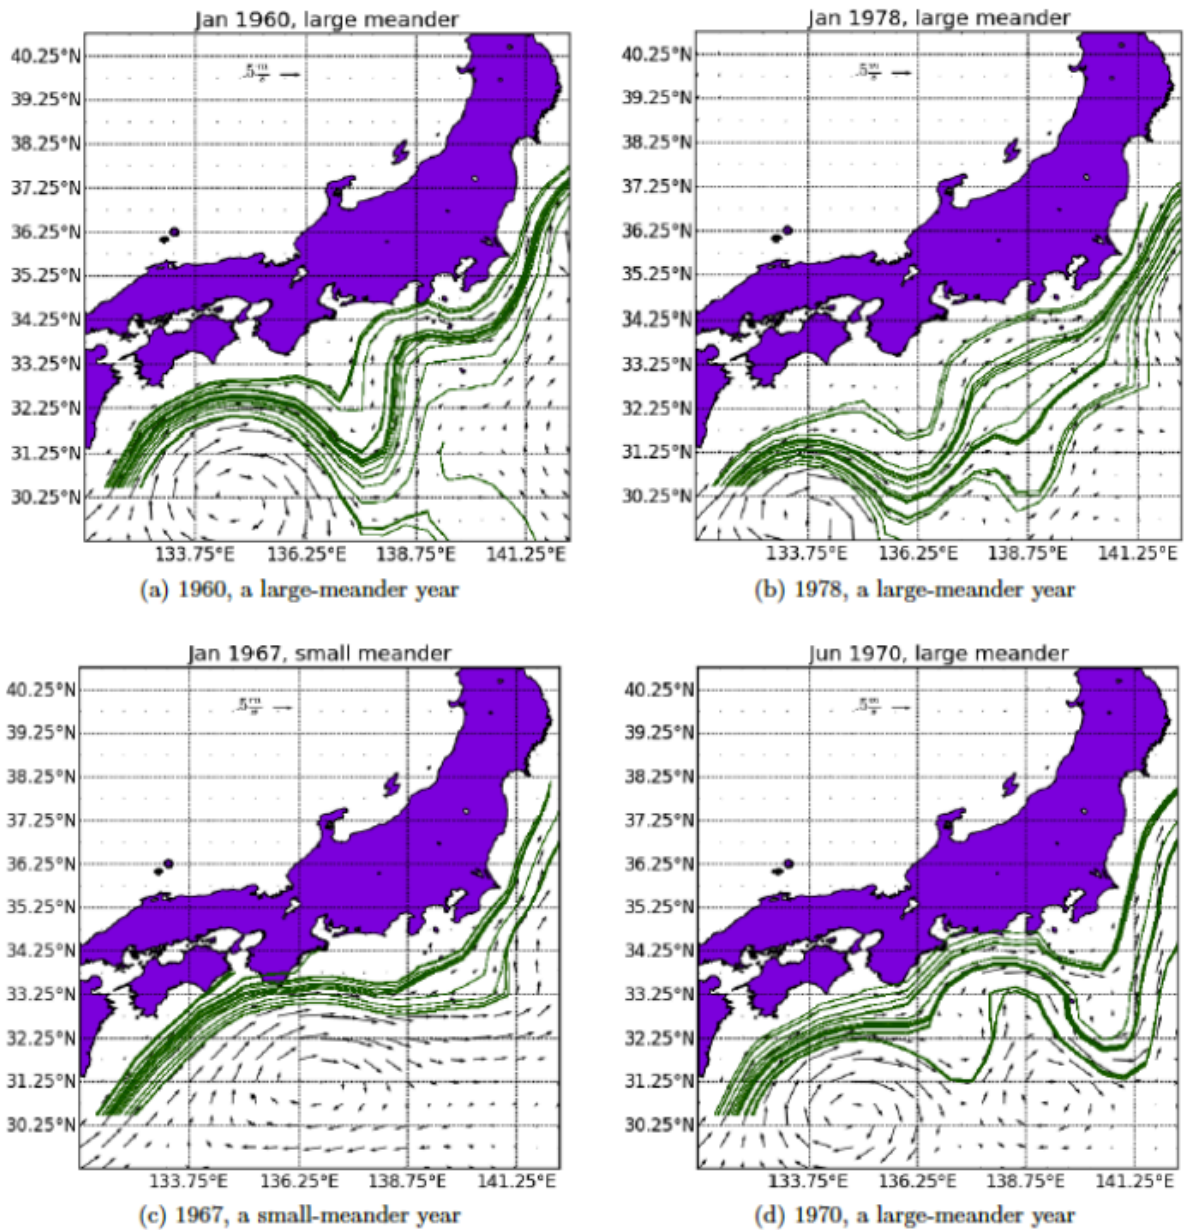


Figure 2.7: Velocity fields and paths seeded in several years of interest. 1960 (panel [a]) and 1978 (panel [b]) are large-meander years according to both the diffusion maps and spectral clustering (DMSC) and the path indices; 1967 (panel [c]) is small-meander according to both; 1970 (panel [d]) is large-meander according to the DMSC index but not the path index.

months. When this region is included, the algorithm is robust to enlargement of the domain to include most of the Pacific basin.

To test whether DMSC simply finds the location of the gyre without identifying the Kuroshio, we ran the algorithm on a modified domain that excludes the region occupied by the gyre but includes the region just downstream of the gyre; this is the region where the current is localized in the small meander and highly variable in the large meander. The resulting modified DMSC index has a correlation of $r = .94$ with the original DMSC index; this suggests that DMSC identifies variability in the Kuroshio and not simply the location of the gyre.

Finally, we check the sensitivity of DMSC to the local scale parameter, ϵ , which is roughly the radius around each data point \mathbf{x}_i such that the nearby data are well approximated with a hyperplane. In the previous analysis, we used $\epsilon = 3 \text{ m s}^{-1}$; this is comparable with the standard deviation of the distances of velocity fields, which is 3.12 (corresponding to $\sim .2 \text{ m/s}$ at each grid point). The results are robust to using $\epsilon \in (2, 6) \text{ m s}^{-1}$. This is consistent with the fact that if ϵ is chosen too small (significantly less than one standard deviation), the transition probability away from any point \mathbf{x}_i becomes negligible; if ϵ is chosen too large (greater than two standard deviations), the random walk is approximately uniform over the data and the underlying structure is lost.

2.3.5 Comparison with AVISO SSH Data

Given that the results presented in Sections 2.3.1 and 2.3.2 implicate the recirculation gyre in the bimodal nature of the Kuroshio, we must confirm that the bimodality of the system is not an artifact of the model used in the SODA reanalysis. To do this, we compare SODA-based indices to ones calculated from AVISO satellite sea surface height (SSH) data for the years 1993-2009.

First, we calculate the DMSC index for the years in which AVISO data are available and compare it with a DMSC index calculated using the AVISO data. The indices agree well

($r = .86$, but neither shows bimodality for this period. This is a result of the AVISO time series being too short to contain any significant excursions into the large meander, rather than a failure either of the data or of our methods. This is supported by the fact that the DMSC, gyre, and path indices calculated using the full SODA time series show no significant excursions into the large meander between 1993 and 2009.

To show that the SODA data are consistent with AVISO data (and that the bimodality we find is therefore not an artifact of the SODA model), we calculate the gyre index using the AVISO data. A comparison of a gyre index calculated using AVISO data with the SODA DMSC index for the years 1993-2009 yields a correlation coefficient of $r = .83$. We therefore conclude that the features found using SVD and DMSC on the SODA reanalysis are not model artifacts.

2.4 Discussion

As noted in Section 2.3.1, the small meander state contains only paths that remain near the coast, while the large meander state contains both paths that remain near the coast and those that exhibit traditional large-meander behavior. From this observation, we conclude that large and small meanders are separated by variation in path distance from the coast more than by the average path distance from the coast: the small meander has low variance in path distance from the coast while the large meander has high variance in path distance. Accurate models of the Kuroshio must explain this change in variability in addition to the previously observed change in the current axis.

Importantly, the gyre moves towards the southwest of Japan during the large meander. This is a region where the Kuroshio is highly localized at all times; thus, the gyre moves closer to the incoming current and spins a portion of it off to the southeast. Because only the easternmost portion of the current is affected by the gyre, most paths that originate at the western portion of the current continue to follow small-meander trajectories. This accounts for the increased variance of path location during large meander years despite a relatively

small shift in the mean path location.

While this study does not explicitly yield information about the transitions between the two states, it provides qualitative evidence for the transition mechanism proposed by Qiu and Miao (2000):

1. Low-potential vorticity (PV) water is carried northward by the Kuroshio.
2. Accumulation of a low PV anomaly strengthens the recirculation gyre to the south of the Kuroshio and presses the current to the coast, leading to a small meander.
3. The strengthening gyre increases shear, causing instability that leads to the large meander.
4. The large-meander path leads to cyclonic eddies that mix PV, weakening the gyre.
5. The system restarts when low-PV water is once more transported from the south.

Our results show that, as in the above mechanism, the small meander is prevalent when inflow (and therefore current velocity) is high, the gyre velocities are strongest, and the gyre is near the northeasternmost part of its range (step 2). The gyre then moves to the southwest and weakens (steps 3 and 4). When this happens, a portion of the current is spun off to the southeast; this leads to the formation of the large meander. As we have shown, the large meander state corresponds to lower inflow and a southwestern gyre location, as predicted by steps 3 and 4 of the mechanism above. Thus, while our analysis does not directly yield information about the dynamics of the system, it shows that the salient features of the small and large meanders agree well with those that would result from the dynamics proposed by Qiu and Miao (2000).

2.5 Conclusions

By using SVD in tandem with DMSC, we are able to characterize both the variability and the bimodality of the Kuroshio. We find that the recirculation gyre south of Japan is both

the dominant form of variability and the dominant bimodal feature in the system. With DMSC, we automatically create a one-dimensional index that characterizes the state of the Kuroshio; this index is highly correlated with the location of the gyre. Further, we find that small and large meanders are well-characterized by small and large path variation, respectively, but are poorly characterized by mean path location. We thus conclude that:

- The Kuroshio is bimodal, with high inflow and low path variability in the small meander and low inflow and high path variability in the large meander.
- The dominant form of variability and the dominant bimodal feature in the system is the location of the recirculation gyre.
- The large and small meanders are better characterized by path variability than by mean path location.

More work is needed to definitively characterize the dynamics of the Kuroshio and the mechanism by which it transitions between the small and large meanders. Further, Pierini (2006) proposes that the Kuroshio Extension east of Japan displays bimodality that is driven by an internal oscillation related to the mechanism responsible for bimodality of the Kuroshio South of Japan. A similar study of the Kuroshio Extension and of the entire Kuroshio system could shed light on Kuroshio Extension variability and its connection of bimodality of the Kuroshio south of Japan.

By rigorously characterizing the states of the Kuroshio, we hope to facilitate further work on the dynamics of the Kuroshio. Information about the mechanism of transition can help in developing better models as well as in further data analysis of the transition itself (Miller et al., 1994; Vanden-Eijnden and Weare, 2012, 2013a).

CHAPTER 3

MAXIMIZING SIMULATED TROPICAL CYCLONE INTENSITY WITH ACTION MINIMIZATION

3.1 Introduction

Despite extensive study, a consensus has yet to emerge about the mechanisms responsible for tropical cyclone (TC) intensification (Montgomery and Smith, 2014). The study of intensification remains challenging, owing to the number of scales and processes involved, e.g. mesoscale vortices, convection, turbulence, and microphysics (Rogers et al., 2013; Kaplan et al., 2015), as well as the expense and danger of obtaining in situ observations. Because category 3, 4, and 5 TCs cause 85% of all TC damage in the United States despite comprising only 24% of storms (Pielke and Landsea, 1998), the study of these relatively less frequent storms is of particular interest. Further, many studies suggest that global warming will increase the frequency of the most intense TCs (Emanuel, 2005; Webster et al., 2005; Knutson et al., 2010). Understanding the intensification mechanisms of the strongest TCs is thus both especially important and especially computationally expensive, since it requires TC simulations that are both highly resolved enough to capture critical processes and long enough to provide sufficient sample size (Murakami et al., 2012; Roberts et al., 2015).

Kaplan and DeMaria (2003) showed that 83% of North Atlantic category 4 and 5 TCs undergo rapid intensification, defined as an increase of at least 15 m s^{-1} in the maximum sustained surface wind speed over 24 hours (as compared to 31% of all TCs). TCs in climate models generally fail to undergo rapid intensification, thus leading to the underprediction of intense storms (Rogers, 2010; Camargo and Wing, 2016). While rapid intensification may not be a distinct process but rather a tail event of the distribution of intensification rates (Kowch and Emanuel, 2015), the inability of climate models to sample this tail presents a challenge to TC simulation.

Investigating the tail of the intensification distribution by direct computer simulation

faces two principle difficulties: first, climate models require higher resolution than currently possible to conclusively differentiate among proposed rapid intensification mechanisms (Gentry and Lackmann, 2010). Second, the rarity of strong TCs (\sim one category 5 storm every five years in the Atlantic basin (MacAdie et al., 2009)) makes running climate models at high enough resolutions and long enough time horizons to study rapid intensification prohibitively computationally expensive. The vast majority of simulation time would be wasted in quiescent states. In this study, we propose and implement a new methodology to study rapid intensification that circumvents these issues; it does so by targeting computational power at the problem of rapid intensification more specifically than can be done with direct simulation.

To better leverage computational resources for the problem of rapid intensification and thereby avoid the need for prohibitively long model runs, we apply an action minimization algorithm to the study of TC rapid intensification. The advantage of action minimization for the study of rapid intensification is that it allows for the targeted study of storms that experience the lowest pressure minima and the highest intensification rates. By targeting these storms, we substantially reduce the computational cost of accessing and studying the tail of the intensification distribution. The algorithm adds a series of perturbations to a model trajectory over time, nudging the model toward states with some characteristic of interest (in this case, an intense TC). Each perturbation is indistinguishable from noise and consists of an adjustment to each value in several two- or three-dimensional physical fields: zonal and meridional wind, temperature, surface pressure, water vapor mixing ratio, and geopotential.

We obtain the set of perturbations over time by solving an optimization problem - namely, minimization of the action functional. The action functional imposes costs both on perturbation size and on TC intensity (measured here by minimum central surface pressure) in the final state. We solve this optimization problem by leveraging the four-dimensional variational analysis (4d-VAR) capability of the Weather Research and Forecasting (WRF)

model (Huang et al., 2009). The optimization we solve seeks perturbations that are small and realistic, thus providing a physically plausible rapid intensification trajectory. By using a background covariance matrix generated by forward model runs, we ensure physically relevant perturbations. In fact, given an initial condition, action minimization yields perturbations that comprise the likeliest trajectory from that initial condition to an intense TC under a stochastically perturbed model. In this study, we take as initial conditions already formed TCs of intermediate intensity; however, in principle it should be possible to apply action minimization to a quiescent atmosphere to study genesis as well as intensification.

While our methodology uses software tools originally built for the purpose of data assimilation and 4d-VAR, we solve a fundamentally different problem. The purpose of 4d-VAR is to improve the agreement between a model trajectory and previously observed data (e.g. sub-states of the full state of the weather at several points in time), thereby improving estimates of weather states that actually occurred (Zupanski, 1996). In contrast, action minimization seeks realistic model trajectories to some hypothetical state (or set of states) of interest, thereby improving our understanding of how a model transitions to such states. In this study, we investigate the processes by which the WRF model undergoes rapid intensification to form intense TCs; we do not aim to provide information about the history of specific storms.

Our work builds on that of Hoffman et al. (2006a) and Hoffman et al. (2006b), which first applied algorithms based on 4d-VAR to optimal TC perturbation. These studies were both done in a geoengineering context; Hoffman et al. (2006b) sought perturbations to reduce cyclone damage based on wind speeds and property values, while Hoffman et al. (2006a) sought perturbations for deflecting cyclone tracks away from population centers. These studies were done at 20 km resolution with a diagonal background covariance matrix, allowed perturbations only to the initial condition, relied on TC bogusing (artificially embedding an axisymmetric vortex into a regional weather model), and in the case of Hoffman et al. (2006a) constrained the model trajectory to be near a specific target state. While related,

our work differs in that it seeks perturbations that provide physical insight into the intensification process. We achieve this by maximizing storm intensity at 6 km resolution, using a background covariance matrix that enforces physically realistic perturbations, allowing perturbations throughout the model trajectory, and using storms that arise from our choice of initial and boundary conditions.

The purpose of this study is to demonstrate the feasibility of using action minimization to study TC rapid intensification and to investigate the action-minimizing perturbations for information about the mechanisms of intensification. We apply action minimization to simulations of two category 3 hurricanes, Danny and Fred, that were near the threshold for rapid intensification. We argue that action minimization causes both storms to undergo rapid intensification via physically realistic pathways. We also compare the computational cost of action minimization to that of studying these intensification pathways with ensemble simulation, showing an approximately ten-fold reduction in computational cost resulting from action minimization. We conclude that the extreme events at the tail of the TC intensification distribution can be usefully interrogated using action minimization.

The optimal perturbations show that reduction of low-level shear, at least above a threshold of approximately 5 m s^{-1} , has a much stronger effect on intensification compared to reduction of high-level shear. Our results support the findings of Wong and Chan (2004) and (Nolan and McGauley, 2012) in that they suggest that there is an upper threshold above which shear is highly detrimental to intensification but not necessarily a lower threshold below which reducing shear is also detrimental to intensification. Further, we show that asymmetric, time-dependent patterns of heating can cause significant TC intensification beyond symmetric, azimuthally-averaged heating. In doing so, we find a regime of non-linear response to asymmetric heating that has not been extensively studied in previous work. The patterns of heating produced by action minimization are consistent with recent observational and modeling studies that posit the importance of relatively low-level latent heating for intensification (Park and Elsberry, 2013; Zagrodnik and Jiang, 2014; Tao and Jiang, 2015).

The action-minimizing heating patterns also support work that argues that TC response to asymmetric heating can be nonlinear and depends significantly on pre-existing vortex asymmetries (Moller and Shapiro, 2005).

3.2 Methods

3.2.1 Action minimization

Algorithms to find the most likely path of a stochastic dynamical system between two states (or sets of states) have been used to great effect in computational chemistry for two decades (Jonsson et al., 1998; E et al., 2002; E and Vanden-Eijnden, 2010). In most of those applications the goal is to study the very infrequent transitions between two long lived states by identifying a trajectory of the process connecting the two states that is exponentially more likely than other potential trajectories in the small temperature limit as detailed by the theory of large deviations (Freidlin and Wentzell, 1984). The method that we apply here is closely related to the method in E et al. (2004), though here we also allow perturbations of the initial condition and we do not perturb the process at every model time step.

In addition to Hoffman et al. (2006b) and Hoffman et al. (2006a), several previous studies have suggested the potential impact of rare event simulation and analysis tools on geophysical applications: a similar path finding approach was considered in the context of rogue ocean waves by Dematteis et al. (2018); methods designed to generate random samples of rare events were considered in the contexts of rare transitions of an ocean current model in Weare (2009) and Vanden-Eijnden and Weare (2013b) and extreme heat waves in Ragone et al. (2018).

Given a model of some physical process, we would like to add small perturbations to true model trajectories in order to achieve a transition into some rare state. The goal of the action minimization algorithm is to minimize the magnitude of the perturbations added to the model, thus obtaining the most physically realistic path from the initial condition

to a desired end state. Similar to action minimization, the 4d-VAR cost function penalizes perturbations away from the model trajectory; however, our final cost draws the model trajectory toward final states with intense TCs, whereas the 4d-VAR final cost penalizes the squared distance from observations of states along the model trajectory.

The following quantities are necessary to perform the minimization:

- \mathbf{x}_i : The state of the system at time index i . Includes horizontal velocities, temperatures, pressures, water vapor mixing ratios, and geopotentials at every gridpoint.
- $F(\cdot)$: The function that integrates the model state to the next time index. In the absence of perturbations, this gives $\mathbf{x}_i = F(\mathbf{x}_{i-1})$.
- $\hat{\eta}_i$: The perturbation at time index i .
- $\Phi(\mathbf{x}_i)$: A function (order parameter) that measures the intensity of a TC at state \mathbf{x}_i .
- N : The number of optimization time steps.
- B : The background covariance matrix.
- R_i, R_f : Constants that balance running cost against final cost.
- \mathbf{a} : The initial state \mathbf{x}_0 is restrained to remain near the state \mathbf{a} .

Importantly, N is the number of perturbations added to a model trajectory; it is not necessarily the same as the model time step. In principle, perturbations can be added every time step; however, we use one perturbation for every hour of model time whereas the time step for model dynamics is 30 seconds. This is done in order to avoid the memory cost of storing perturbations at each time step. We demonstrate in section 3.3.7 that this choice of perturbation frequency does not qualitatively affect the results.

We seek to minimize the perturbations added to a true model trajectory in order to obtain a trajectory that results in a lower value of our order parameter $\Phi(\mathbf{x}_N)$ at the final

state. We define the action functional J as:

$$\begin{aligned}
 J &= \sum_{i=0}^{N-1} \frac{1}{2R_i^2} \hat{\eta}_i^T B^{-1} \hat{\eta}_i + \frac{1}{R_f^2} \Phi(\mathbf{x}_N) \\
 \mathbf{x}_i &= F(\mathbf{x}_{i-1}) + \hat{\eta}_i, \quad 0 < i < N \\
 \mathbf{x}_0 &= \mathbf{a} + \hat{\eta}_0
 \end{aligned} \tag{3.1}$$

J is large both if the perturbations $\hat{\eta}_i$ are large (high running cost) or if the order parameter $\Phi(\mathbf{x}_N)$ is large (high final cost). The order parameter Φ we choose is the minimum surface pressure at the final state. By minimizing J , we seek a model trajectory with both small perturbations (indicating a physically relevant intensification process) and a final state that has a center of low minimum surface pressure. Given a domain over the tropical North Atlantic, such a final state should correspond to an intense TC. Constants $\{R_i\}_{i=0}^{N-1}$ and R_f determine the relative importance of the running cost at each time step and of the final cost. We choose them heuristically so as to get the maximal intensification without causing the model to crash.

The covariance matrix B is obtained using the National Meteorological Center (NMC) method from Parrish and Derber (1992), where the covariance is thus the time averaged outer product of the differences between 12- and 24-hour forecasts produced with the WRF model over the course of several months (here chosen to correspond with the Atlantic hurricane season). As noted in Xiao et al. (2009), this introduces some approximation since the correlations obtained in this manner are accurate for climatological data but not necessarily for TC conditions.

To improve the conditioning of equation 3.1, we substitute $\eta_i = B^{-\frac{1}{2}} \hat{\eta}_i$, which yields:

$$\begin{aligned}
J &= \sum_{i=0}^{N-1} \frac{1}{2R_i^2} \eta_i^T \eta_i + \frac{1}{R_f^2} \Phi(\mathbf{x}_N) \\
\mathbf{x}_i &= F(\mathbf{x}_{i-1}) + B^{\frac{1}{2}} \eta_i, \quad 0 < i < N \\
\mathbf{x}_0 &= \mathbf{a} + B^{\frac{1}{2}} \eta_0
\end{aligned} \tag{3.2}$$

Minimizing J requires the derivatives of both the running cost and the final cost with respect to the set of perturbations η_i :

$$\frac{\partial J}{\partial \eta_i} = \frac{1}{R_i^2} \eta_i + \frac{1}{R_f^2} \left(\frac{\partial \mathbf{x}_N}{\partial \eta_i} \right)^T \frac{\partial \Phi(\mathbf{x}_N)}{\partial \mathbf{x}_N} . \tag{3.3}$$

where $\eta = \{\eta_i\}_{i=0}^{N-1}$ is the set of all perturbations at all times. We apply the chain rule to obtain the derivative of the final cost:

$$\begin{aligned}
\frac{\partial \mathbf{x}_N}{\partial \eta_i} &= \frac{\partial \mathbf{x}_N}{\partial \mathbf{x}_i} \frac{\partial \mathbf{x}_i}{\partial \eta_i} \\
\frac{\partial \mathbf{x}_i}{\partial \eta_i} &= B^{\frac{1}{2}} .
\end{aligned} \tag{3.4}$$

We thus need $\frac{\partial \mathbf{x}_N}{\partial \mathbf{x}_i}$, i.e. the derivative of the final state with respect to each previous physical state \mathbf{x}_i . We formulate a recursion that simplifies calculation of the necessary derivatives:

$$\begin{aligned}
\mathbf{x}_0 &= \mathbf{a} + \hat{\eta}_0 \\
\mathbf{x}_1 &= \hat{\eta}_1 + F(\mathbf{x}_0) \\
\mathbf{x}_2 &= \hat{\eta}_2 + F(\mathbf{x}_1) = \hat{\eta}_2 + F(\hat{\eta}_1 + F(\mathbf{x}_0)) \\
\mathbf{x}_3 &= \hat{\eta}_3 + F(\mathbf{x}_2) = \hat{\eta}_3 + F[\hat{\eta}_2 + F(\hat{\eta}_1 + F(\mathbf{x}_0))] .
\end{aligned} \tag{3.5}$$

Using equation 3.5, we differentiate each \mathbf{x}_j with respect to each \mathbf{x}_i , which makes clear the general form of $\frac{\partial \mathbf{x}_N}{\partial \mathbf{x}_i}$:

$$\begin{aligned} \frac{\partial \mathbf{x}_3}{\partial \mathbf{x}_2} &= F'(\mathbf{x}_2) \\ \frac{\partial \mathbf{x}_3}{\partial \mathbf{x}_1} &= \frac{\partial \mathbf{x}_3}{\partial \mathbf{x}_2} \frac{\partial \mathbf{x}_2}{\partial \mathbf{x}_1} = F'(\mathbf{x}_2)F'(\mathbf{x}_1) \\ \Rightarrow \frac{\partial \mathbf{x}_N}{\partial \mathbf{x}_i} &= F'(\mathbf{x}_{N-1}) \dots F'(\mathbf{x}_{i+1})F'(\mathbf{x}_i) \end{aligned} \quad (3.6)$$

$F'(\mathbf{x})$ is the derivative of the flow; that is, F' measures how a change in initial state \mathbf{x} affects the state obtained by integrating \mathbf{x} forward with $F(\mathbf{x})$. We substitute into equation 3.4 to obtain the derivative of the final state \mathbf{x}_N with respect to the perturbations η_i :

$$\left(\frac{\partial \mathbf{x}_N}{\partial \eta_i} \right)^T = B^{\frac{1}{2}} [F'(\mathbf{x}_i)]^T [F'(\mathbf{x}_{i+1})]^T \dots [F'(\mathbf{x}_{N-1})]^T$$

Substituting back into equation 3.3 yields the full derivative of action functional J with respect to perturbations η_i :

$$\frac{\partial J}{\partial \eta_i} = \frac{1}{R_i^2} \eta_i + \frac{1}{R_f^2} B^{\frac{1}{2}} [F'(\mathbf{x}_i)]^T [F'(\mathbf{x}_{i+1})]^T \dots [F'(\mathbf{x}_{N-1})]^T \left(\frac{\partial \Phi(\mathbf{x}_N)}{\partial \mathbf{x}_N} \right) \quad (3.7)$$

We use equation 3.7 and a conjugate gradient method to find the action-minimizing perturbations η_i^* of equation 3.2. The trajectory of physical states $\{\mathbf{x}_i^*\}_{i=0}^N$ that corresponds to this solution is given then given by:

$$\mathbf{x}_i^* = \begin{cases} \mathbf{a} + B^{1/2}\eta_0^* & i = 0 \\ F(\mathbf{x}_{i-1}^*) + B^{1/2}\eta_i^* & 0 < i < N \end{cases} \quad (3.8)$$

That is, the initial state is given by adding the $t = 0$ perturbation to the physical initial condition. Subsequent states \mathbf{x}_i^* are obtained by integrating the previous state \mathbf{x}_{i-1}^* and adding the corresponding optimal perturbation η_i^* .

The optimal trajectory of states $\{\mathbf{x}_i^*\}_{i=1}^{N_t}$ can be considered as a particular realization of a stochastic version of the original forward model. This implied stochastic model has Gaussian noise with covariance BR_i^2 added at each time index i . Such a model formulation is nearly equivalent to that used in 3d- and 4d-VAR (Huang et al., 2009), including for the study of TCs (Singh et al., 2011), and in other common applications including ensemble simulation (Pereira and Berre, 2006). This implied model will be used to evaluate the efficacy of action minimization as compared to ensemble simulation in section 3.2.4. The use of small stochastic perturbations in ensemble simulations serves to diversify the set of simulated trajectories and is often justified as representing subgrid effects and other model errors (Oouchi et al., 2006; Leutbecher and Palmer, 2008). The trajectory produced by action minimization is therefore physical in the same sense that those produced by an ensemble simulation of the same stochastic model are physical. In fact, our action minimization results in the maximum likelihood trajectory within the stochastic model conditioned on trajectories having equal or lower terminal values of the order parameter, i.e. on trajectories x_i with $\Phi(x_N) \leq \Phi(x_N^*)$. To see this note that the likelihood of any trajectory under the stochastic model is proportional to

$$e^{-\sum \hat{\eta}^T B^{-1} \hat{\eta} / 2R_i^2}.$$

If a trajectory other than x_i^* has both higher likelihood and equal or lower value of $\Phi(x_N)$ then it would have a smaller action than x_i^* .

3.2.2 WRFPLUS

Equation 3.7 shows that the derivative of the model flow, $F'(\mathbf{x}_i)$, is necessary to compute the derivative of the action functional; this in turn requires a model with adjoint capabilities. Among open source models with this capability, the WRF Variational Data Assimilation (WRFDA) model (Huang et al., 2009) is state-of-the-art. While WRFDA is normally used for four-dimensional variational data assimilation (Barker et al., 2012), the module WRFPLUS contains stand-alone tangent linear and adjoint versions of WRF as obtained by automatic differentiation and subsequent by-hand correction. We use this adjoint model to obtain derivatives of the model flow.

The adjoint model takes as input a perturbation vector η and state \mathbf{x} and returns sensitivity $F'(\mathbf{x})^T \cdot \eta$, which is valid at the previous time index. Sensitivities are available for a subset of the physical variables, and the optimization is performed over perturbations with respect to these variables. Specifically, we optimize over perturbations to \mathbf{u} (zonal wind), \mathbf{v} (meridional wind), \mathbf{p}_{sfc} (surface pressure), \mathbf{T} (temperature), \mathbf{r} (water vapor mixing ratio), and ϕ (geopotential, which describes the gravitational potential energy per unit mass at a given height).

3.2.3 Experimental design

To test the ability of the algorithm to force TCs within WRF to undergo stronger rapid intensification, we chose two category 3 Atlantic TCs as test cases: 1) Hurricane Danny (2015), and 2) Hurricane Fred (2009). We use these storms as they were of similar intensity and stayed over open ocean, allowing for a simpler proof of concept. Simulations for both storms are done at 6 km resolution with a 36 hour time horizon chosen such that peak intensity is attained at the end of the simulation.

Hurricane Danny formed as the result of a tropical wave on August 14, 2015. Thanks to abnormally low environmental wind shear, it relatively quickly reached a minimum pressure of 960 hPa and maximum wind speed of 56 m s^{-1} at 1200 UTC on August 21 (Stewart,

2016). Hurricane Danny was an abnormally small storm, with tropical-storm-force winds extending only to a radius of 60 km from its center. Figure 3.1 shows the storm track for Hurricane Danny from the National Oceanic and Atmospheric Administration National Hurricane Center (NOAA NHC). The domain size we use for action minimization of Hurricane Danny is 3000 km \times 1800 km (500 \times 300 grid points) and is centered at 15°N and 42°W.

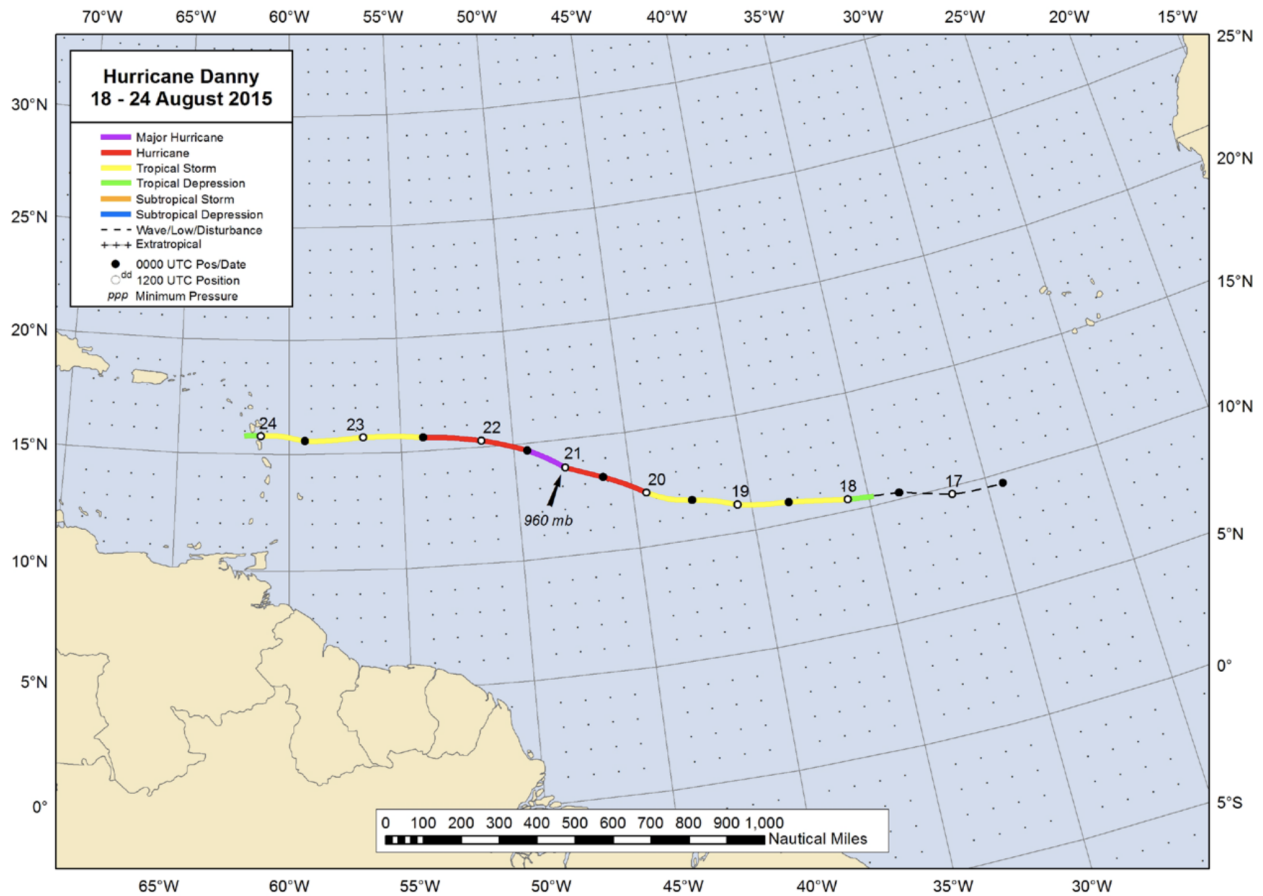


Figure 3.1: NOAA NHC track Hurricane Danny (Stewart, 2016).

Hurricane Fred resulted from a tropical depression that formed after a tropical wave crossed the west coast of Africa September 6. This depression intensified into Tropical Storm Fred on September 8 and Hurricane Fred on September 9, reaching a minimum pressure of 958 hPa and maximum wind speed of 54 m s⁻¹ later that day (Berg and Avila, 2011). Figure 3.2 shows the storm track for Hurricane Fred from the National Oceanic and Atmospheric Administration National Hurricane Center (NOAA NHC) (Brennan, 2009). The domain size

we use for action minimization of Hurricane Fred is $2250 \text{ km} \times 2250 \text{ km}$ (375×375 grid points) and is centered at 15°N and 28°W .

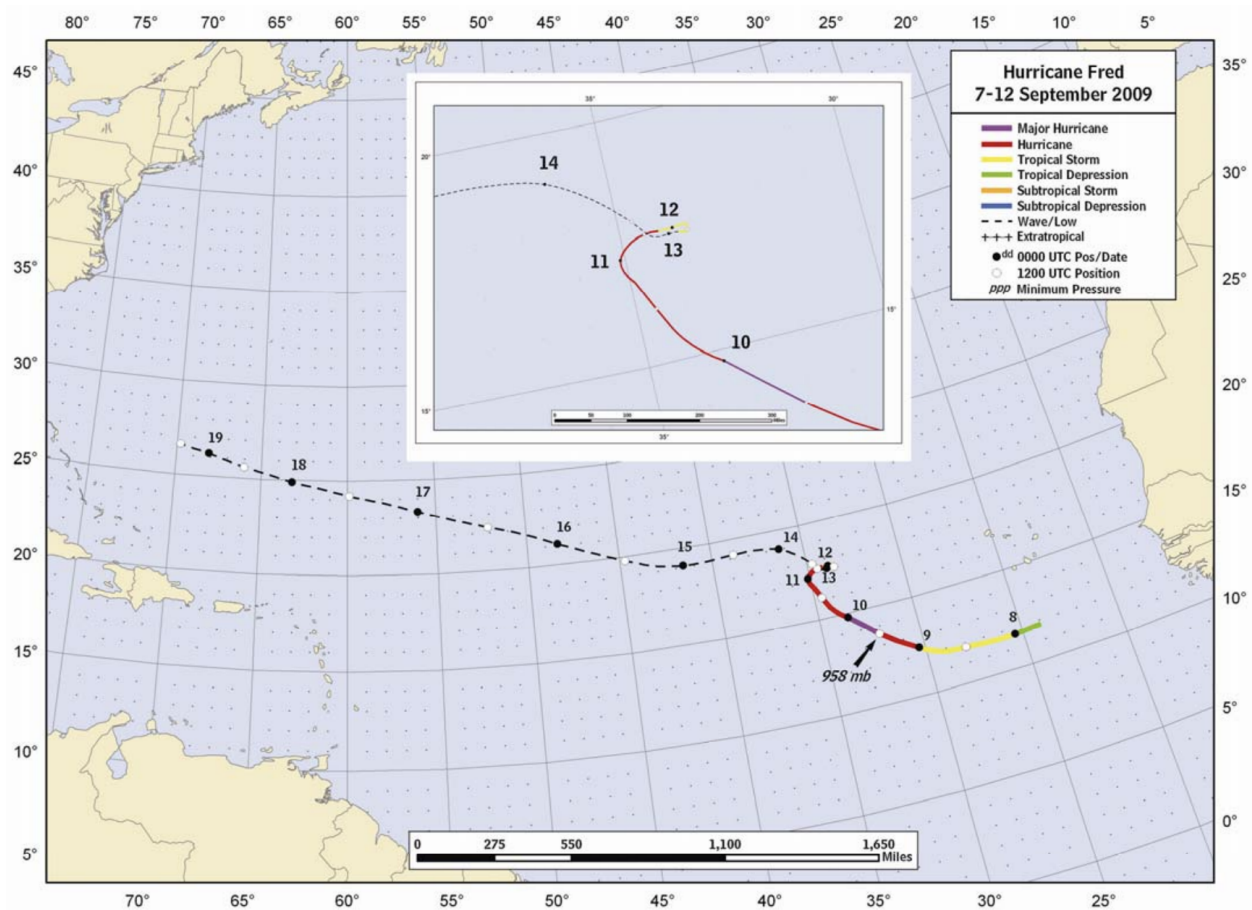


Figure 3.2: NOAA NHC track Hurricane Fred (Brennan, 2009).

The physics schemes follow the work of Fierro et al. (2009) and Judt et al. (2016). We use the WRF Single-Moment 6-Class Microphysics scheme (Hong and Lim, 2006) to simulate moist processes, the Yonsei University scheme (Hong et al., 2006) for planetary boundary-layer processes, and the Kain-Fritsch cumulus parameterization scheme (Kain and Fritsch, 1993). The parameters used for the action minimization algorithm are: running cost weight $R_i = 0.6$ for each time step i and final cost weight $R_f = 1.0$. These constants are empirically chosen to give maximal weight to the final cost without causing the model to enter unphysical states and crash. Perturbations are added once every hour of integration time, as opposed to at every time step, so as not to incur excessive memory cost associated with

storing perturbations at every time step. Low resolution tests suggest that this perturbation frequency is sufficient, as decreasing the perturbation frequency to once every four hours yields nearly identical results to those presented in 3.3. Sensitivity of results to both time horizon and perturbation frequency are discussed in section 3.3.7.

In all cases, the minimization of equation 3.2 is done using a Fletcher-Reeves conjugate gradient method (Nocedal and Wright, 2006). Initial and boundary conditions for all runs are obtained from the NOAA GFS-ANL 0.5° dataset.

3.2.4 Action minimization for the study of rapid intensification

While ensemble simulation is a natural tool for the study of cyclone genesis, intensification, and predictability (e.g. Van Sang et al., 2008), action minimization provides a more computationally efficient way to explore the tail of the cyclone intensity distribution. We test this by comparing the computational cost of using action minimization to cause further intensification of Hurricane Danny with the computational cost of generating an equally intense storm via ensemble simulation. As noted in section 3.2.1, the background covariance matrix B defined in that section implies a stochastic model wherein the noise is drawn once every hour from a Gaussian distribution with covariance B . Action minimization output is then equivalent to a particular realization of this stochastic physical model. As noted in section 3.2.1, this model formulation is one that has previously been used for ensemble study and data assimilation; further, the distributions of the prognostic variables (dry column air mass, horizontal and vertical velocities, potential temperature, and geopotential) under this model are indistinguishable from those obtained from the unperturbed WRF model. We run this model 100 times using the same initial condition \mathbf{a} as in the action minimization run for Hurricane Danny.

Figure 3.3 shows a histogram of the minimum pressures obtained with the stochastic model and compares them to the minimum pressure (962 hPa) obtained with action minimization. None of the ensemble members reaches as low a minimum pressure as the action

minimization output; the closest ensemble member intensifies to 964 hPa. The unperturbed minimum pressure is 970 hPa, while the distribution of ensemble pressure minima has a mean of 971 hPa and a standard deviation of 3.2 hPa. This means that the minimum pressure obtained via action minimization is 2.8 standard deviations lower than the ensemble mean. Assuming a normal distribution of ensemble pressures (this is an approximation since the ensemble distribution appears slightly right-skewed), obtaining an equally low minimum pressure using ensemble simulation is approximately a one in 400 event.




Figure 3.3: Histogram of minimum pressures for Hurricane Danny obtained by ensemble simulation (blue crosses) and the minimum pressure obtained by action minimization (red cross).

The cost advantage for action minimization relative to ensemble simulation should increase rapidly as one seeks samples further into the tails of the ensemble distribution. Indeed, if the probability of observing a trajectory $\{\mathbf{x}_i\}_{i=0}^N$ with $\Phi(\mathbf{x}_N) \leq L$ under the stochastic model is p , then the expected number of trajectories required to observe the event $\Phi(\mathbf{x}_N) \leq L$

is exactly $1/p$. For most stochastic models, one would expect that p will decrease very rapidly as L decreases (e.g. exponentially under our Gaussian perturbation model). To assess the efficiency of our implementation of action minimization to access the lower tail of the pressure distribution, we compare the computational cost of running action minimization with the cost of running a 400-member ensemble. The action minimization algorithm uses a conjugate gradient method, which in our experience takes 3 - 7 iterations to converge. Each iteration involves one forward run and one adjoint run to calculate the derivative in equation 3.7, and the cost of each adjoint run is no more than three times the cost of a forward run (Farrell et al., 2013). Each iteration also involved a line search in the direction of the gradient, which takes at most 5 forward runs. The computational cost of each iteration is therefore at most 9 times that of a forward run, so the cost of running action minimization is at most equivalent to that of 63 forward runs (operationally, we find that the typical cost is more similar to that of 30 forward runs). Thus, even in the worst case scenario, action minimization is over six times less costly for accessing the tail of TC intensity distribution as compared to ensemble simulation. We expect that there is substantial room to reduce the cost of action minimization by improving upon our initial implementation.

We conclude that while ensemble simulation may be the preferred tool in a number of contexts, the extreme events at the tail of the TC intensification distribution are best accessed using action minimization. Action minimization is thus a complementary tool to ensemble simulation in the sense that the latter should be used to characterize the bulk of the intensification distribution, while the former can be used to more efficiently fill out the tails.

3.3 Results

3.3.1 Comparison with maximum potential intensity

It is possible to estimate the maximum potential intensity (MPI) of a TC, usually defined as the minimum theoretically attainable surface pressure. Emanuel (1988) presents a theoretical limit based on the assumption that a TC functions as a Carnot heat engine; Holland (1997) presents a thermodynamic limit based on empirical relationships among relative humidity, central pressure, and wind speed. Most TCs do not attain maximum potential intensity, with the typical pressure depression magnitude approximately 55% of that predicted by MPI (DeMaria and Kaplan, 1994). The reasons for TCs failing to attain MPI could include dynamical controls, e.g. environmental flow and vertical shear (Zeng et al., 2007). It is therefore natural to ask whether action minimization can modify the dynamics in such a way as to bring TCs closer to MPI.

Figure 3.4 shows a comparison of the minimum pressures of the action minimization output and simple forward integration, as well as Emanuel and Holland MPI for Hurricanes Danny and Fred. We obtain the Emanuel MPI by solving equation 16 in Emanuel (1991), which is an implicit relationship for the central pressure derived from balancing the work done by a TC (to drive surface wind) against frictional dissipation in the boundary layer. A correction for dissipative heating is done according to Bister and Emanuel (1998). Solving this equation requires quantities easily obtained from WRF output, including temperature at the sea surface and at the top of the troposphere, ambient sea level pressure, and relative humidity at both the center and outer radius of the storm. We obtain Holland MPI by substituting into equation 2 of Holland (1997), which requires the same quantities.

The optimization for Hurricane Danny results in a decrease in minimum central pressure from 970.5 hPa to 962.0 hPa, accompanied by an increase in maximum surface wind speed from 54.4 m s⁻¹ to 60.9 m s⁻¹. For Hurricane Fred, the pressure decreases from 953.1 hPa to 943.3 hPa and the maximum wind speed increases from 57.6 m s⁻¹ to 64.4 m s⁻¹.

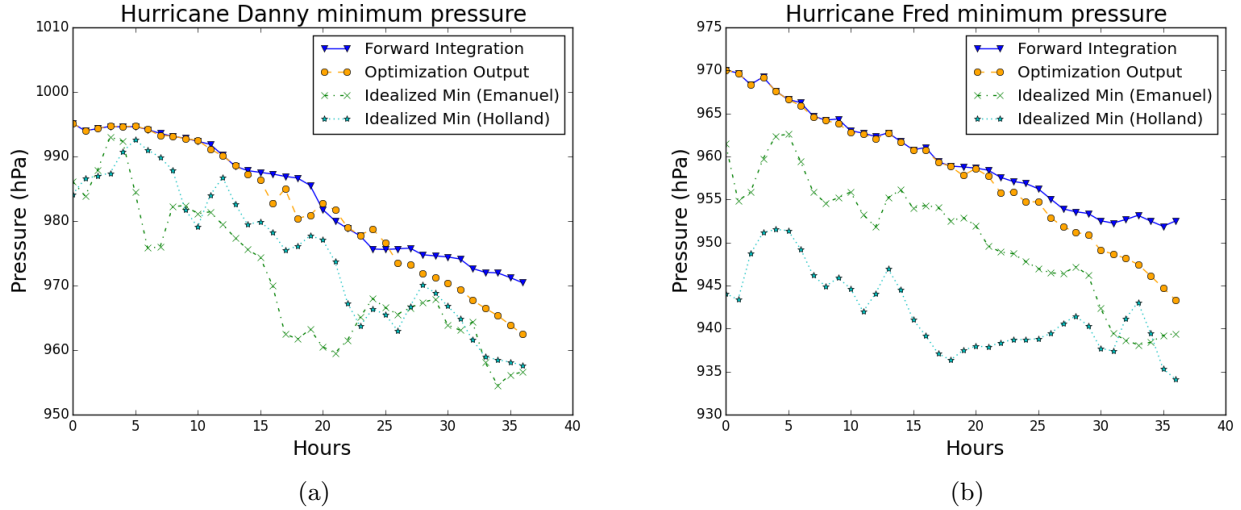


Figure 3.4: Minimum pressure as a function of time for forward integration (blue triangles) and action minimization output (orange circles) of Hurricanes Danny (panel a) and Fred (panel b). The green crosses and cyan stars are maximum potential intensities according to Emanuel (1991) and Holland (1997), respectively.

Thus, action minimization causes both storms to intensify from category 3 to category 4. Figure 3.5 shows the maximum surface winds as a function of time. The shaded regions show how fast surface wind must be to meet the RI criterion of intensification by at least 15 m s^{-1} over 24 hours. The unperturbed forward integrations of both Hurricanes Danny and Fred are on the borderline of the RI criterion; the action minimization output clearly meets the RI criterion in both cases. The action minimization algorithm is therefore able to ensure that RI occurs for these storms.

Although the optimization runs show similar decreases in pressure and increases in wind speed, the nature of the divergence of the optimized run from the forward integration is not the same in the two cases. While both optimization runs diverge from the forward integration during the second half of the time horizon, the optimized pressure trajectory for Hurricane Danny falls slightly above the pressure trajectory of the forward integration at several points. This is in contrast with the optimized pressure trajectory for Hurricane Fred, which remains below the forward integration trajectory once the two begin to diverge. We speculate that this is a result of differences in heating patterns discussed in section 3.3.3.

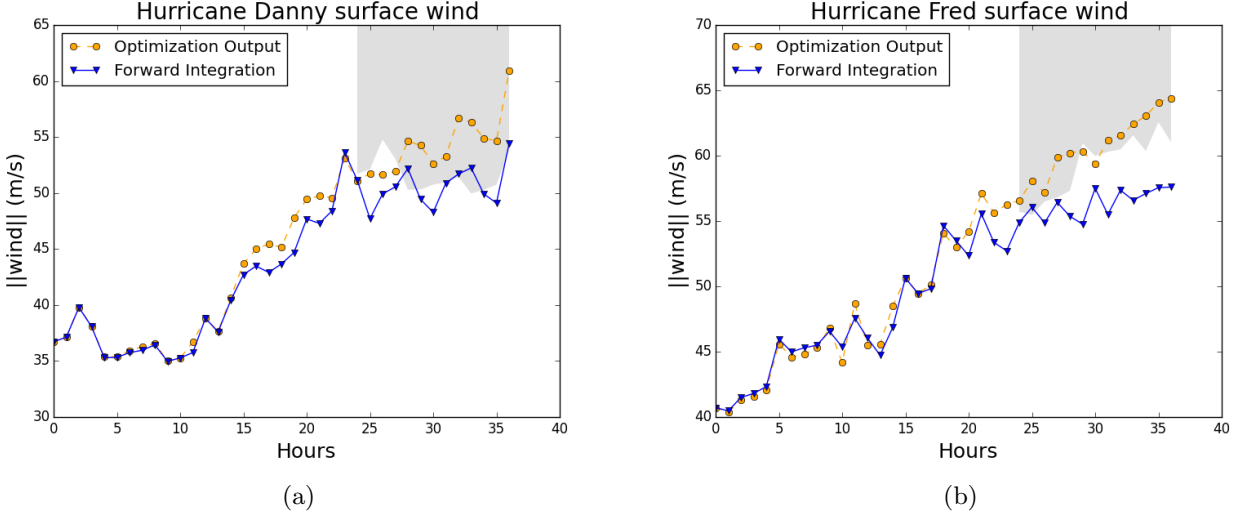


Figure 3.5: Maximum surface winds as a function of time for forward integration (blue triangles) and action minimization output (orange circles) of Hurricanes Danny (panel a) and Fred (panel b). The shaded regions mark wind speeds that must be reached for RI to occur.

Importantly, neither run results in lower central pressure than that indicated by maximum potential intensities, suggesting that the more intense storms resulting from action minimization are thermodynamically feasible. We verify this by integrating both storms forward from both the perturbed and unperturbed states at the final time in Figure 3.4. In both cases, the central pressures rise as the storms pass peak intensity, with those from the perturbed runs eventually approaching those from the forward integration. Further, we conduct tests at lower resolution in which we restart the optimization after 36 hours; that is, we use the final state of the original action minimization output as the initial state of a new action minimization problem with the subsequent 36 hours as its time horizon. This yields storms that stay near MPI for a longer time, but even in these simulations the minimum central pressures do not cross the MPI by more than 1-3 hPa or for longer than one optimization time step. The perturbed simulations thus do not violate thermodynamic constraints on TC intensity.

3.3.2 Comparison of perturbed and unperturbed storm structure

Figure 3.6 shows contour plots of the surface wind and the outgoing longwave radiation (OLR) at the final time for the forward integration and the optimization output for Hurricane Danny. Both fields show no qualitative structural inconsistencies between the forward integration and the optimization output. In particular, the optimization contains the key expected features including rain bands and an eye-wall. The radius of maximum wind decreases from approximately 38 km in the unperturbed forward integration to approximately 36 km in the action minimization output. Both surface wind and OLR fields for Hurricane Fred (not shown) exhibit similar behavior, with action minimization once again causing a slight reduction in the radius of maximum wind but no major qualitative changes. In both cases, the reduction in the radius of maximum wind accounts for slightly under half of the increase in maximum surface wind speed via conservation of angular momentum; thus, a significant portion of the intensification results from other mechanisms.

3.3.3 Temperature perturbations

To determine the mechanisms responsible for intensifying the optimized storm over the forward integration, we investigate the perturbations $\hat{\eta}_{\mathbf{i}}$ added to the forward trajectory. Figure 3.7 shows the 850 hPa temperature perturbations $\hat{\eta}_{\mathbf{T}}$ added 13, 9, 5, and 1 hours before the minimum pressure is attained for both Hurricane Danny (panels a-d) and Hurricane Fred (panels e-h).

Consistent with previous studies demonstrating the importance of warm cores in TC intensification (Schubert and Hack, 1982; Simpson et al., 1997), the action minimization algorithm adds a ring of low level warm temperature to the outside of the developing storm. The ring radius is initially approximately 250 km for Hurricane Danny and 350 km for Hurricane Fred, which is consistent with the radii of the two storms. As the TCs intensify, the warm air at the outside of the storms moves toward the storm centers, resulting in a final time warm core with 40 km and 60 km radii for Danny and Fred, respectively.

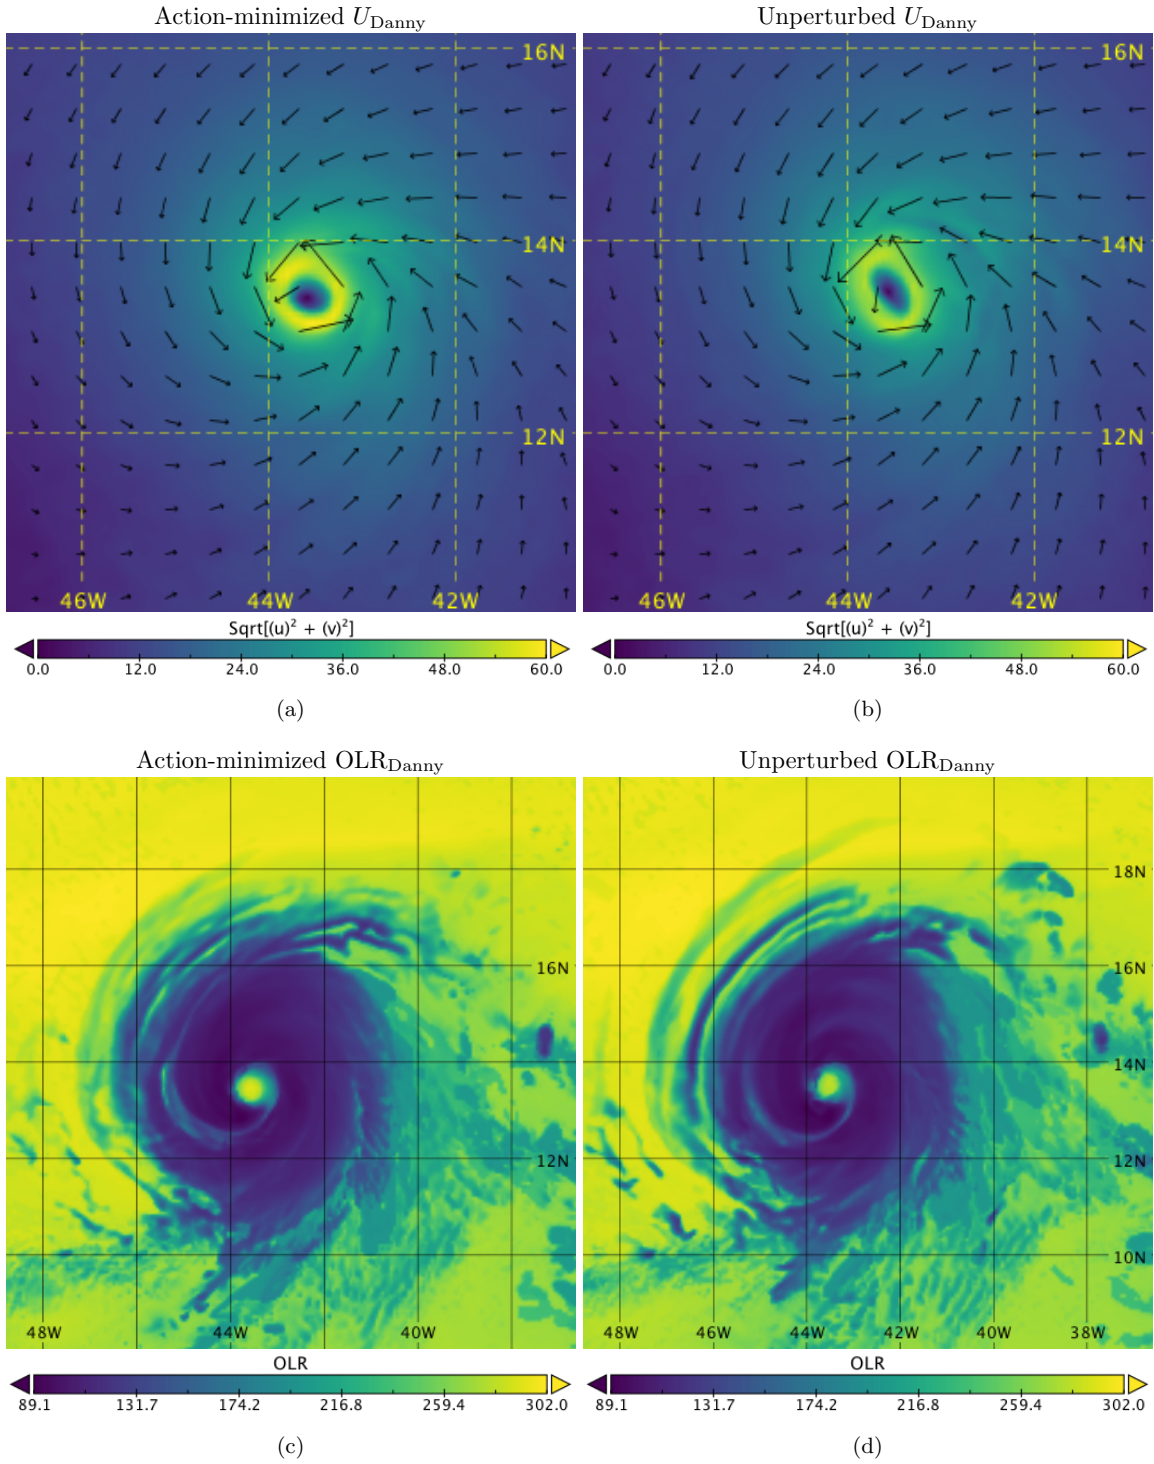


Figure 3.6: Surface wind magnitude and OLR snapshots at final integration time step for Hurricane Danny: optimization output (panels a and c) and forward integration (panels b and d).

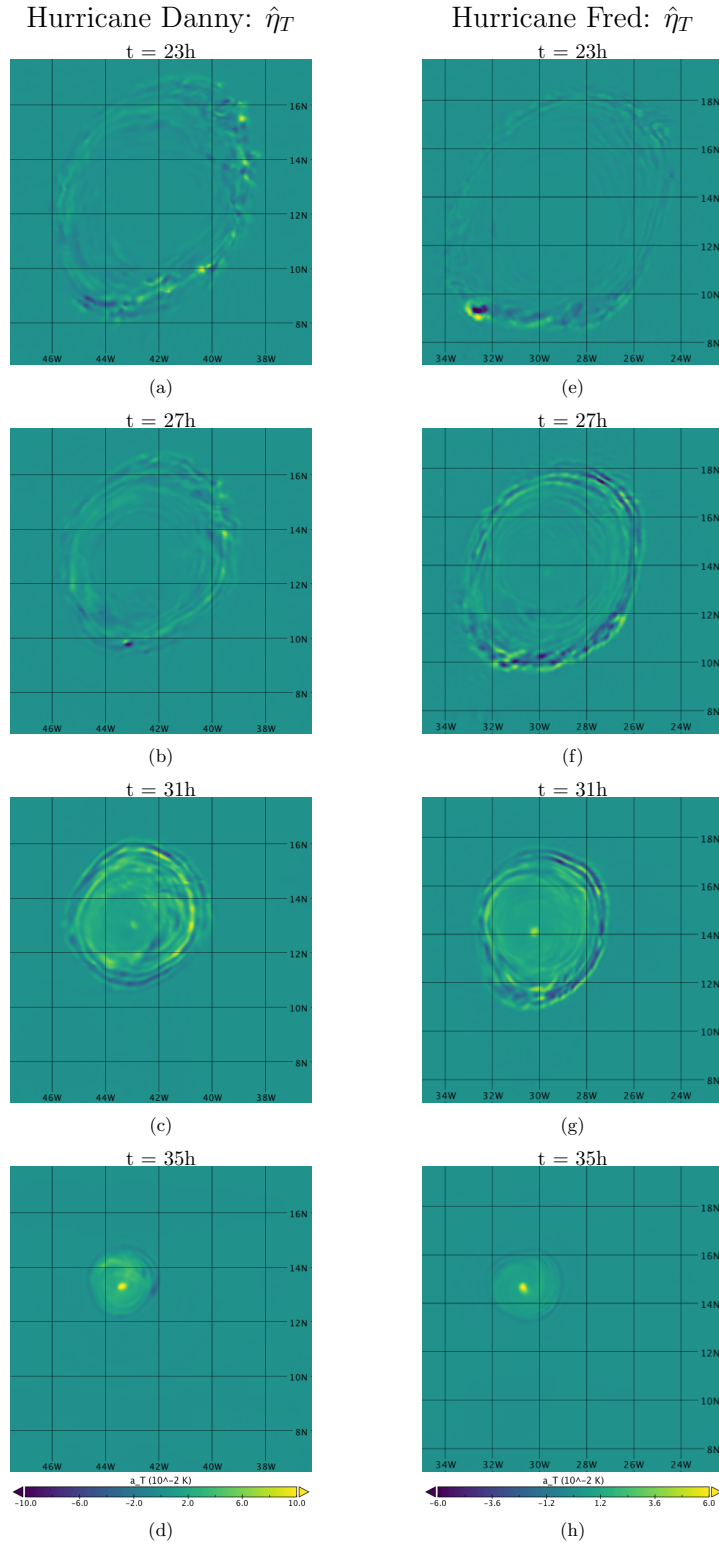


Figure 3.7: Temperature perturbations $\hat{\eta}_T$ at 850 hPa height for action minimization of Hurricanes Danny (panels a-d) and Fred (panels e-h).

This is consistent with Stern and Nolan (2012), who used the WRF model to demonstrate a warm core horizontal length scale of $\mathcal{O}(50\text{km})$ in a mature storm. In both cases, the entirety of the ring of warm air consistently penetrates up to 650 hPa ($z \approx 3 \text{ km}$) with maxima in the 800 – 850 hPa range, but also contains regions that penetrate up to 400 hPa toward the end of the simulations (figure 3.8). The relatively shallow structure of the temperature perturbations provides evidence for the paradigm of Tao and Jiang (2015), who used Tropical Rainfall Measuring Mission Precipitation Radar data to argue that shallow convection, potentially at relatively large radii, often acts as a trigger for RI, while deep convection near the radius of maximum wind arises later as a response to RI. This contrasts with some previous studies, (e.g. Chen and Zhang, 2013) that argue that deep convective bursts precede or even trigger RI.

While individual perturbations are instantaneous and so cannot be advected, we hypothesize that the temperature perturbations are consistent with advection toward the center of the storm by converging boundary layer air in the following sense:

1. A perturbation is added at time t_i at radius r_i .
2. Air parcels at t_i, r_i are advected to a new radius r_{i+1} at time t_{i+1} .
3. The perturbation at time t_{i+1} is added at this new radius r_{i+1} .

In the Lagrangian perspective, the algorithm heats the same parcels of air throughout the time horizon of the simulation; specifically, it heats those parcels that reach the radius of maximum wind at the end of the simulation. Figure 3.9 shows Hovmuller diagrams of 900 hPa ($z \approx 1.5 \text{ km}$) temperature perturbations with radius and time for Hurricanes Danny and Fred. The orange triangles indicate the positions of parcels of air being advected by the radial flow (calculated by integrating the velocity field). In both cases, the alignment between the trajectory of the perturbations and the trajectory of the advected parcels shows that action minimization is choosing to heat roughly the same parcels as they are advected inward toward the center of the storm. During the last five hours of each run, temperature

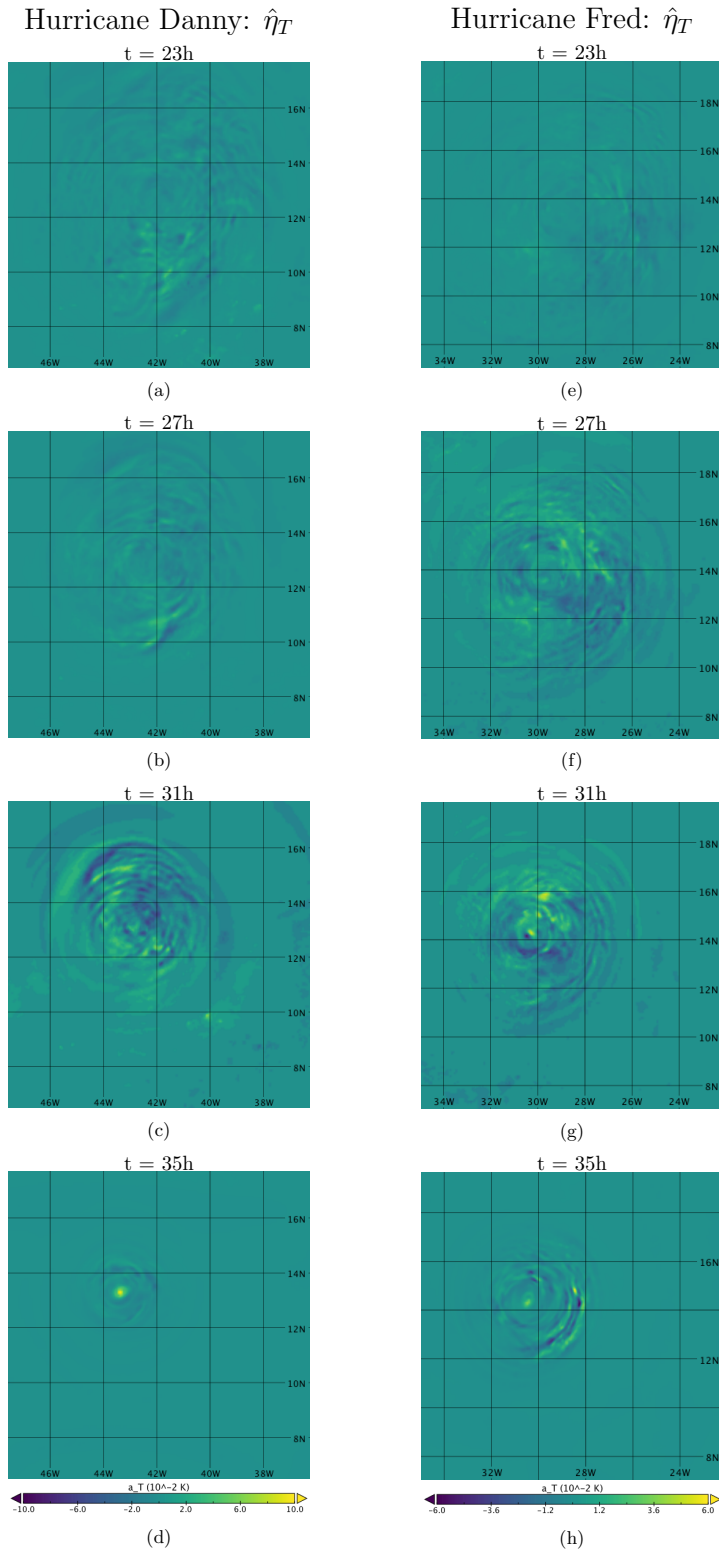


Figure 3.8: As in figure 3.7, but at 500 hPa instead of 850 hPa.

perturbations are dominated by heating near the center of the storm (not shown in figure 3.9 for clarity of the color bar). Thus, even though heating during the early stages of the storm may be relatively inefficient since it occurs outside of the radius of maximum wind (Vigh and Schubert, 2009), the action minimization algorithm finds a path to transport this heat to the center of the storm and efficiently cause intensification toward the end of the simulation. This may explain why the pressure trajectories shown in figure 3.4 show larger divergence from the unperturbed runs during the later stages of the action minimization output.




Figure 3.9: Hovmuller diagrams of 900 hPa temperature perturbations with radius and time for Hurricanes Danny and Fred. Orange triangles indicate the position of an air parcel advected by mean radial flow.

Numerous studies have investigated the roles of symmetric and asymmetric heating in TC intensification (Carr and Williams, 1989; Nolan and Farrell, 1999; Nolan and Montgomery, 2002; Nolan and Grasso, 2003; Nolan, 2007; Shapiro, 2000; Persing et al., 2013). By decomposing temperature perturbations into an azimuthally averaged symmetric component and a mean-zero asymmetric component, Nolan and Grasso (2003) found that the response of a TC wind field can be closely approximated by the purely symmetric response to the symmetric component of the heating. Nolan et al. (2007) confirmed this result, but did find that there are times when vortices with relatively small radii can be slightly strengthened by the asymmetric heating.

While these studies used some temperature and velocity perturbations with time-varying azimuthal structure, the radial structure of the perturbations was constant with time. Further, Nolan et al. (2007) used temperature perturbations with maxima at $z = 6$ km and that decayed to 0 at $z = 3$ km. Such perturbations are consistent with some observational studies of latent heating distributions (Yang et al., 2002) and with modeling studies in which deep convection is the source of heating Hendricks et al. (2004); Montgomery et al. (2006). However, both observational and modeling studies suggest that latent heating in certain storms can be significant at much lower altitudes (Park and Elsberry, 2013; Zagrodnik and Jiang, 2014; Tao and Jiang, 2015). Moller and Shapiro (2005) found that, depending on the nature of pre-existing asymmetries in a vortex, asymmetric heating can cause TC intensification of a similar magnitude (both in pressure and maximum wind) and over a similar time horizon as in this study. Importantly, while there are varied and sometimes contradicting studies regarding the impact of asymmetry and asymmetric heating on TC intensification, these studies all use prescribed perturbations to the temperature and/or PV fields; their conclusions are thus limited to the particular perturbations under consideration. By searching over the full space of possible perturbations with action minimization, this study markedly advances this line of inquiry.

Figure 3.10 shows Hovmuller diagrams of 900 hPa temperature perturbations, averaged over $r \in [50, 150]$ km, with angle (relative to the center of the storm) and time for Hurricanes Danny (panel a) and Fred (panel b). Perturbations for Hurricane Danny take on a clear asymmetric structure 20 hours into the simulation, as indicated by the area of strong positive perturbations that begin rotating around the storm at this time. This contrasts sharply with perturbations for Hurricane Fred, which show little dependence on angle.

To test the significance of the asymmetric component of the temperature perturbations to Hurricane Danny, we first establish that the asymmetry of the perturbations does not serve simply to symmetrize the temperature field. We do so by decomposing the 900 hPa temperature field into symmetric and asymmetric components and then calculating the av-

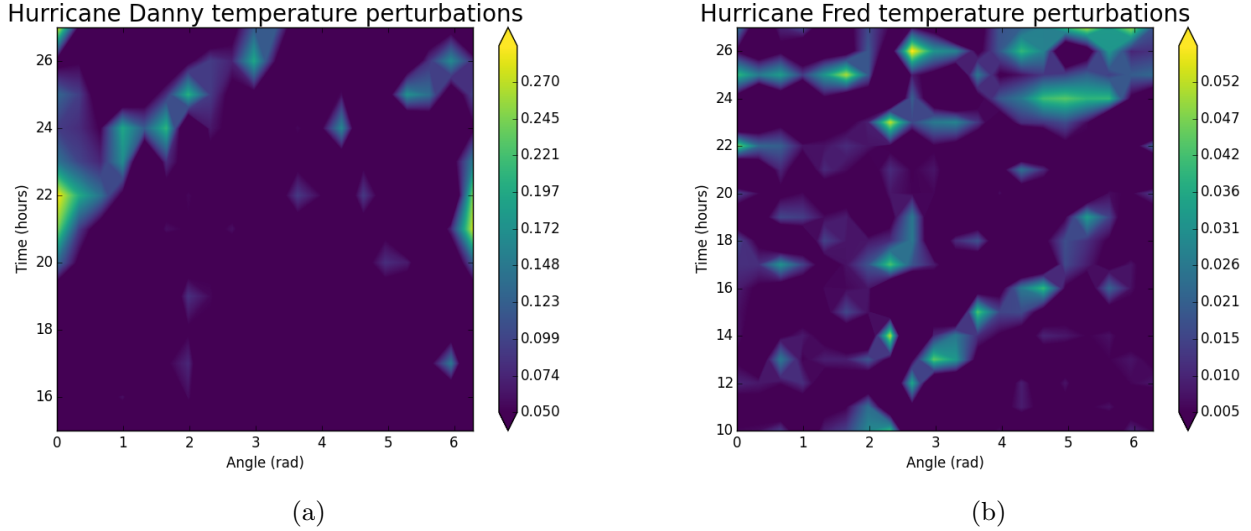


Figure 3.10: Hovmuller diagrams of 900 hPa temperature perturbations, averaged over $r \in [50, 150]$ km, with angle and time for Hurricanes Danny and Fred.

erage magnitude of the deviation from symmetry as a function of radius from the center of the storm. Figure 3.11 shows these deviations from symmetry for both the action minimization output and the unperturbed forward run for Hurricane Danny at $t = 26$ hours (panel a) and $t = 36$ hours (panel b). It shows while that the temperature field becomes more symmetric in the optimization at the final time, it is less symmetric near the radius of maximum wind as compared to the forward run at prior times. The relative asymmetry of the action minimization output is strongest between $t = 20$ hours and $t = 28$ hours; notably, this corresponds to the time at which the minimum pressure in the optimization output is higher than that of the forward integration.

This suggests that the asymmetric component of the temperature perturbations may contribute to intensification in the action minimization output. To test this, we run a forward model and add only the symmetric component of the temperature perturbations to the temperature field at each time step. Figure 3.12 shows that, for Hurricane Danny, this configuration yields a minimum pressure of 964.5 hPa. Given that the full action minimization output and unperturbed run yielded minimum pressures of 962 hPa and 970.5 hPa, respectively, removing the asymmetry results in a 30% reduction in intensification. Notably,

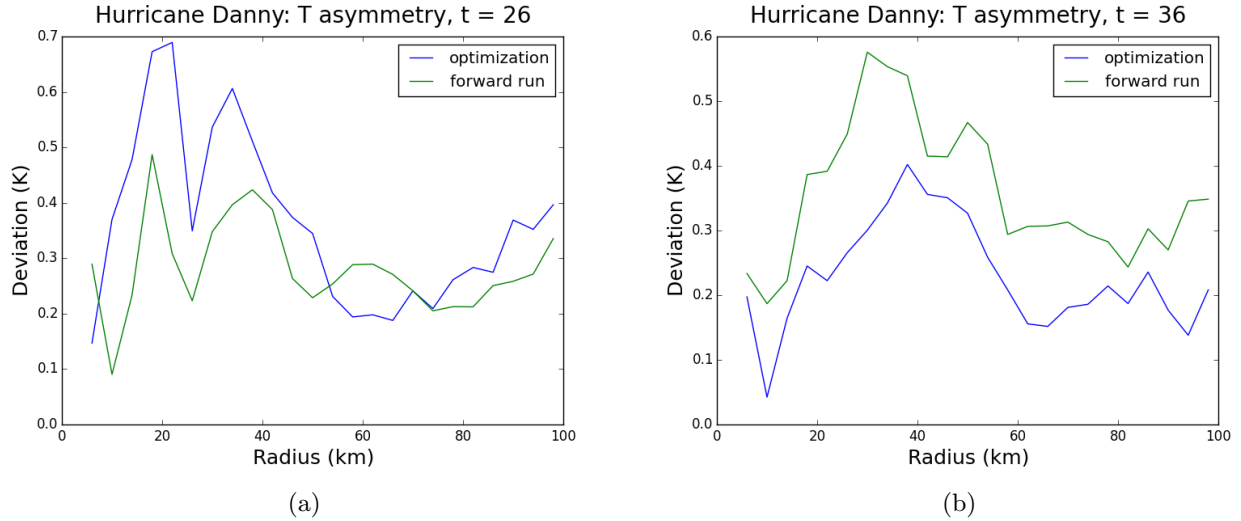


Figure 3.11: Average deviation from symmetry of the 900 hPa temperature field for Hurricane Danny at $t = 26$ hours (panel a) and $t = 36$ hours (panel b).

repeating this test for Hurricane Fred results in no significant change to the minimum pressure relative to the optimization output. We also check the effect of the purely asymmetric component of the action-minimizing perturbations; figure 3.12 shows that forward integration with purely asymmetric perturbations yields a small (< 1 hPa) weakening of Hurricane Danny. This is consistent with Nolan and Grasso (2003) and Nolan (2007). However, the fact that the linear combination of the symmetric and asymmetric perturbations produces stronger intensification than the symmetric perturbations alone suggests a significant non-linearity in storm response to perturbations that was not observed in previous studies.

3.3.4 Velocity response to heating

As with temperature, we decompose the radial and azimuthal wind fields for Hurricanes Danny and Fred into symmetric and asymmetric components. Figure 3.13 shows the average of the magnitude of deviations from symmetry of the azimuthal velocities for both storms as a function of radius. Both storms show more symmetric azimuthal wind fields in the optimization output than in the forward integration. The maximum average asymmetry at the final time occurs near the radius of maximum wind for both storms, which is approx-

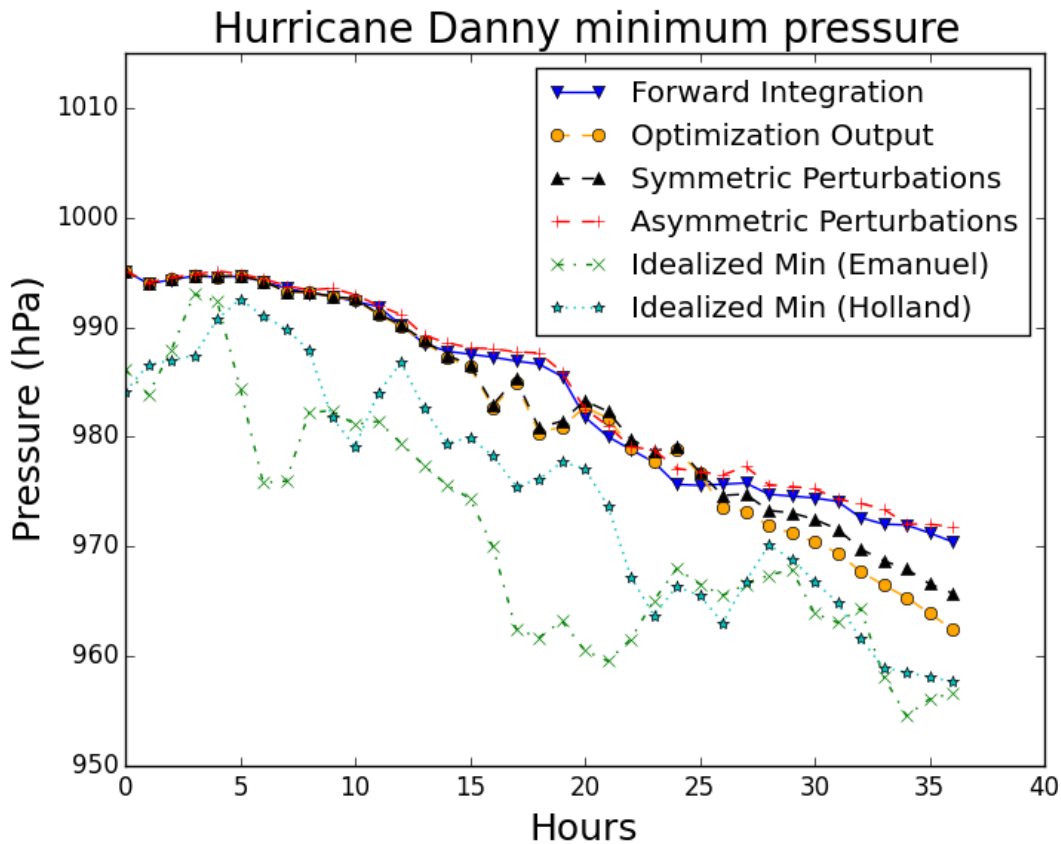
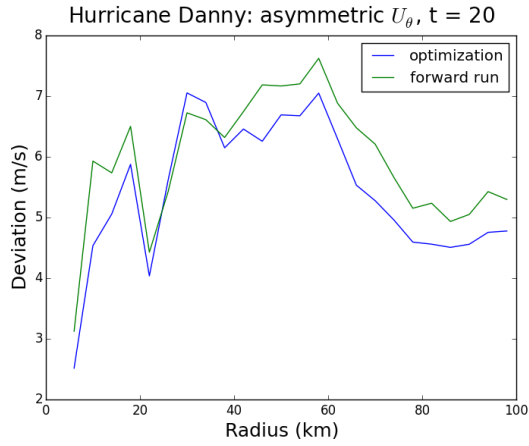


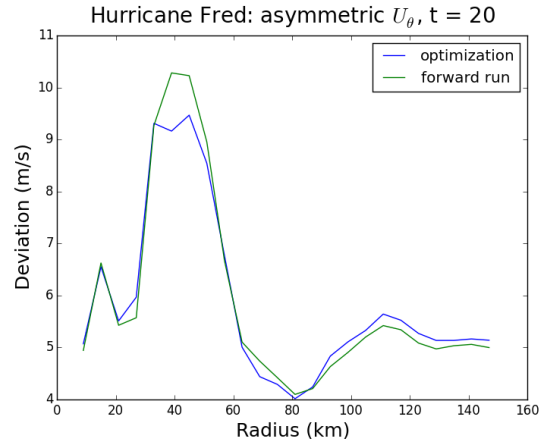
Figure 3.12: As in figure 3.4a, but with the purely symmetric portion of action-minimizing perturbations (black triangles) and with purely asymmetric portion of the perturbations (red crosses) for Hurricane Danny.

imately 30 km for Danny and 50 km for Fred. The deviations from symmetry of the radial wind fields (not shown) have very similar responses to those of the azimuthal fields.


Previous studies have also shown that vertical wind shear is a crucial factor inhibiting TC intensification (Merrill, 1988). DeMaria (1996) notes that the effect of vertical shear is to advect upper level heat and moisture away from the low level circulation, thereby disrupting the vertical structure and suppressing storm development. Jones (1995) argues that vertical shear displaces upper-level potential vorticity (PV) relative to low level PV, which inhibits convection and intensification. Riemer et al. (2010) and Tang and Emanuel (2010) argue that vertical wind shear ventilates mid-level storm cores and introduces cold, dry air, as well as reducing the thermodynamic efficiency of TCs.



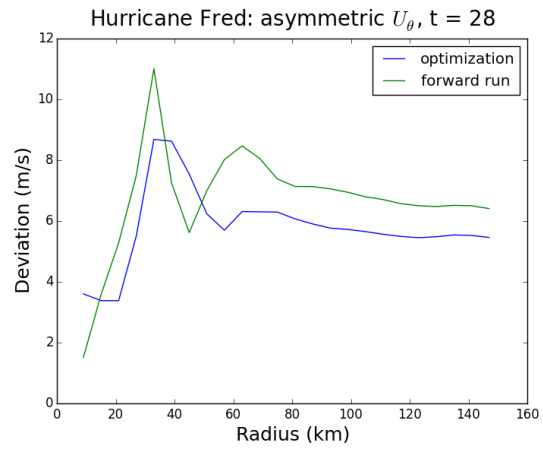
(a)




(d)




(b)



(e)



(c)



(f)

Figure 3.13: Average deviation from symmetry of the azimuthal surface wind field as a function of radius for Hurricane Danny (panels a-c) and Hurricane Fred (panels d-f) at $t = 20, 28,$ and 36 hours.

Following Rappin and Nolan (2012) and Zhang and Tao (2013), we define low-level vertical shear as the difference between horizontal winds at 900- and 500-hPa levels averaged over a circle marking the edge of the storm, where the edge is the radius of 8 m s^{-1} wind (Schenkel et al., 2017); we define high-level shear analogously but using the 850- and 250-hPa levels (Frank and Ritchie, 2001; Corbosiero and Molinari, 2003). Figure 3.14 shows time series of both low- and high-level vertical shear in the action minimization output and unperturbed forward integration for Hurricanes Danny and Fred. Both storms show a reduction in low-level shear in the action minimization output with the strongest effect towards the beginning of the run. The reduction is approximately five times larger for Hurricane Fred: initial time shear is reduced from 6.11 m s^{-1} to 5.08 m s^{-1} for Hurricane Fred and from 3.79 m s^{-1} to 3.60 m s^{-1} for Hurricane Danny. This is consistent with Wong and Chan (2004), who found that TC intensity is most sensitive to changes in shear above a threshold that lies between $2 - 4 \text{ m s}^{-1}$. Hurricane Danny, which formed in an abnormally low-shear environment, is already at or near this threshold in the unperturbed run, so the algorithm does not waste energy (running cost) reducing its shear. Further, the action-minimizing perturbations reduce the shear as defined by the difference between 900- and 500-hPa level velocities, but not as defined by the difference between 850- and 200-hPa level velocities; this indicates that the low-level shear may be more important for TC intensification in WRF. The effect of the reduction in low-level shear is to reduce the tilt, as defined by the difference between the location of the storm center at the 900- and 500-hPa levels (Tao and Zhang, 2014), from 9.6 km to 6.1 km.

Finally, we check that the reduction in wind shear does not immediately dissipate after the perturbations are added by calculating the magnitude of the shear every 15 minutes in order to provide data between perturbations; we do this both for low-level (900-500 hPa) and high-level (850-250 hPa) shear. Were the perturbations transient, we would expect to see a sawtooth pattern in the perturbed wind shear in figure 3.14, where shear gets nudged lower each hour but then quickly returns back to the shear observed in the forward integration.

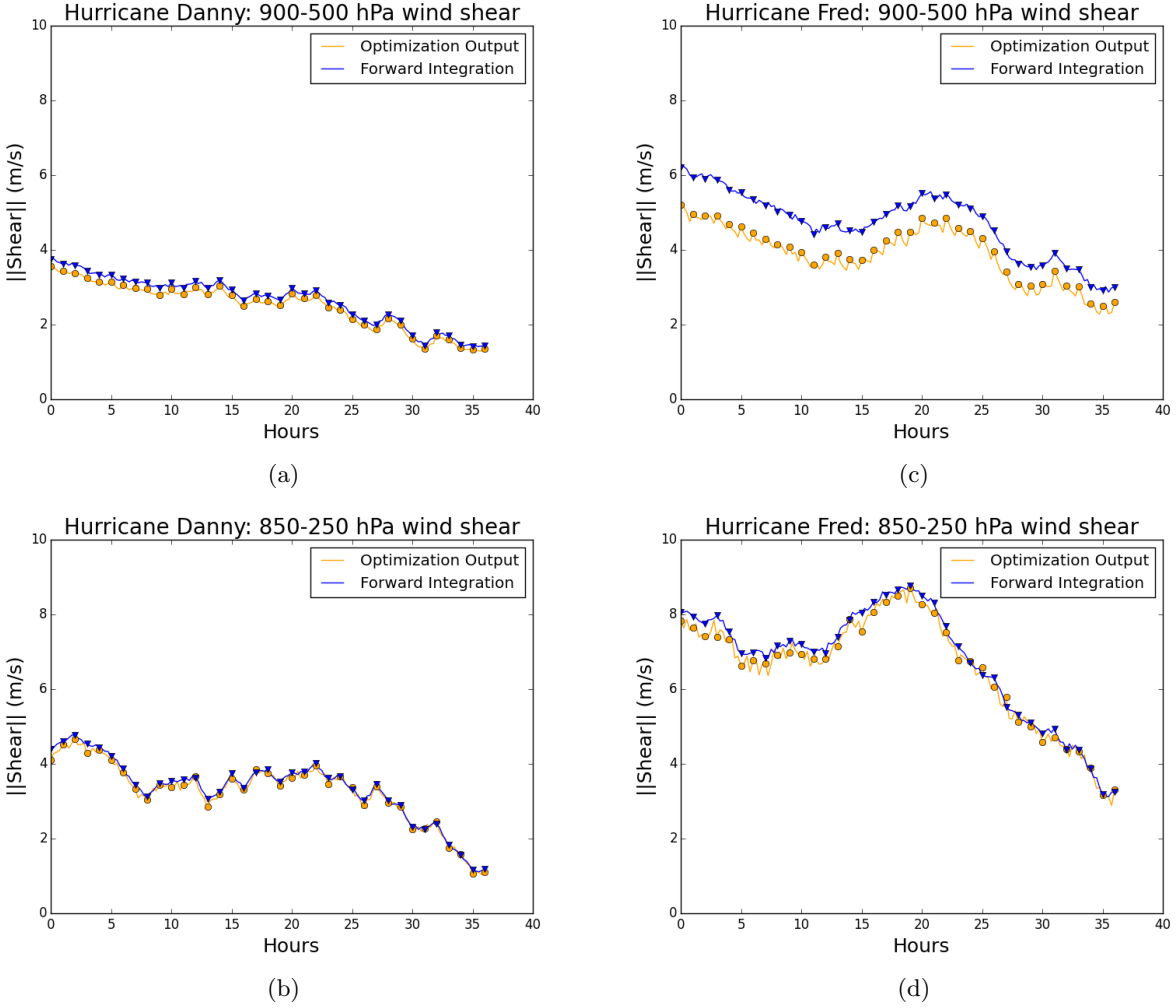


Figure 3.14: Magnitude of 900 - 500 hPa and 850 - 250 hPa wind shear as a function of time for Hurricanes Danny (panels a and b) and Fred (panels c and d) for forward integration (blue curve) and optimization output (orange curve). Markers denote times at which perturbations are added by action minimization.

Instead, low-level shear is consistently lower in the optimization output, confirming that the perturbations to shear are not transient in nature.

3.3.5 Water vapor perturbations

Panels a and b of figure 3.15 show initial time water vapor perturbations and initial time relative humidity at the surface for Hurricane Danny, respectively. As expected, the perturbations are generally positive and seem to be larger where relative humidity is smaller. In fact,

the Pearson product-moment correlation coefficient between $\hat{\eta}_{q_{vapor}}$ and RH is $r = -.61$, confirming that the two are strongly anticorrelated (Rodgers and Nicewander, 1988b). Combined with the fact that $\hat{\eta}_{q_{vapor}}$ falls off by nearly two orders of magnitude at distances greater than 500 km from the storm center, this suggests that the major function of these perturbations is to moisten the near-storm environment. The water vapor perturbations are qualitatively similar for Hurricane Fred both in that they are confined to the boundary layer and in that they are strongly anticorrelated with relative humidity ($r = -.54$).

Panel a of figure 3.15 suggests that water vapor perturbations are slightly larger downstream (to the west) of the storm center. One may expect this to be the case, as perturbations that preferentially moisten the part of the domain that the storm will shortly traverse can lead to more latent heating once the storm reaches that part of the domain. Although covariance matrix B contains only instantaneous correlations, the derivative $F'(\cdot)$ can contain dynamic information including the fact that water vapor at an early time in the simulation may become latent heating later on.

To check whether or not action minimization preferentially adds water vapor downstream of the storm, we calculate the average relative humidity in a 500×500 km box around the storm center in both the action minimization output and the unperturbed integration (Chen and Gopalakrishnan, 2015) and take the difference of the two as a function of time. We then repeat this, but using only the downstream half of the domain. Panels c and d of figure 3.15 show these differences in near-storm relative humidity for Hurricanes Danny and Fred, respectively. In both cases, the downstream only average relative humidity shows a larger increase compared to the unperturbed run at most times, suggesting that moistening this portion of the domain is relatively more important for intensification.

It is important to note that, compared with temperature and wind perturbations, the action-minimizing water vapor perturbations $\hat{\eta}_{q_{vapor}}$ are fairly small. They comprise approximately 5% and 3% of total running cost for Hurricanes Danny and Fred, respectively (by comparison, temperature perturbations comprise 63% and 45% and wind perturbations

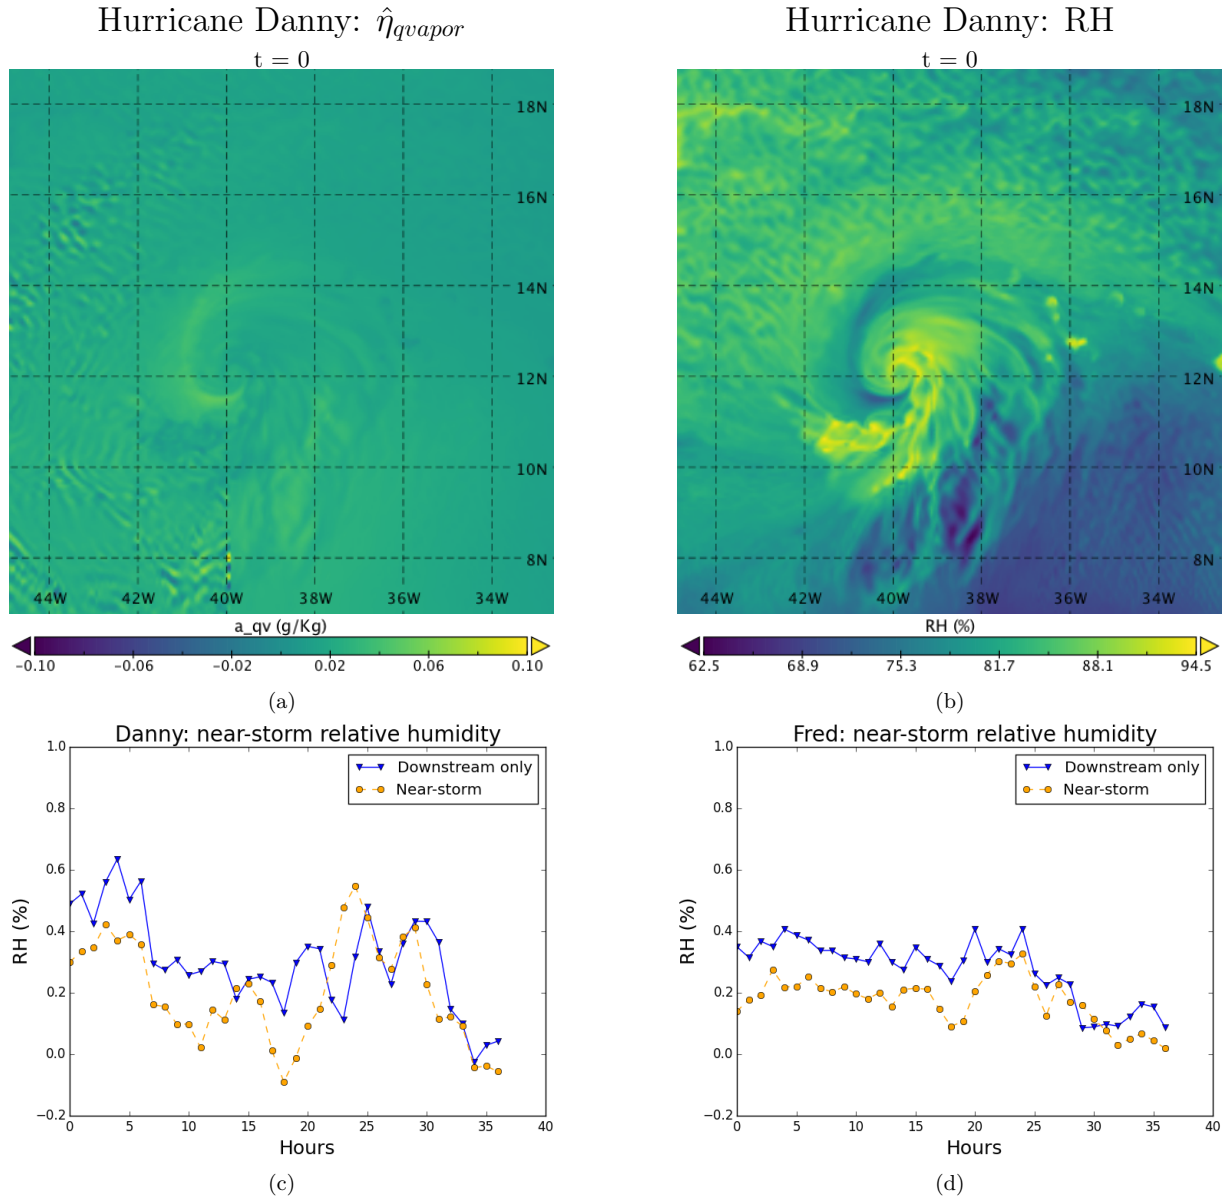


Figure 3.15: Initial time water vapor perturbations (panel a) and initial time relative humidity (panel b) at the surface for Hurricane Danny, as well as time series of near-storm relative humidity differences (optimized minus unperturbed) for the full near-storm domain and only the downstream portion of the domain for Hurricane Danny (panel c) and Hurricane Fred (panel d).

comprise 17% and 28% of running cost from Danny and Fred, respectively). This is likely the case because the perturbations are limited almost entirely to the boundary layer (under the 900 hPa level). Only very small amounts of water vapor could be added higher in the atmosphere before reaching saturation (beyond which point the additional water vapor

would immediately rain out without affecting storm development); however, such perturbations would incur non-negligible cost since the covariance matrix B accounts for water vapor content decreasing with height. We hypothesize that the algorithm avoids this cost in favor of confining water vapor perturbations to the boundary layer.

To test the importance of the water vapor perturbations, we integrate the model but set $\hat{\eta}_{q_{\text{vapor}}} = 0$ at all times. This results in final minimum pressures of 962.7 hPa for Hurricane Danny and 943.7 hPa for Hurricane Fred (compared to 962.0 hPa and 943.3 hPa, respectively, when using the full perturbations). In both cases, omitting the water vapor perturbations reduces intensification by well under 10% compared to the full perturbation case. This shows that, at least for our choice of covariance matrix B , perturbations to water vapor are relatively less efficient than perturbations to the temperature and wind fields.

3.3.6 *A note on energetics*

In the preceding discussion, we have used small perturbations to temperature aloft as a proxy for convective activity. While this interpretation is consistent with numerous studies that impose temperature perturbations in the context of TC simulation (e.g. Nolan and Grasso, 2003; Nolan, 2007; Shapiro, 2000), it may be energetically suspect. This is due to the fact that temperature perturbations themselves do not necessarily conserve energy (i.e. there could be a net input of energy), while convection is a conservative process that simply converts latent heat to sensible heat.

To investigate the energetics at play, we calculate the energy content of both the action minimization output and the unperturbed storm for Hurricane Danny. The results are displayed in table 3.1, which includes the kinetic energy ($\frac{1}{2}|\mathbf{v}|^2$), sensible heat energy ($C_p T$), and latent heat energy ($L_v q$) at the final time for both the action-minimized storm and the unperturbed storm. Gravitational potential energy is omitted as it does not significantly differ in the two cases. Importantly, table 3.1 shows only the energy content of the hurricane, defined horizontally by the radius of 8 m s^{-1} wind and vertically by the 200 hPa height.

Table 3.1: Energy budget for simulations of Hurricane Danny ($J \times 10^{-16}$).

	Action minimization	Original storm	Difference	AM input
$\frac{1}{2} \mathbf{v} ^2$	426	408	18	8
$C_p T$	13500	13485	15	37
$L_v q$	73022	72993	29	3
Total	86948	86886	62	48

First, we note that kinetic energy associated with wind comprises only about one half of one percent of the total energy content of Hurricane Danny; this is consistent with previous findings (Emanuel, 1999). Despite this, the difference in kinetic energy between the action-minimized and unperturbed storms comprises approximately 30% of the total energy difference between these storms. This suggests that action minimization is relatively efficient at decreasing the minimum pressure and therefore increasing the wind speed of Hurricane Danny; that is, action minimization allocates energy in a way that is effective at increasing the pressure gradient and driving winds.

It is plausible that action minimization simply strengthens Hurricane Danny simply by adding energy to the system. While action minimization does add 4.8×10^{17} J to the final state, it is important to note that this does not account for the entire energy difference between the action-minimized and unperturbed storms. Instead, there are 1.4×10^{17} J of additional energy content in the action-minimized storm that are not accounted for by direct energy input from the algorithm itself. Further, when one considers the entire domain at the final time, rather than just the domain covered by the storm, the energy difference between the action minimization output and the unperturbed simulation decreases from 4.8×10^{17} J to 2.9×10^{17} J. This implies that action minimization does not only add energy to the system, but it also reorganizes the energy within the domain in such a way as to strengthen Hurricane Danny. Additional energy content in the action-minimized storm may also be a result of increased water vapor flux from the ocean surface due to warming of boundary-layer air via temperature perturbations.

Finally, it is important to note that the WRF model does not perfectly conserve energy, especially for moist processes (Kniewicz, 2006). In the unperturbed simulation of Hurricane Danny, the domain loses approximately 6.5×10^{16} J over the course of the 36-hour time horizon. This is less than one order of magnitude less than the energy differences discussed in the previous paragraph; that discussion is thus necessarily tempered by the fact that the relevant energy differences are not much larger than the model error.

In sum, action minimization does not simply strengthen Hurricane Danny by adding energy. Rather, it also reorganizes the energy in the domain around Hurricane Danny in order to improve conditions for intensification. Further, it is plausible that action minimization does increase convective activity by increasing the water vapor content of boundary layer air; however, the interpretation of temperature perturbations as a proxy for convection should be re-examined in future work. A thorough analysis of the energetics of action minimization would benefit from further development that would improve the WRF model's ability to conserve energy during moist processes.

3.3.7 Parameter sensitivity

The most important parameter choices in this study are: 1) the frequency of perturbations, 2) the relative weights, R_i and R_f , of the running cost and the final cost, and 3) the time horizon. We choose to apply perturbations once an hour over the course of a 36 hour time horizon; however, the results are robust to changes in both of these choices. We show this by varying both time horizon and perturbation frequency at 30 km resolution and showing that, when the parameters are varied within a reasonable range, both the magnitude and mechanism of intensification remain very similar.

Figure 3.16 shows Hovmuller diagrams of 900 hPa temperature perturbations with radius and time for 30 km resolution simulations of Hurricane Danny; the perturbation frequency varies from 90 minutes to 15 minutes. While the diagrams are not identical, they are qualitatively similar; namely, all show the movement of warm temperature perturbations into the

center of the storm with roughly the same velocity as an advected air parcel. The magnitude of the heating at each time step decreases with perturbation frequency so that the total amount of heating remains approximately constant. Further, all of these simulations result in intensification of the pressure minimum by 9.1 – 9.4 hPa. This suggests that both the magnitude and mechanism of intensification is robust to changes in perturbation frequency.




Figure 3.16: Hovmuller diagrams of 900 hPa temperature perturbations with radius and time for Hurricane Danny at various perturbation frequencies. Orange triangles indicate the position of an air parcel advected by mean radial flow.

We also test the sensitivity of the algorithm to the weights R_i and R_f . Since running cost goes as $1/R_f^2$, higher choices of R_f mean relatively less weight is placed on the running

cost (and therefore more on the final cost), meaning that we expect higher values of R_i/R_f to result in more intensification. Panel a of figure 3.17 shows the pressure drop of the action minimization output compared to unperturbed forward integration as a function of the ratio of weights R_i/R_f . As expected, the pressure drop increases roughly as the square of R_i/R_f . In practice, the high resolution results were obtained by choosing R_i/R_f to be as large as possible while maintaining model stability.

Finally, we run the model for several different time horizons (with the same end time but varying start times). Panel b of figure 3.17 shows the pressure drop of the action minimization output compared to unperturbed forward integration as a function of time horizon. The magnitude of the pressure drop is nearly constant for time horizons of 36 hours or longer; at 24 hours and below, intensification decreases drastically with time horizon. Further, temperature perturbations for varying time horizon runs (not shown) have the same qualitative behavior as those in figure 3.16. This suggests that there is a critical time horizon of approximately 36 hours below which the intensification mechanisms found by action minimization do not have time to maximize intensification; above this threshold, the intensification mechanism and magnitude are relatively unchanged as a function of time horizon.

3.4 Discussion

In this study, we applied an action minimization algorithm to nudge the WRF model into forming more intense TCs than it otherwise would have. For Hurricanes Danny and Fred, our method results in reductions of the minimum pressure of nearly 10 hPa and corresponding increases in maximum surface wind speed of $6 - 7 \text{ m s}^{-1}$. Action minimization a) causes both storms go from category 3 to category 4 on the Saffir-Simpson scale, and b) ensures that both storms undergo RI. In both cases, the nudged model trajectories approach but do not exceed the maximum potential intensities predicted by thermodynamic constraints. This shows that action minimization can consistently result in more intense TCs in weather model




Figure 3.17: Sensitivity of intensification to the ratio of weights R_i/R_f (panel a), and sensitivity of intensification to time horizon (panel b).

output in varying ambient conditions and that these TCs are thermodynamically feasible.

Relative to category 3 TCs, category 4 TCs are approximately 4 times less common and category 5 storms are nearly 20 times less common. In normal climate/weather simulations, this presents a challenge to the study of category 4 and 5 TCs owing to the large amount of simulation time (and therefore computational expense) needed to observe these storms. An algorithm that consistently causes category 2-3 TCs to further intensify to category 4 or 5 can greatly reduce the computational expense of studying the most intense TCs. This, in turn, would allow for more systematic study of category 4 and 5 TCs in climate models. Since we have shown that action minimization reduces the computational cost of simulating events at the tail of TC intensification distribution by approximately an order of magnitude, we argue that action minimization is a viable tool for studying the most intense TCs.

For both Hurricanes Danny and Fred, the action-minimizing perturbations to the forward model introduce low-level warm temperature anomalies that move into the warm cores of the storms in a manner consistent with advection. These anomalies have maxima at approximately $z = 2$ km, and suggest that low-level convection may be a preferred mechanism of TC intensification in the WRF model. This provides evidence for recent observational

studies that argue that shallow to moderate convection can be a trigger for RI, while deep convection is often a response to RI.

The temperature perturbations for Hurricane Danny are asymmetrical in such a way as to make the temperature field more asymmetrical during the middle stages of the optimization. As expected, the vortex axisymmetrizes by the end of the optimization time horizon. However, asymmetry of perturbations contributes to intensification since integrating the model with only azimuthally averaged perturbations results in an approximately 30% reduction in intensification as compared to integrating using the full perturbations. Further, and in agreement with Nolan (2007), integrating with only the purely asymmetric component of the perturbations results in a slight weakening of the vortex. Our results differ markedly from that study because, at least for the specific set of perturbations produced by action minimization, both asymmetry and non-linearity of the response of the vortex to warming significantly increase intensification.

Our results for Hurricane Danny may differ from those of previous studies due to the specific nature of perturbations applied to the vortex. The significance of nonlinearity may arise from the fact that the asymmetric velocities in this study are above the “limit of where we can expect linear calculations to be accurate” described in Nolan (2007). Further, the action-minimizing perturbations differ from those used in previous studies in that they a) are centered at lower altitudes, b) have time-dependent radial structure, and c) have azimuthal structure that is not constrained by a particular functional form. This is a major contribution of the action minimization algorithm, as it is able to search over the space of all possible perturbations whereas previous studies and their conclusions were limited to the specific perturbations that they imposed. Finally, the fact that asymmetry of perturbations does not significantly impact the intensification of Hurricane Fred suggests that the role of asymmetric heating depends on pre-existing asymmetry and storm structure; this result is consistent with Moller and Shapiro (2005).

For both Hurricanes Danny and Fred, the action-minimizing perturbations also reduce low-level environmental wind shear. The decrease in low-level shear is approximately $\sim 1 \text{ m s}^{-1}$ for Fred at the beginning of the integration and is about five times larger than the reduction in shear for Hurricane Danny. This may be because the unperturbed Danny integration already starts below a postulated threshold near 4 m s^{-1} below which reduction in shear has little effect. While previous studies (e.g. Chen and Zhang, 2013) suggest that environments where shear is too low ($< 5 \text{ m s}^{-1}$) are actually detrimental to TC intensification as compared to moderate shear ($5 - 10 \text{ m s}^{-1}$) environments, our results are more in line with those of Nolan and McGauley (2012), who show that extremely low-shear environments are no less favorable to intensification than are moderate-shear environments but that the effect of reducing shear gets smaller as the magnitude of the shear gets smaller.

Finally, action-minimizing perturbations to water vapor moisten the near-storm environment. These perturbations are strongly anticorrelated with relative humidity, and are almost entirely confined to the boundary layer. Further, relative humidity is preferentially increased downstream of the storm in the action minimization output. Somewhat surprisingly, the perturbations to water vapor are relatively small compared to perturbations to the wind and temperature fields, both in terms of the running cost allocated to them and in terms of their effect on the final minimum pressure.

While these results show that the trajectories resulting from action minimization can elucidate the physical trajectories by which TCs undergo rapid intensification in the WRF model, our work has several important limitations. First, action minimization requires a specific model and can only contribute to understanding of trajectories/processes within that model. If a model is incapable of simulating the true process by which an event occurs, action minimization will not yield that true trajectory; instead, action minimization should be thought of as a method for efficiently accessing and studying tail events that a model can produce.

Second, the choice of covariance matrix B may constrain the output that action mini-

mization can produce. Specifically, we generate B from long model runs that mostly access quiescent states. The flows in these states are in balance and relatively smooth, so it is possible that our choice of B imposes smoothness and overly penalizes large gradients. Such gradients are uncommon in most weather states but prevalent under TC conditions. While the choice of B poses some difficulty, we also note that any ensemble algorithm also requires an analogous choice in how to generate ensemble members that diverge from one another. This is true even for ensembles where only the initial condition is varied, as choosing the different initial conditions requires drawing them from a distribution that is characterized by some covariance matrix B .

Third, and as with ensemble methods, diagnosing the causes of rapid TC intensification requires an analysis of trajectories. If one has a putative mechanism in mind, one can assess the degree to which the action-minimized path agrees with that mechanism, as in our analysis of the effects of asymmetries in the perturbations. Alternatively, one can systematically investigate a potential mechanism by modifying the background covariance matrix to penalize certain classes of perturbations (e.g. asymmetric perturbations) and compare the resulting trajectory to the trajectory found with the original background covariance. Identification of new mechanisms from the trajectories resulting from action minimization (or ensemble simulation) is a more difficult problem. It can, in principle, be approached using standard data analysis tools like PCA (Bretherton et al., 1992b) or diffusion maps (Coifman et al., 2008; Giannakis and Majda, 2012; Plotkin et al., 2014). However, the data are both high dimensional and time dependent and effective analysis would benefit from the development of new tools capable of distinguishing important features of the trajectories.

Fourth, the optimal perturbations could depend on both resolution and on choice of model. While the results presented in section 3.3.7 show little qualitative dependence on resolution, the action-minimizing perturbations could change at a resolution that allows the model to resolve convection. In both this study and in previous work, temperature perturbations have often been considered as a proxy for latent heating that results from

convection. We are not aware of studies that impose convective perturbations in a more direct way, so it is difficult to predict how doing so would affect the results. Further, the choice of model and of parametrizations (convective and boundary layers schemes, etc.) are significantly limited in this study, since WRF is one of few publicly available climate or weather models with adjoint capabilities. Within the WRFPLUS adjoint scheme, only several state variables can be perturbed and only a small subset of all parametrizations are available. It is possible that applying action minimization to a different model with different parametrizations, for example to a coupled GCM, would yield different results. This study is necessarily constrained by the imperfections of the WRF model and by the approximations inherent in the adjoint scheme.

The introduction of action minimization into the repository of tools for the study of rapid intensification calls for several directions of future work. First, we intend to apply action minimization to a larger sample of historical TCs to better study intensification mechanisms. Such a sample will show what it would have taken to cause historical storms to further intensify and will allow us to quantify whether or not certain pathways of intensification are consistently likelier than others. By shedding light on the preferred intensification pathway(s), we will also improve understanding of the crucial physics that may be missing from climate models and precluding them from forming sufficiently intense TCs.

By asking what it takes to cause further intensification, both this study and the ones proposed above seek to elucidate the processes that are sufficient for rapid intensification. Instead, action minimization can also be used to investigate necessary processes/conditions. To do so, we can modify the action functional used in this study to instead prefer final states with quiescent weather (instead of intense TCs). By doing so, we ask the converse of the question that we study here, namely: what would it take to prevent intensification? By combining the study of both questions, we can provide insight into both the necessary and sufficient processes required for rapid intensification in the WRF model.

Finally, action minimization can be used to study TCs under climate change. With the

advent of GCMs with adjoint capabilities (Marotzke et al., 1999), it is possible to apply action minimization to climate models. Specifically, Vanden-Eijnden and Weare (2013b) present algorithms that allow ensembles created with action minimization to yield statistical information about the unperturbed system dynamics; this is done by transforming the distributions (e.g. of intensification rates) arising from the action minimization ensembles to distributions that would be observed in an ensemble created via forward integration. By favoring cyclogenesis and rapid intensification, a GCM with action minimization could create a large ensemble of TCs without spending the vast majority of computational time simulating quiescent/non-TC states. This would reduce the computational expense of using GCMs to create large ensembles of strong TCs under different climate scenarios. The benefits of this approach are twofold: first, the computational savings could be used to increase resolution. Second, such ensembles would provide sufficient sample sizes to generate robust statistics about the response of TC frequency and intensity distributions to climate change.

3.5 Conclusions

In this study, we apply an action minimization algorithm to Hurricanes Danny and Fred using the WRF and WRFPLUS models to reach the following conclusions:

1. Action minimization can consistently cause TCs in the WRF model to intensify beyond their maximum intensities attained via forward integration.
2. Compared to ensemble simulation, action minimization yields computational savings of approximately an order magnitude in accessing extreme events at the tail of the TC intensification distribution.
3. Model trajectories output by action minimization are consistent with thermodynamic constraints, as evidenced by the fact that these trajectories have TCs intensifying close to but not beyond their maximum potential intensities.

4. Intensification pathways include low-level heating, possibly consistent with shallow convection, that moves toward the center of the storm. Depending on pre-existing storm structure, the asymmetry of the heating can play a non-negligible role in intensification. Further, intensification responds nonlinearly to the symmetric and asymmetric components of this heating.
5. Action minimization preferentially reduces low-level vertical shear compared to high-level vertical shear, especially when low-level shear is above a threshold of $\sim 5 \text{ m s}^{-1}$. This leads to intensification by reducing storm tilt.
6. Action minimization has several applications in the study of TCs, including a) the study of rapid intensification mechanisms and b) reduction of the computational cost required to generate high resolution ensembles measuring the response of TCs to climate change.

In sum, we have demonstrated the feasibility of action minimization as a tool for the study of TC rapid intensification. Further, we have provided evidence for the roles low-level wind shear and of low-level, asymmetric heating in rapid intensification. Through direct continuation of this study and through further application of this algorithm to climate models, action minimization could aid in the investigation of rapid intensification.

CHAPTER 4

CONCLUSIONS

In this thesis, we have applied two distinct rare events methods to ocean and atmosphere systems for the first time. In Chapter 2, we used diffusion maps and spectral clustering to improve our understanding of the bimodal states of the Kuroshio current off the coast of Japan. Specifically, we found that:

- The Kuroshio current is bimodal, but the nature of the bimodality is different from what has been described in previous studies. Specifically, the small meander and large meander are best described by path variability and not by path location. The small meander is localized near the coast, while the large meander spreads out and allows for highly variable paths.
- Diffusion maps and spectral clustering provides an index that can be used to measure whether the Kuroshio system is in the small or large meander at any given time (and how strongly it is in a given state). By comparing this to other indices, we can correlate the meander state of the Kuroshio with other physical features of the system. We find that the small meander state corresponds to a nearby recirculation gyre moving to the northeast, while the large meander corresponds to that gyre moving to the southwest. This provides evidence for a meander transition mechanism based on potential vorticity transport causing gyre strengthening and instability, thereby leading to the large meander state. Further, we find that the small meander corresponds to high inflow and that the large meander corresponds to low inflow.
- By comparing the diffusion maps and spectral clustering index to the coefficients produced by principle component analysis, we find that the bimodality of the Kuroshio is also the dominant form of dynamic variability in the system.

In Chapter 3, we apply action minimization to the WRF model in order to study the rapid intensification of tropical cyclones. Action minimization is a new method within the

geophysical sciences that allows a model to be systematically perturbed into more readily accessing rare states of interest. By applying action minimizations to simulations of Hurricanes Danny (2015) and Fred (2009), we show that:

- Action minimization can consistently cause tropical cyclones in the WRF model to intensify beyond the typical maximum intensities attained via forward integration. Specifically, it causes both Hurricanes Danny and Fred, which were both on the borderline of meeting the rapid intensification criterion in unperturbed forward integration, to clearly undergo rapid intensification. When compared to ensemble simulation with the stochastic model implied by action minimization, action minimization yields a computational savings of approximately an order of magnitude in accessing extreme events at the tail of the TC intensification distribution.
- The intensification pathways selected by action minimization, including low-level heating and reduction of low-level wind shear, are physically realistic. Further, model trajectories output by action minimization are consistent with thermodynamic constraints, as evidenced by the fact that these trajectories have tropical cyclones intensifying close to, but not beyond, their maximum potential intensities.
- Asymmetric heating can play a non-negligible role in intensification. Specifically, removing the asymmetric component of the action-minimizing temperature perturbations for Hurricane Danny results in a $\sim 30\%$ reduction in intensification. Further, we discover a regime of non-linear storm response to asymmetric heating that has not previously been studied.
- Action minimization suggests that low-level vertical shear plays a larger role in inhibiting tropical cyclone intensification as compared to high-level vertical shear, especially when low-level shear is above a threshold of $\sim 5 \text{ m s}^{-1}$. It further suggests that there is not necessarily a lower threshold below which vertical shear is actually beneficial to the intensification process.

By applying machine learning and rare event methods to problems in weather and climate, we have set the stage for further development and application of these methods in the field. Algorithms such as diffusion maps and spectral clustering can be applied to any system that exhibits rare transitions, while action minimization can be applied to any such system for which there is an accompanying dynamical model. Potential applications in weather and climate are numerous, including but not limited to the study of extreme heat, droughts, floods, extreme cold, abrupt climate change, and the appearance and disappearance of zonal jets in the Jovian atmosphere. Given the wealth and importance of potential applications of rare events methods in the geophysical sciences, we hope that this thesis sets a precedent for the widespread use of these techniques in future studies.

APPENDIX

DIFFUSION MAPS AND SPECTRAL CLUSTERING

DMSC is a method that reduces the dimensionality of a high-dimensional time series by considering the data as a random walk. Specifically, each point in the time series is stripped of the time dimension and is then plotted in state space. The data are then treated as a random walk, with the transition probability between any two points \mathbf{x}_i and \mathbf{x}_j given by

$$p_{ij} \propto e^{-\frac{|\mathbf{x}_i - \mathbf{x}_j|^2}{2\epsilon_i \epsilon_j}}. \quad (\text{A.1})$$

where ϵ_i and ϵ_j are local scale parameters. The transition probability between any two points decays exponentially with the distance between the points. In this study, we normalize the variance of the list of pairwise distances between data points, d_{ij} , to unity in order to allow the DMSC parameter to scale well with domain size for a constant choice of the local scale parameter ϵ .

The exponential decay of transition probability with distance ensures that the random walk is likely to persist for a relatively long time on sets of points with small pairwise separations. These sets need not correspond to obvious clusters - groups of points close to one of several centroids - but can be spatially separated sets with similar centroids, e.g. concentric rings in R^2 . Note that, in this case, standard clustering algorithms, such as k-means, would fail to identify the concentric rings because they are centered at the same point. The random walk constructed in diffusion maps would persist within each ring for a long time before transitioning to the other ring.

To exploit this feature of the random walk, we construct the transition matrix P with entries p_{ij} and compute its eigenvalues. By the Perron-Frobenius theorem, the largest eigenvalue of P is $\lambda_0 = 1$; the associated eigenvector \mathbf{v}_0 corresponds to the stationary distribution (Seneta, 1973) The next largest eigenvalues and their associated eigenvectors correspond to the most persistent features (those with long relaxation times to the stationary distribu-

tion) of the random walk. The presence of a spectral gap among the first few eigenvalues allows us to estimate the amount of persistent features. That is, a large drop-off from λ_i to λ_{i+1} indicates that there are i relatively persistent features of the random walk (Rohrdanz et al., 2011). These features correspond to clusters in state space.

The spectral clustering portion of the algorithm associates points in the original time series with the clusters present in the random walk. We construct a low-dimensional data set comprising the first i “important” eigenvectors of P . Each point in this data set is $(\mathbf{v}_{1j}, \mathbf{v}_{2j}, \dots, \mathbf{v}_{ij})$, where $j \in \{1, \dots, N\}$ and N is the number of points in the original data. The new data encode information about the spatial separation between points, so running a clustering algorithm on it finds the persistent features of the random walk. We then assign points in the original data to corresponding clusters in the high-dimensional state space to find the physical significance of the persistent features of the original time series; in this study, we compute the average of each physical cluster in order to illustrate the qualitative difference between the two states.

The full diffusion maps and spectral clustering algorithm is:

1. Compute the transition matrix P according to Eq A1 using the original time series $\{\mathbf{x}_i\}_{i=1}^N$. Normalize so that the entries in each row and column sum to 1.
2. Compute the eigenvalues of P . Determine whether there is a spectral gap to isolate persistent features.
3. Construct a low-dimensional dataset $\{\hat{\mathbf{x}}_i\}_{i=1}^N$ using the first j eigenvectors, where j is the number of large eigenvalues preceding the spectral gap.
4. Run a clustering algorithm on the low-dimensional data to obtain clusters $\{\hat{C}_i\}_{i=1}^j$.
5. Map the original time series to clusters in the high-dimensional state space using the rule

$$\mathbf{x}_i \in C_k \iff \hat{\mathbf{x}}_i \in \hat{C}_k \quad (\text{A.2})$$

6. Compute the averages and/or centroids of the physical clusters C_i .

While the advantage of DMSC is its specificity in picking up on multimodality, its results must be interpreted carefully for several reasons. First, unlike SVD, DMSC does not necessarily find features explaining a significant amount of the variance in a system; it is therefore imperative to verify the dynamic relevance of DMSC results (as we do in this study by comparing to SVD results). Second, while DMSC can be used to characterize the modes of a system, it does not yield explicit dynamical information about the pathway(s) the system takes among modes. Finally, DMSC can be sensitive to the choice of ϵ , so robustness to this parameter should be checked prior to drawing conclusions.

REFERENCES

- Akitomo, K., Awaji, T., and Imasato, N. (1991). Kuroshio path variation South of Japan: 1. Barotropic inflow-outflow model. *Journal of Geophysical Research-Oceans*, 96(C2):2549–2560.
- Barker, D., Huang, X.-Y., Liu, Z., Auligne, T., Zhang, X., Rugg, S., Ajjaji, R., Bourgeois, A., Bray, J., Chen, Y., Demirtas, M., Guo, Y.-R., Henderson, T., Huang, W., Lin, H.-C., Michalakes, J., Rizvi, S., and Zhang, X. (2012). The Weather Research and Forecasting model’s community variational/ensemble data assimilation system WRFDA. *Bulletin of the American Meteorological Society*, 93(6):831–843.
- Belkin, M. and Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396.
- Beniston, M., Stephenson, D. B., Christensen, O. B., Ferro, C. A. T., Frei, C., Goyette, S., Halsnaes, K., Holt, T., Jylha, K., Koffi, B., Palutikof, J., Schoell, R., Semmler, T., and Woth, K. (2007). Future extreme events in European climate: an exploration of regional climate model projections. *Climatic Change*, 81(1):71–95.
- Berg, R. J. and Avila, L. A. (2011). Atlantic hurricane season of 2009. *Monthly Weather Review*, 139(4):1049–1069.
- Bister, M. and Emanuel, K. (1998). Dissipative heating and hurricane intensity. *Meteorology and Atmospheric Physics*, 65(3-4):233–240.
- Borden, K. A. and Cutter, S. L. (2008). Spatial patterns of natural hazards mortality in the United States. *International Journal of Health Geographics*, 7.
- Brennan, M. J. (2009). Hurricane Fred. *NOAA/National Hurricane Center Tropical Cyclone Report*.
- Bretherton, C., Smith, C., and Wallace, J. (1992a). An intercomparison of methods for finding coupled patterns in climate data. *Journal of Climate*, 5(6):541–560.
- Bretherton, C., Smith, C., and Wallace, J. (1992b). An intercomparison of methods for finding coupled patterns in climate data. *Journal of Climate*, 5(6):541–560.
- Camargo, S. J. and Wing, A. A. (2016). Tropical cyclones in climate models. *Wiley Interdisciplinary Reviews-Climate Change*, 7(2):211–237.
- Carr, L. and Williams, R. (1989). Barotropic vortex stability to perturbations from axisymmetry. *Journal of the Atmospheric Sciences*, 46(20):3177–3191.
- Carton, J. A. and Giese, B. S. (2008). A reanalysis of ocean climate using Simple Ocean Data Assimilation (SODA). *Monthly Weather Review*, 136(8):2999–3017.
- Chao, S. (1984). Bimodality of the Kuroshio. *Journal of Physical Oceanography*, 14(1):92–103.

- Chen, H. and Gopalakrishnan, S. G. (2015). A Study on the Asymmetric Rapid Intensification of Hurricane Earl (2010) Using the HWRF System. *Journal of the Atmospheric Sciences*, 72(2):531–550.
- Chen, H. and Zhang, D.-L. (2013). On the Rapid Intensification of Hurricane Wilma (2005). Part II: Convective Bursts and the Upper-Level Warm Core. *Journal of the Atmospheric Sciences*, 70(1):146–162.
- Coifman, R. R., Kevrekidis, I. G., Lafon, S., Maggioni, M., and Nadler, B. (2008). Diffusion maps, reduction coordinates, and low dimensional representation of stochastic systems. *Multiscale Modeling & Simulation*, 7(2):842–864.
- Corbosiero, K. and Molinari, J. (2003). The relationship between storm motion, vertical wind shear, and convective asymmetries in tropical cyclones. *Journal of the Atmospheric Sciences*, 60(2):366–376.
- Dellago, C., Bolhuis, P., and Geissler, P. (2002). Transition path sampling. In Prigogine, I and Rice, SA, editor, *Advances in Chemical Physics, Vol 123*, volume 123 of *Advances in Chemical Physics*, pages 1–78. John Wiley & Sons INC, 111 RIVER ST, HOBOKEN, NJ 07030 USA.
- DeMaria, M. (1996). The effect of vertical shear on tropical cyclone intensity change. *Journal of the Atmospheric Sciences*, 53(14):2076–2087.
- DeMaria, M. and Kaplan, J. (1994). Sea-surface temperature and the maximum intensity of Atlantic tropical cyclones. *Journal of Climate*, 7(9):1324–1334.
- Dematteis, G., Grafke, T., and Vanden-Eijnden, E. (2018). Rogue waves and large deviations in deep sea. *Proceedings of the National Academy of Sciences of the United States of America*, 115(5):855–860.
- E, W., Ren, W., and Vanden-Eijnden, E. (2002). String method for the study of rare events. *Physical Review B*, 66(5).
- E, W., Ren, W., and Vanden-Eijnden, E. (2004). Minimum action method for the study of rare events. *Communications on Pure and Applied Mathematics*, 57(5):637–656.
- E, W. and Vanden-Eijnden, E. (2010). Transition-path theory and path-finding algorithms for the study of rare events. In Leone, SR and Cremer, PS and Groves, JT and Johnson, MA and Richmond, G, editor, *Annual Review of Physical Chemistry*, volume 61 of *Annual Review of Physical Chemistry*, pages 391–420. Annual Reviews.
- Easterling, D., Evans, J., Groisman, P., Karl, T., Kunkel, K., and Ambenje, P. (2000). Observed variability and trends in extreme climate events: A brief review. *Bulletin of the American Meteorological Society*, 81(3):417–425.
- Ebuchi, N. and Hanawa, K. (2003). Influence of mesoscale eddies on variations of the Kuroshio path south of Japan. *Journal of Oceanography*, 59(1):25–36.

- Emanuel, K. (1999). The power of a hurricane: An example of reckless driving on the information superhighway. *Weather*, 54(4):107–108.
- Emanuel, K. A. (1988). The maximum intensity of hurricanes. *Journal of the Atmospheric Sciences*, 45(7):1143–1155.
- Emanuel, K. A. (1991). The theory of hurricanes. *Annual Review of Fluid Mechanics*, 23:179–196.
- Emanuel, K. A. (2005). Increasing destructiveness of tropical cyclones over the past 30 years. *Nature*, 436(7051):686–688.
- Farrell, P. E., Ham, D. A., Funke, S. W., and Rognes, M. E. (2013). Automated derivation of the adjoint of high-level transient finite element programs. *SIAM Journal on Scientific Computing*, 35(4):C369–C393.
- Fierro, A. O., Rogers, R. F., Marks, F. D., and Nolan, D. S. (2009). The impact of horizontal grid spacing on the microphysical and kinematic structures of strong tropical cyclones simulated with the WRF-ARW model. *Monthly Weather Review*, 137(11):3717–3743.
- Frank, W. and Ritchie, E. (2001). Effects of vertical wind shear on the intensity and structure of numerically simulated hurricanes. *Monthly Weather Review*, 129(9):2249–2269.
- Frei, C. and Schar, C. (2001). Detection probability of trends in rare events: Theory and application to heavy precipitation in the Alpine region. *Journal of Climate*, 14(7):1568–1584.
- Freidlin, M. I. and Wentzell, A. D. (1984). *Random Perturbations*, pages 15–43. Springer US, New York, NY.
- Gentry, M. S. and Lackmann, G. M. (2010). Sensitivity of simulated tropical cyclone structure and intensity to horizontal resolution. *Monthly Weather Review*, 138(3):688–704.
- Giannakis, D. and Majda, A. J. (2012). Nonlinear Laplacian spectral analysis for time series with intermittency and low-frequency variability. *Proceedings of the National Academy of Sciences of the United States of America*, 109(7):2222–2227.
- Gretener, P. (1967). Significance of the rare event in geology. *AAPG Bulletin*, 51(11):2197–2206.
- Hannachi, A. (2004). A primer for EOF analysis of climate data. Reading: University of Reading.
- Hasumi, H., Tatebe, H., Kawasaki, T., Kurogi, M., and Sakamoto, T. T. (2010). Progress of North Pacific modeling over the past decade. *Deep-Sea Research Part II-Topical Studies in Oceanography*, 57(13-14):1188–1200.
- Hendricks, E., Montgomery, M., and Davis, C. (2004). The role of “vortical” hot towers in the formation of tropical cyclone Diana (1984). *Journal of the Atmospheric Sciences*, 61(11):1209–1232.

- Hoffman, R., Henderson, J., Leidner, S., Grassotti, C., and Nehrkorn, T. (2006a). Using 4d-var to move a simulated tropical cyclone in a mesoscale model. *Computers and Mathematics with Applications*, 52(8):1193 – 1204. Variational Data Assimilation and Optimal Control.
- Hoffman, R. N., Henderson, J. M., Leidner, S. M., Grassotti, C., and Nehrkorn, T. (2006b). The response of damaging winds of a simulated tropical cyclone to finite-amplitude perturbations of different variables. *Journal of the Atmospheric Sciences*, 63(7):1924–1937.
- Holland, G. (1997). The maximum potential intensity of tropical cyclones. *Journal of the Atmospheric Sciences*, 54(21):2519–2541.
- Hong, S.-Y. and Lim, J.-O. J. (2006). The WRF Single-Moment 6-Class Microphysics Scheme (WSM6). *Journal of the Korean Meteorological Society*, 42(2):129–151.
- Hong, S.-Y., Noh, Y., and Dudhia, J. (2006). A new vertical diffusion package with an explicit treatment of entrainment processes. *Monthly Weather Review*, 134(9):2318–2341.
- Huang, X.-Y., Xiao, Q., Barker, D. M., Zhang, X., Michalakes, J., Huang, W., Henderson, T., Bray, J., Chen, Y., Ma, Z., Dudhia, J., Guo, Y., Zhang, X., Won, D.-J., Lin, H.-C., and Kuo, Y.-H. (2009). Four-dimensional variational data assimilation for WRF: formulation and preliminary results. *Monthly Weather Review*, 137(1):299–314.
- Hurlburt, H., Wallcraft, A., Schmitz, W., Hogan, P., and Metzger, E. (1996). Dynamics of the Kuroshio/Oyashio current system using eddy-resolving models of the North Pacific Ocean. *Journal Of Geophysical Research-Oceans*, 101(C1):941–976.
- Jones, S. (1995). The evolution of vortices in vertical shear .1. Initially barotropic vortices. *Quarterly Journal of the Royal Meteorological Society*, 121(524, B):821–851.
- Jonsson, H., Mills, G., and Jacobsen, K. W. (1998). Nudged elastic band method for finding minimum energy paths of transitions. In Berne, B., Ciccotti, G., and Coker, D., editors, *Classical and Quantum Dynamics in Condensed Phase Simulations*, pages 385–404. World Scientific.
- Judt, F., Chen, S. S., and Berner, J. (2016). Predictability of tropical cyclone intensity: scale-dependent forecast error growth in high-resolution stochastic kinetic-energy backscatter ensembles. *Quarterly Journal of the Royal Meteorological Society*, 142(694, A):43–57.
- Kain, J. S. and Fritsch, J. M. (1993). *The Representation of Cumulus Convection in Numerical Models*. American Meteorological Society.
- Kaplan, J. and DeMaria, M. (2003). Large-scale characteristics of rapidly intensifying tropical cyclones in the North Atlantic basin. *Weather and Forecasting*, 18(6):1093–1108.
- Kaplan, J., Rozoff, C. M., DeMaria, M., Sampson, C. R., Kossin, J. P., Velden, C. S., Cione, J. J., Dunion, J. P., Knaff, J. A., Zhang, J. A., Dostalek, J. F., Hawkins, J. D., Lee, T. F., and Solbrig, J. E. (2015). Evaluating environmental impacts on tropical cyclone rapid intensification predictability utilizing statistical models. *Weather and Forecasting*, 30(5):1374–1396.

- Katz, R. and Brown, B. (1992). Extreme events in a changing climate - variability is more important than averages. *Climatic Change*, 21(3):289–302.
- Kawabe, M. (1980). Sea level variations around the Nansei Islands and the large meander in the Kuroshio south of central Japan. *Journal of the Oceanographic Society of Japan*.
- Kawabe, M. (1995). Variations of current path, velocity, and volume transport of the Kuroshio in relation with the large meander. *Journal of Physical Oceanography*, 25(12):3103–3117.
- Knierim, J. (2006). Numerical weather prediction (nwp) and the wrf model. In *ATEC Forecasters' Conference*.
- Knutson, T. R., McBride, J. L., Chan, J., Emanuel, K., Holland, G., Landsea, C., Held, I., Kossin, J. P., Srivastava, A. K., and Sugi, M. (2010). Tropical cyclones and climate change. *Nature Geoscience*, 3(3):157–163.
- Kowch, R. and Emanuel, K. (2015). Are special processes at work in the rapid intensification of tropical cyclones? *Monthly Weather Review*, 143(3):878–882.
- Kurogi, M. and Akitomo, K. (2006). Effects of stratification on the stable paths of the Kuroshio and on their variation. *Deep-Sea Research Part I-Oceanographic Research Papers*, 53(9):1564–1577.
- Laio, A. and Gervasio, F. L. (2008). Metadynamics: a method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science. *Reports on Progress in Physics*, 71(12).
- Leutbecher, M. and Palmer, T. N. (2008). Ensemble forecasting. *Journal of Computational Physics*, 227(7):3515–3539.
- MacAdie, C. J., Landsea, C. W., Neumann, C. J., David, J. E., Blake, E., and Hammer, G. R. (2009). Tropical cyclones of the North Atlantic Ocean, 1851–2006. Technical report, National Climatic Data Center, TCP/National Hurricane Center.
- Maltrud, M. and McClean, J. (2005). An eddy resolving global 1/10 degrees ocean simulation. *Ocean Modelling*, 8(1-2):31–54.
- Marotzke, J., Giering, R., Zhang, K., Stammer, D., Hill, C., and Lee, T. (1999). Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity. *Journal of Geophysical Research-Oceans*, 104(C12):29529–29547.
- Masuda, A. (1982). An Interpretation of the bimodal character of the stable Kuroshio path. *Deep-Sea Research Part I-Oceanographic Research Papers*, 29(4):471–484.
- Maximenko, N. (2002). Index and composites of the Kuroshio meander south of Japan. *Journal of Oceanography*, 58(5):639–649.
- Merrill, R. (1988). Environmental Influences on Hurricane Intensification. *Journal of the Atmospheric Sciences*, 45(11):1678–1687.

- Miller, R., Ghil, M., and Gauthiez, F. (1994). Advanced data assimilation in strongly nonlinear dynamical systems. *Journal of the Atmospheric Sciences*, 51(8):1037–1056.
- Moller, J. and Shapiro, L. (2005). Influences of asymmetric heating on hurricane evolution in the MM5. *Journal of the Atmospheric Sciences*, 62(11):3974–3992.
- Montgomery, M., Nicholls, M., Cram, T., and Saunders, A. (2006). A vortical hot tower route to tropical cyclogenesis. *Journal of the Atmospheric Sciences*, 63(1):355–386.
- Montgomery, M. T. and Smith, R. K. (2014). Paradigms for tropical cyclone intensification. *Australian Meteorological and Oceanographic Journal*, 64(1):37–66.
- Murakami, H., Wang, Y., Yoshimura, H., Mizuta, R., Sugi, M., Shindo, E., Adachi, Y., Yukimoto, S., Hosaka, M., Kusunoki, S., Ose, T., and Kitoh, A. (2012). Future changes in tropical cyclone activity projected by the new high-resolution MRI-AGCM. *Journal of Climate*, 25(9):3237–3260.
- Nakano, H. and Hasumi, H. (2005). A series of zonal jets embedded in the broad zonal flows in the pacific obtained in eddy-permitting ocean general circulation models. *Journal of Physical Oceanography*, 35(4):474–488.
- NCAR (2013). *The Climate Data Guide: AVISO: Satellite derived Sea Surface Height above Geoid*. National Center for Atmospheric Research.
- Nitani, H. (1975). Variation of the Kuroshio south of Japan. *Journal of the Oceanographical Society of Japan*, 31(4):154–173.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer New York, 2 edition.
- Nolan, D. and Farrell, B. (1999). Generalized stability analyses of asymmetric disturbances in one- and two-celled vortices maintained by radial inflow. *Journal of the Atmospheric Sciences*, 56(10):1282–1307.
- Nolan, D. and Grasso, L. (2003). Nonhydrostatic, three-dimensional perturbations to balanced, hurricane-like vortices. Part II: Symmetric response and nonlinear simulations. *Journal of the Atmospheric Sciences*, 60(22):2717–2745.
- Nolan, D. and Montgomery, M. (2002). Three-dimensional, nonhydrostatic perturbations to balanced, hurricane like vortices. Part I: Linearized formulation, stability, and evolution. *Journal of the Atmospheric Sciences*, 59:2989–3020.
- Nolan, D. S. (2007). What is the trigger for tropical cyclogenesis? *Australian Meteorological Magazine*, 56(4):241–266.
- Nolan, D. S. and McGauley, M. G. (2012). *Tropical cyclogenesis in wind shear: Climatological relationships and physical processes*, pages 1–36. Nova Science Publishers, Inc.
- Nolan, D. S., Moon, Y., and Stern, D. P. (2007). Tropical cyclone intensification from asymmetric convection: Energetics and efficiency. *Journal of the Atmospheric Sciences*, 64(10):3377–3405.

- Oouchi, K., Yoshimura, J., Yoshimura, H., Mizuta, R., Kusunoki, S., and Noda, A. (2006). Tropical cyclone climatology in a global-warming climate as simulated in a 20 km-mesh global atmospheric model: Frequency and wind intensity analyses. *Journal of the Meteorological Society of Japan*, 84(2):259–276.
- Park, M.-S. and Elsberry, R. L. (2013). Latent Heating and Cooling Rates in Developing and Nondeveloping Tropical Disturbances during TCS-08: TRMM PR versus ELDORA Retrievals. *Journal of the Atmospheric Sciences*, 70(1):15–35.
- Parrish, D. F. and Derber, J. C. (1992). The National Meteorological Center’s spectral statistical interpolation analysis system. *Monthly Weather Review*, 120:1747–1763.
- Pereira, M. and Berre, L. (2006). The use of an ensemble approach to study the background error covariances in a global NWP model. *Monthly Weather Review*, 134(9):2466–2489.
- Persing, J., Montgomery, M. T., McWilliams, J. C., and Smith, R. K. (2013). Asymmetric and axisymmetric dynamics of tropical cyclones. *Atmospheric Chemistry and Physics*, 13(24):12299–12341.
- Pielke, R. and Landsea, C. (1998). Normalized hurricane damages in the United States: 1925-95. *Weather and Forecasting*, 13(3, 2):621–631.
- Pierini, S. (2006). A Kuroshio Extension system model study: Decadal chaotic self-sustained oscillations. *Journal of Physical Oceanography*, 36(8):1605–1625.
- Pierini, S., Dijkstra, H. A., and Riccio, A. (2009). A Nonlinear Theory of the Kuroshio Extension Bimodality. *Journal of Physical Oceanography*, 39(9):2212–2229.
- Plotkin, D. A., Weare, J., and Abbot, D. S. (2014). Distinguishing meanders of the Kuroshio using machine learning. *Journal of Geophysical Research-Oceans*, 119(10):6593–6604.
- Qiu, B. and Miao, W. (2000). Kuroshio path variations south of Japan: Bimodality as a self-sustained internal oscillation. *Journal of Physical Oceanography*, 30(8):2124–2137.
- Ragone, F., Wouters, J., and Bouchet, F. (2018). Computation of extreme heat waves in climate models using a large deviation algorithm. *Proceedings of the National Academy of Sciences*, 115(1):24–29.
- Rappin, E. D. and Nolan, D. S. (2012). The effect of vertical shear orientation on tropical cyclogenesis. *Quarterly Journal of the Royal Meteorological Society*, 138(665, B):1035–1054.
- Riemer, M., Montgomery, M. T., and Nicholls, M. E. (2010). A new paradigm for intensity modification of tropical cyclones: thermodynamic impact of vertical wind shear on the inflow layer. *Atmospheric Chemistry and Physics*, 10(7):3163–3188.
- Roberts, M. J., Vidale, P. L., Mizieliński, M. S., Demory, M.-E., Schiemann, R., Strachan, J., Hodges, K., Bell, R., and Camp, J. (2015). Tropical cyclones in the UPSCALE ensemble of high-resolution global climate models. *Journal of Climate*, 28(2):574–596.

- Rodgers, J. L. and Nicewander, A. W. (1988a). Thirteen Ways to Look at the Correlation Coefficient. *The American Statistician*, 42(1):59–66.
- Rodgers, J. L. and Nicewander, A. W. (1988b). Thirteen Ways to Look at the Correlation Coefficient. *The American Statistician*, 42(1):59–66.
- Rogers, R., Aberson, S., Aksoy, A., Annane, B., Black, M., Cione, J., Dorst, N., Dunion, J., Gamache, J., Goldenberg, S., Gopalakrishnan, S., Kaplan, J., Klotz, B., Lorsolo, S., Marks, F., Murillo, S., Powell, M., Reasor, P., Sellwood, K., Uhlhorn, E., Vukicevic, T., Zhang, J., and Zhang, X. (2013). NOAA’s hurricane intensity forecasting experiment: A Progress Report. *Bulletin of the American Meteorological Society*, 94(6):859–882.
- Rogers, R. F. (2010). Convective-scale structure and evolution during a high-resolution simulation of tropical cyclone rapid intensification. *Journal of the Atmospheric Sciences*, 67(1):44–70.
- Rohrdanz, M. A., Zheng, W., Maggioni, M., and Clementi, C. (2011). Determination of reaction coordinates via locally scaled diffusion map. *Journal of Chemical Physics*, 134(12).
- Schenkel, B. A., Lin, N., Chavas, D., Oppenheimer, M., and Brammer, A. (2017). Evaluating outer tropical cyclone size in reanalysis datasets using QuikSCAT data. *Journal of Climate*, 30(21):8745–8762.
- Schmeits, M. and Dijkstra, H. (2001). Bimodal behavior of the Kuroshio and the Gulf Stream. *Journal of Physical Oceanography*, 31(12):3435–3456.
- Schubert, W. and Hack, J. (1982). Inertial stability and tropical cyclone development. *Journal of the Atmospheric Sciences*, 39(8):1687–1697.
- Seager, R., Kushnir, Y., Naik, N., Cane, M., and Miller, J. (2001). Wind-driven shifts in the latitude of the Kuroshio-Oyashio Extension and generation of SST anomalies on decadal timescales. *Journal of Climate*, 14(22):4249–4265.
- Sekine, Y. (1990). A numerical experiment on the path dynamics of the Kuroshio with reference to the formation of the large meander path south of Japan. *Deep-Sea Research Part I-Oceanographic research papers*, 37(3):359–380.
- Seneta, E. (1973). *Nonnegative Matrices*. Allen & Unwin.
- Shapiro, L. (2000). Potential vorticity asymmetries and tropical cyclone evolution in a moist three-layer model. *Journal of the Atmospheric Sciences*, 57(21):3645–3662.
- Simpson, J., Ritchie, E., Holland, G., Halverson, J., and Stewart, S. (1997). Mesoscale interactions in tropical cyclone genesis. *Monthly Weather Review*, 125(10):2643–2661.
- Singh, R., Kishtawal, C. M., Pal, P. K., and Joshi, P. C. (2011). Assimilation of the multi-satellite data into the WRF model for track and intensity simulation of the Indian Ocean tropical cyclones. *Meteorology and Atmospheric Physics*, 111(3-4):103–119.

- Stern, D. P. and Nolan, D. S. (2012). On the height of the warm core in tropical cyclones. *Journal of the Atmospheric Sciences*, 69(5):1657–1680.
- Stewart, S. R. (2016). Hurricane Danny. *NOAA/National Hurricane Center Tropical Cyclone Report*.
- Taft, B. A. (1972). Characteristics of the flow of the Kuroshio south of Japan. In Stommel, H. and Yoshida, K., editors, *Kuroshio—Its Physical Aspects*, pages 165–216. University of Tokyo Press.
- Tang, B. and Emanuel, K. (2010). Midlevel Ventilation’s Constraint on Tropical Cyclone Intensity. *Journal of the Atmospheric Sciences*, 67(6):1817–1830.
- Tao, C. and Jiang, H. (2015). Distributions of Shallow to Very Deep Precipitation-Convection in Rapidly Intensifying Tropical Cyclones. *Journal of Climate*, 28(22):8791–8824.
- Tao, D. and Zhang, F. (2014). Effect of environmental shear, sea-surface temperature, and ambient moisture on the formation and predictability of tropical cyclones: An ensemble-mean perspective. *Journal of Advances in Modeling Earth Systems*, 6(2):384–404.
- Trenberth, K. (1997). The definition of El Nino. *Bulletin of the American Meteorological Society*, 78(12):2771–2777.
- Van Sang, N., Smith, R. K., and Montgomery, M. T. (2008). Tropical-cyclone intensification and predictability in three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 134(632, A):563–582.
- Vanden-Eijnden, E. and Weare, J. (2012). Rare Event Simulation of Small Noise Diffusions. *Communications on Pure and Applied Mathematics*, 65(12, SI):1770–1803.
- Vanden-Eijnden, E. and Weare, J. (2013a). Data Assimilation in the Low Noise Regime with Application to the Kuroshio. *Monthly Weather Review*, 141(6):1822–1841.
- Vanden-Eijnden, E. and Weare, J. (2013b). Data assimilation in the low noise regime with application to the Kuroshio. *Monthly Weather Review*, 141(6):1822–1841.
- Vigh, J. L. and Schubert, W. H. (2009). Rapid Development of the Tropical Cyclone Warm Core. *Journal of the Atmospheric Sciences*, 66(11):3335–3350.
- Weare, J. (2009). Particle filtering with path sampling and an application to a bimodal ocean current model. *Journal of Computational Physics*, 228(12):4312 – 4331.
- Webster, P., Holland, G., Curry, J., and Chang, H. (2005). Changes in tropical cyclone number, duration, and intensity in a warming environment. *Science*, 309(5742):1844–1846.
- White, W. and McCreary, J. (1976). Formation of Kuroshio Meander and its Relationship to Large-Scale Ocean Circulation. *Deep-Sea Research*, 23(1):33–47.

- Wong, M. and Chan, J. (2004). Tropical cyclone intensity in vertical wind shear. *Journal of the Atmospheric Sciences*, 61(15):1859–1876.
- Xiao, Q., Zhang, X., Davis, C., Tuttle, J., Holland, G., and Fitzpatrick, P. J. (2009). Experiments of hurricane initialization with airborne doppler radar data for the Advanced Research Hurricane WRF (AHW) Model. *Monthly Weather Review*, 137(9):2758–2777.
- Xu, R. and Wunsch, D. (2005). Survey of clustering algorithms. *IEEE Transactions on Neural Networks*, 16(3):645–678.
- Yang, S., Olson, W., Kummerow, C., Chen, S., Haddad, Z., and Smith, E. (2002). Latent heating structures of hurricanes from TRMM measurements. In *Preprints, 25th Conf. on Hurricanes and Tropical Meteorology, San Diego, CA, Amer. Meteor. Soc., CD-ROM*, volume 13, page A3.
- Yoshida, K (1961). On the variations of Kuroshio and cold water mass off Enshu Nada. *Hydrographic Bulletin*, 67:54–57.
- Zagrodnik, J. P. and Jiang, H. (2014). Rainfall, Convection, and Latent Heating Distributions in Rapidly Intensifying Tropical Cyclones. *Journal of the Atmospheric Sciences*, 71(8):2789–2809.
- Zeng, Z., Wang, Y., and Wu, C.-C. (2007). Environmental dynamical control of tropical cyclone intensity - An observational study. *Monthly Weather Review*, 135(1):38–59.
- Zhang, F. and Tao, D. (2013). Effects of vertical wind shear on the predictability of tropical cyclones. *Journal of the Atmospheric Sciences*, 70(3):975–983.
- Zhang, Q., Hou, Y., and Yan, T. (2012). Inter-annual and inter-decadal variability of Kuroshio heat transport in the East China Sea. *International Journal of Climatology*, 32(4):481–488.
- Zupanski, D. (1996). A general weak constraint applicable to operational 4DVAR data assimilation systems. *Monthly Weather Review*, 125.