

©Copyright 2021

Ariana Mendible

Data-Driven Reduced-Order Modeling for Symmetrical Systems

Ariana Mendible

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

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Steven Brunton, Chair

J. Nathan Kutz, Chair

Duane Storti

Program Authorized to Offer Degree:
Mechanical Engineering

University of Washington

Abstract

Data-Driven Reduced-Order Modeling
for Symmetrical Systems

Ariana Mendible

Co-Chairs of the Supervisory Committee:

Associate Professor Steven Brunton
Mechanical Engineering

Professor J. Nathan Kutz
Applied Mathematics

Spatiotemporal dynamical systems are ubiquitous across all areas of science and engineering. While some systems of interest can be studied through first-principles governing equations, often, the underlying dynamics are unknown. Furthermore, even for problems where equations are known, analytical solutions are rare. Computational solutions are then plagued by issues of nonlinearity, multiple relevant scales of time and space, chaos, and high-dimensionality. When parametric studies are needed, these simulations become prohibitively expensive, even under the most advanced computational architectures.

Often, even the most complex dynamical systems contain inherent low-dimensional coherent structure. Reduced-order models (ROMs) attempt to ease the computation burden of such simulations by uncovering this low-rank subspace, reducing the overall size of such models while retaining the critical dynamics. Many ROMs leverage the proper orthogonal decomposition (POD) to produce a linear dimensionality reduction, where a dominant set of correlated modes provide a subspace in which to project the dynamics. Dimensionality reduction techniques then enable downstream tasks such as prediction, estimation, and low-latency control. However, many dimensionality reduction techniques rely on separation of

time and space variables. This assumption does not hold in systems with underlying symmetry, because the time and space variables are inherently coupled. Namely translation, or traveling waves, scaling, and rotation inhibit the effectiveness of dimensionality reduction techniques.

The focus of this thesis is to develop an approach to make spatiotemporal systems with underlying symmetry amenable to traditional ROM architectures. Three objectives are addressed (1) developing a method to address translations in one spatial dimension, (2) demonstrating the effectiveness of such a method on complex dynamics in experimental data, and (3) expanding the method to approach higher-dimensional data with scaling symmetries. The Unsupervised Traveling Wave Identification with Shifting and Truncation (UnTWIST) method can be applied to pure data, yielding interpretable models for wave speeds. Several examples are explored, showing that even for systems with multiple waves, non-constant wave speeds, and nonlinear wave phenomenon, UnTWIST enables the use of traditional methods like POD for dimensionality reduction. Further, the method is demonstrated on experimental data from a rotating detonation engine, presenting complex nonlinear dynamics. UnTWIST enables the discovery of interpretable linear and nonlinear models for the traveling detonation speeds, and yields meaningful spatial modes to describe the wave shapes. Lastly, the method is augmented to accommodate higher-dimensional data and an arbitrary number of symmetries. Two-dimensional data with a scaling symmetry are investigated, and accurate models are found to improve the dimensionality reductions. Lastly, limitations to the method in the presence of interacting waves are explored, and improvements and expansions are discussed.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	v
Chapter 1: Introduction	1
1.1 Motivation	1
1.2 Methods for Dimensionality Reduction	4
1.2.1 Proper Orthogonal Decomposition (POD)	4
1.2.2 Robust Principal Components Analysis (RPCA)	5
1.2.3 Dynamic Mode Decomposition (DMD)	5
1.3 Related Dimensionality Reduction for Translational Symmetry	6
Chapter 2: Developing an Optimization for Model Discovery	10
2.1 Formulating the Optimization Framework	11
2.1.1 The UnTWIST Minimization	13
2.1.2 Implementation of UnTWIST	14
2.2 Results on Example Systems: Model Discovery	17
2.2.1 Korteweg-de Vries Equation	17
2.2.2 Two Crossing Solitons: Non-periodic and Non-Constant Speed	19
2.2.3 Burgers' Equation	19
2.2.4 Nonlinear Schrödinger Equation	21
2.3 Results on Example Systems: Dimensionality Reduction	23
Chapter 3: UnTWIST on Experimental Data: Rotating Detonation Engine	28
3.1 Experimental Data: Rotating Detonation Engine	28
3.2 Detonation Wave Tracking with UnTWIST	33
3.2.1 UnTWIST Method	33

3.2.2	UnTWIST Applied to RDE Data	36
3.3	Data-Driven Models of Rotating Detonation Front Dynamics	39
3.3.1	Linear Dynamics: Dynamic Mode Decomposition	39
3.3.2	Nonlinear Dynamics: Lotka-Volterra Model	40
3.4	Discussion	43
Chapter 4:	Higher-Dimensional Symmetries	45
4.1	Extensions to UnTWIST: Arbitrary Dimensions and Symmetries	45
4.1.1	Modifications to the Initialization: Symmetry Detection	45
4.1.2	Modifications to the UnTWIST Algorithm	48
4.2	Example Problem: Growing Two-Dimensional Gaussian	50
4.3	High-Altitude Nuclear Explosions	55
4.4	Three-Dimensional Euler Equations	55
4.5	One Wave in Two Dimensions	57
4.6	Two Waves in Two Dimensions	62
4.7	Discussion	70
Chapter 5:	Conclusions	71
Appendix A:	74

LIST OF FIGURES

Figure Number	Page
1.1 Invariance effect on dimensionality reduction	3
2.1 Schematic of the optimization-based shift mechanism	11
2.2 Model discovery on the Korteweg-de Vries Equation	18
2.3 Model discovery on a challenging example system	20
2.4 Model discovery on Burgers' Equation	21
2.5 Model discovery on the Nonlinear Schrödinger Equation	23
2.6 Comparison of dimensionality reduction methods for generated example data	25
2.7 Comparison of dimensionality reduction methods for NLS data	26
3.1 Schematic of rotating detonation engine and data processing	32
3.2 Initialization of UnTWIST for RDEs	34
3.3 Preprocessing step using UnTWIST on RDE data	37
3.4 Refinement step using UnTWIST on RDE data	38
3.5 Linear model of RDE wave front interactions	40
3.6 Nonlinear RDE data set	41
3.7 Nonlinear model of RDE wave front interactions	42
4.1 Time series of example 2D Gaussian problem	46
4.2 Schematic of symmetry identification in higher dimensions	48
4.3 Shifting and scaling on 2D Gaussian example data.	53
4.4 1-mode reconstructions of 2D Gaussian example data	54
4.5 Example of one- and two-wave Euler simulations	57
4.6 Time series of single detonations with various initial energies	58
4.7 Model discovery and singular values for one-wave Euler simulations	59
4.8 1-mode reconstructions of one-wave Euler simulations	61
4.9 Euler simulations with two waves	65
4.10 Model discovery for two-wave Euler simulations	66

4.11 Singular values for two-wave Euler simulations 67
4.12 2-mode reconstructions of closely-spaced two-wave Euler simulations 68
4.13 2-mode reconstructions of farther-spaced two-wave Euler simulations 69

LIST OF TABLES

Table Number	Page
A.1 Hyperparameters, iterations, errors, and resulting models for one-wave Euler simulations	74
A.2 Hyperparameters, iterations, and errors for two-wave Euler simulations . . .	75
A.3 Models for two-wave Euler simulations	76

ACKNOWLEDGMENTS

This dissertation is the result of five years of work at the University of Washington. Throughout this time, I have been supported and encouraged by a multitude of people to whom I owe my gratitude. My greatest thanks are to my advisors Steve Brunton and Nathan Kutz for taking me on as an “orphaned” student and for sharing their passion and expertise. They have modeled exceptional teaching, scholarship, and mentorship simultaneously; I can only hope to follow in their example. Thank you for your invaluable guidance and equally treasured espresso.

I would like to thank my committee members Duane Storti and Kevin Jamieson for their thoughtful reflections and perspectives on this work. I would also like to thank my collaborators and coauthors, whose data and knowledge base provided the platform on which to build this dissertation. Thank you to Sasha Aravkin for your challenging optimization course and your enthusiastic work on developing UnTWIST, to Wes Lowrie for being a dependable collaborator and for providing much of the data for Chapter 4, to James Koch for the experimental data for Chapter 3 and helpful pep talks throughout our working together, and to Henning Lange for your expertise in neural networks and computing, all which greatly strengthened our work. I would also like to acknowledge the mentors who helped me early on in my path to graduate school, John Carter and Teodora Shuman.

Much of this work was strengthened and shaped through thoughtful conversations and supportive lab mates along the way. Thank you to Benjamin Hermann, Krithika Manohar, Kathleen Champion, Sam Rudy, Seth Hirsch, Ben Strom, Thomas Mohren, Jared Callahan, Kadi Kaheman, Michelle Hickner, Isabel Scherl and all the Brunton and Kutz lab members for helpful feedback on talks and figures and for creating a vibrant group. Thank you

to Emma Cotter, Curty Rusch, Trent Dillon, Seth Hirsch, and especially to Isabel Scherl for collaboration and commiseration in tough courses. Thank you to everyone involved in MESAR for eye-opening conversations on racial justice, which will certainly shape my future in research and education.

Getting through graduate school would not have been possible without the support of friends and family along the way. Thank you to Becca, Mel and Cat for being a refuge outside of school, to Izzy for cooking projects and impromptu century rides, to Mike and Hannah for being excellent half-time roommates, to Curty Rusch and Sarah Edwards for helping maintain sanity during the pandemic, and to Marissa Miramontes, Angela Straccia, Laurel Marsh, and all the department's women for honest conversations and creating a positive environment together. Most especially, thank you to Trevor Harrison for being an exceptionally supportive and patient partner. Sharing this experience with you has been a blessing. Thank you to all my family for encouraging me from all over the world, most especially to my parents for supporting me at every step in my education.

Funding: This work was supported by the Army Research Office under PECASE W911NF-19-1-0045, the Defense Threat Reduction Agency under HDTRA1-18-1-0038, and the Graduate Opportunities and Minority Achievement Program through the Presidential Graduate Fellowship.

Chapter 1

INTRODUCTION

1.1 Motivation

Many open problems in science and engineering involve dynamical systems where both space and time are relevant variables, known as a spatiotemporal system. This broad category of systems exists across disciplines, from fluid dynamics to neurobiology, and modeling and understanding them represents a large share of research endeavors. Typically, the starting point for modeling and understanding these systems lies in governing equations derived from first principles or known physical relationships. However, complete governing equations for such systems are seldom known, and analytical solutions are even more rare. These governing physics, often in the form of PDEs are rarely amenable to analytic, closed-form solutions, thus requiring recourse to computational solutions. In light of this, computational solutions and experimental data are indispensable for studying and predicting the behaviors of these systems [1, 2]. Both computational solutions and experimental data collection are prone to added noise and uncertainty in real data. In addition, computational approaches often require fine numerical discretization in order to accurately resolve multiple time and space scales and underlying nonlinearities [3], which generates high-dimensional solutions. This is in contrast with the observation that the underlying dynamics often exhibit low-dimensional structure [4]. This is especially true in fluid systems, which manifest many coherent structures dynamics such as vortices, large scale eddies, and traveling waves [1, 4]. Reduced-order models (ROMs) exploit the intrinsic, low-rank structure of the data in order to create more tractable models for the spatiotemporal evolution dynamics of the PDE. Historically, ROMs arose in the fluid dynamics community in order to study the low-dimensional structures which are canonical to complex flows [4]. Even high-dimensional turbulence can be more

efficiently represented and analyzed within such a reduction framework. Indeed, the majority of ROM simulations today for time-dependent PDEs are focused on fluid dynamic simulations, typically at high Reynolds numbers where improved simulation architectures are required for Monte-Carlo simulations for parametric studies and uncertainty quantification. Although emerging techniques in machine learning are being used to identify more effective coordinate systems and models for systems in fluid mechanics [5–10], the majority of the time, low-rank structure used in ROMs is extracted using more established methods. Many ROMs leverage the *proper orthogonal decomposition* (POD) [4, 11–13], which is typically computed via the *singular value decomposition* (SVD). This tool produces a linear dimensionality reduction [3, 14], whereby a dominant set of correlated modes provide a subspace in which to project the PDE dynamics [2, 15, 16]. Typically, low-energy modes are truncated, leaving a small number of high-energy modes to approximate the original system. If the governing equations are available, Galerkin projection [4, 11, 12, 17], or the improved Petrov-Galerkin projection [18, 19], can then be used to create an interpretable model of the few modes that contain the majority of the system’s energy (variance). While POD-Galerkin projections have been quite successful in many cases [20, 21], this approach requires access to the governing equations, and there are well-known issues with the POD subspace [22]. In addition to POD-based methods, data-driven methods for ROMs have gained traction, such as the *dynamic mode decomposition* (DMD) [23–26] which creates a best-fit linear operator to advance spatiotemporal snapshots forward in time. Dimensionality reduction and modal decomposition approaches have been well-studied and are extremely efficient [2, 15, 16, 27–29]. They enabling downstream tasks such as prediction, estimation and low-latency control. However, many dimensionality reduction techniques rely on the inherent separation of the time and space. This separation, described in equation (1.1), parses the time variance, a , and spatial behavior ϕ .

$$u(x, t) = a(t)\phi(x) \tag{1.1}$$

Inherent to this approach is the assumption that these two variables can be completely decoupled. In systems with time correlation of spatial phenomenon, such as in invariances–

namely translation, scaling, and rotation– the SVD, and thus many ROMs techniques, is not appropriate. To illustrate this canonical phenomenon, the POD/SVD of stationary, traveling, and growing and traveling soliton waves are compared in Figure 1.1. The stationary wave is perfectly captured by one spatial mode, as seen in the rapid singular value decay, whereas a wave translating or growing in time cannot be represented in few modes with POD/SVD. This space-time separation problem has received surprisingly little attention, but research is growing [30–33].

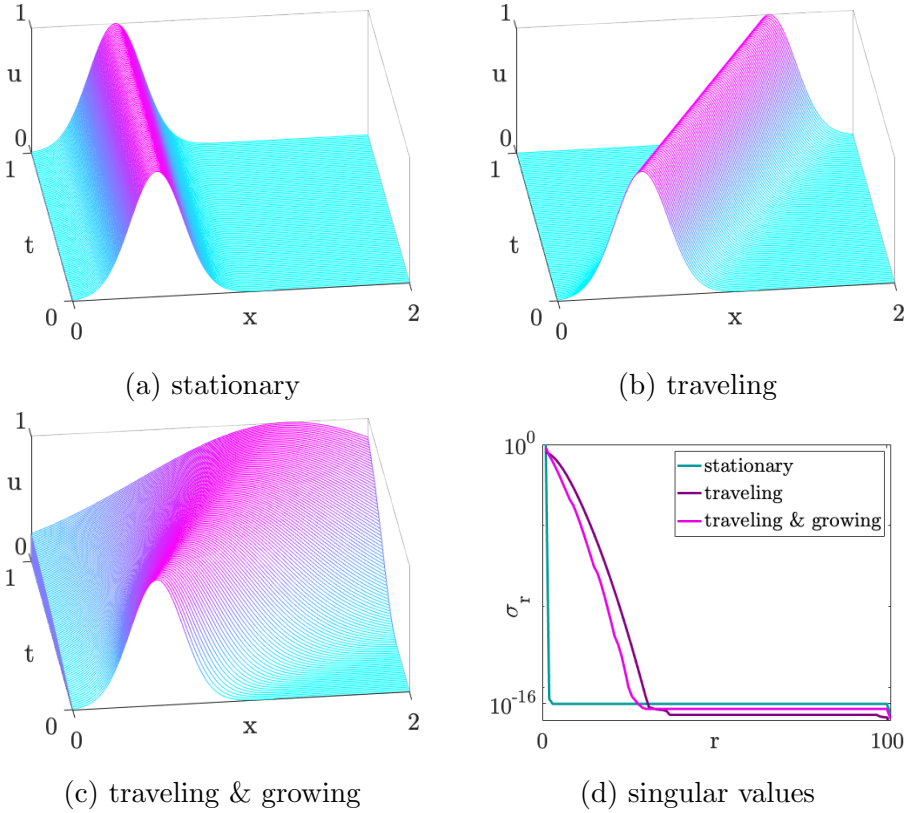


Figure 1.1: A stationary wave contains no time-invariances, traveling wave exhibits a translational invariance, and traveling & growing wave contains a translational and scaling invariance. Singular value spectra for each show that these invariances inhibit the effectiveness of the dimensionality reduction– energy is spread into many modes, therefore more are needed to approximate the original system.

The focus of this thesis is to ameliorate the issues presented by these symmetries in order to utilize traditional and well-studied techniques for ROMs.

1.2 Methods for Dimensionality Reduction

There are a number of widely used methods for dimensionality reduction, or modal decomposition. These methods utilize the principle of separation of variables to break down spatiotemporal data into a set of spacial modes with a corresponding time evolution. These modes capture the information contained in the system in order of decreasing variance or importance. Therefore, the system can be well approximated by retaining these high-importance modes and truncating later modes. This reduces the amount of retained information (i.e. fewer modes), while keeping the maximum amount of detail. In general, the modes may have the useful property of being orthogonal, and in some methods, additional properties imposed by their construction.

1.2.1 Proper Orthogonal Decomposition (POD)

The proper orthogonal decomposition (POD), also known as Principal Components Analysis, is perhaps the most common data-driven method for dimensionality reduction, and is widely used in the fluid dynamics community [2, 11–13]. POD performs the decomposition given in (1.1) by computing a set of basis modes ϕ to describe spatial variation in the data. POD is computed on data organized into a *snapshot matrix*, given as

$$\mathbf{X} = \begin{bmatrix} | & | & | & \dots & | \\ \mathbf{x}_1 & \mathbf{x}_2 & \mathbf{x}_3 & \dots & \mathbf{x}_m \\ | & | & | & \dots & | \end{bmatrix} \quad (1.2)$$

where a snapshot $\mathbf{x}_i = \mathbf{x}(t_i) \in \mathbb{C}^n$ is the state of the n -dimensional system a point in time $t_i, i = 1 \dots m$. The singular value decomposition (SVD) is then used to rapidly compute a low-rank representation of the snapshot matrix.

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \quad (1.3)$$

where columns of \mathbf{U} are the spatial modes. The singular values are given in the diagonal matrix $\mathbf{\Sigma}$, indicating the relative weighting of each mode. The columns of \mathbf{V}^T are the time series of each mode to reconstruct \mathbf{X} . These are analogous to Equation (1.1), with columns of \mathbf{U} equivalent to spacial modes ϕ and columns of \mathbf{V}^T giving $a(t)$ corresponding to each mode. This decomposition guarantees the best ℓ_2 -norm approximation of \mathbf{X} and modes descend in order of importance or energy. Truncating the system by only using the first r modes, \mathbf{X}_r gives the best rank- r approximation to \mathbf{X} .

$$\mathbf{X}_r = \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r^T \quad (1.4)$$

1.2.2 Robust Principal Components Analysis (RPCA)

Robust Principal Components Analysis (RPCA) is a useful extension to POD/PCA that improves robustness to outliers and noise [34]. The method uses an optimization to separate \mathbf{X} into a low-rank matrix, \mathbf{L} , and a sparse matrix, \mathbf{S} . It is computed as

$$\min_{\mathbf{L}, \mathbf{S}} \|\mathbf{L}\|_* + \lambda_0 \|\mathbf{S}\|_1 \text{ s.t. } \mathbf{L} + \mathbf{S} = \mathbf{X} \quad (1.5)$$

where $\|\cdot\|_*$ is the nuclear norm, or the sum of the singular values, $\lambda_0 = \lambda/\sqrt{\max(m, n)}$, where λ is a parameter to tune the sparsity of the \mathbf{S} matrix. This formulation allows the noise to be contained in the sparse matrix and separated from the underlying data, which is low-rank. Then, a POD performed on the low-rank matrix reveals structures within the dynamics, free of noise and outliers. In general, a well-tuned application of RPCA can allow for a simultaneous denoising and dimensionality reduction. In addition, it is useful in filtering out highly varying, or traveling, phenomena in the “background” of a data set.

1.2.3 Dynamic Mode Decomposition (DMD)

The dynamic mode decomposition (DMD) [23–26] is an alternative to the the POD reduction typically used in ROMs. It not only correlates spatial activity, but also enforces that various low-rank spatial modes be correlated in time, essentially merging the favorable aspects of

POD in space and the Fourier transform in time. In its simplest form [25], one can consider two snapshots of measurement data

$$\mathbf{X} = \begin{bmatrix} | & | & & | \\ \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_{m-1} \\ | & | & & | \end{bmatrix} \quad \text{and} \quad \mathbf{X}' = \begin{bmatrix} | & | & & | \\ \mathbf{x}_2 & \mathbf{x}_3 & \cdots & \mathbf{x}_m \\ | & | & & | \end{bmatrix} \quad (1.6)$$

where the primed data is advanced Δt into the future compared to its unprimed counterpart. Exact DMD computes the leading eigendecomposition of the best-fit linear operator \mathbf{A} relating the data

$$\mathbf{A} = \mathbf{X}'\mathbf{X}^\dagger. \quad (1.7)$$

where \dagger represents the Moore-Penrose pseudo-inverse. This gives a least-square fit to the best linear model fitting the data whose solution is

$$\mathbf{x}_k = \sum_{j=1}^n \phi_j \lambda_j^k b_j = \Phi \Lambda^k \mathbf{b} \quad (1.8)$$

where ϕ_j and λ_j are the eigenvectors and eigenvalues of the matrix \mathbf{A} , and the coefficients b_j are the coordinates of the initial condition \mathbf{u}_0 in the eigenvector basis. The eigenvalues λ of \mathbf{A} determine the temporal dynamics of the system. It is often convenient to convert these eigenvalues to continuous time, $\omega = \log(\lambda)/\Delta t$, so the real parts of the eigenvalues ω determine growth and decay of the solution, and the imaginary parts determine oscillatory behaviors and their corresponding frequencies. The eigenvalues and eigenvectors are critically enabling for producing interpretable diagnostic features of the dynamics.

1.3 Related Dimensionality Reduction for Translational Symmetry

Much of the current literature seeks to overcome the difficulties posed by translational invariances in dimensionality reduction. This work also focuses on translational invariances, or traveling waves, for Chapters 2 and 3. Scaling invariance will be introduced in Chapter 4, where extensions for rotational and other invariances are also discussed.

Previous work [35–38] has focused on using modal decomposition techniques, such as POD, in single transport phenomenon. Others have also explored symmetry groups and coordinate transforms [35, 38–40] for decompositions, but these methods all require knowledge of the underlying equations, which are not always readily available. Instead, it is often desirable to focus on methods that will be effective on either simulated equations or experimental data with no closed-form solution.

Many approaches to creating ROMs for traveling waves have focused on determining sliding coordinate frames in which the translating quantities move. By finding this coordinate frame, the data can be transformed such that the traveling waves appear stationary, and therefore amenable to traditional dimensionality reduction methods. There are a number of novel approaches to determining this coordinate frame.

Many utilize a template fitting approach as demonstrated by Kirby and Armbruster [30] to shift waves into a transport-invariant coordinate frame. In this template fitting approach, the shift amount $c(t)$ is determined by projecting $u(x, t)$ onto a template $u_0(x)$, often taken to be the initial condition. This approach is well-suited to single waves which maintain a near-constant waveform. When data contains multiple traveling quantities, or waves exhibit breathing or dispersion phenomenon, the shift speeds cannot be effectively captured. In 2000, Rowley and Marsden [31] developed a method of reconstructing systems with translational symmetry by pre-shifting the POD modes by a speed $c(t)$ discovered through template fitting with the initial condition, resulting in:

$$u(x, t) = \sum_{r=1}^N a_r(t) \phi_r(x - c(t)). \quad (1.9)$$

While this approach may handle even complicated periodic systems well, such as the Kuramoto-Sivashinsky equation, the model creation is based inherently in the known dynamics of the system. This approach does not generalize well to systems with unknown governing equations, similar to methods developed by others in [35, 38–40]. Work by Iollo and Lombardi [35] accounted for diffusive phenomenon in traveling wave decomposition by relying on the governing equations to separate advective and diffusive components. Without

relying on equations of motion, Lucia et al. [36] were able to create ROMs for shock waves by separating out and reducing the transport-dominated domain. This approach may introduce an undesirable bias in the selection of the domains to decompose, and here we aim for an unsupervised approach.

Rim et al. [32] use the template fitting approach to determine the so-called shift numbers in which the wave is stationary. They perform a transport reversal to shift each time step of the snapshot matrix into the moving coordinate frame of the transport phenomenon. They also propose solutions to many of the difficulties introduced by a periodic shift operator and quick evolution of the transported structures. They present a scaling vector to account for changing wave heights, a cut-off map for non-periodic boundaries, regularization for a smooth and interpretable transport vector, and a greedy algorithm to account for multiple or rapidly-changing transported quantities. The authors show examples of the linear wave equation for demonstration, Burgers' equation to show shock formation, and the acoustic equation with heterogeneous materials to exhibit variable wave speeds. These innovations are promising computational tools that demonstrate creative solutions to many of the issues presented by the pre-shifting approach. However, interpretability of the shift speeds and underlying dynamics of the system can be lost with such an approach.

Most closely related to this work is the work of Reiss et al. [33], who introduced a data-driven decomposition for multiple traveling waves. The *shifted Proper Orthogonal Decomposition* (sPOD) views the snapshot matrix $u(x_i, t_n)$ of discretized spatiotemporal data in N_s moving coordinate frames, decomposing each with POD:

$$u(x_i, t_n) \approx \sum_{k=1}^{N_s} T^{c_k} \left(\sum_r \alpha_r^k(t_n) \phi_r^k(x_i) \right). \quad (1.10)$$

The sPOD algorithm uses the discrete shift operator T^{c_k} , corresponding to the transport speeds c_k , to pre-shift the data. After aligning data into a moving coordinate frame, POD is iteratively used to determine the modes for each transport speed until a prescribed error tolerance is met. The shift operator is found in one of two ways: using a thresholding algorithm or using an SVD-based search in a prescribed interval. In the thresholding method,

the threshold must be chosen ad hoc, and a potentially noisy or non-smooth set of shift values is returned. In the SVD-based shift search, many candidate constant shift values c must be supplied, and the SVD computation is performed for each in order to determine which leads to the maximum leading singular values (or quickest decaying singular value spectrum). This leads to a potentially high computational cost and a lack of robustness to non-constant wave speeds. In addition, both of these supervised methods yield black box shift values that are not robust to drastically evolving wave shapes or non-periodic boundaries. Regardless of these limitations, Reiss et al. [33] provides an exceptional step forward for handling traveling waves in PDEs. More recent work [41] employs this method on a pulsed detonation chamber to isolate shock waves within the combustor. Central to many of these methods is the concept of determining traveling coordinate frames into which the data can be shifted to align translating structures.

Chapter 2

DEVELOPING AN OPTIMIZATION FOR MODEL DISCOVERY

While methods described in Chapter 1 have provided a framework for the shift-discovery approach, there are key drawbacks from them: the need for underlying knowledge, and the lack of interpretability. In this work, we aim to create a method to determine these wave speeds without knowledge of the underlying system, and also in a way that will bring to light new, interpretable insights about the driving dynamics. Therefore the three main goals of the algorithm we develop are to (1) determine the wave speeds, (2) do so with as few “expert” inputs as possible, and (3) discover some interpretable information about the underlying physics in the system.

Our mathematical architecture provides a method for pre-shifting traveling wave data in order to improve the performance of standard dimensionality reduction techniques for ROMs. We develop an optimization framework for discovering interpretable wave speed models for multiple, non-constant speed waves, called UnTWIST (*Unsupervised Traveling Wave Identification with Shifting and Truncation*).

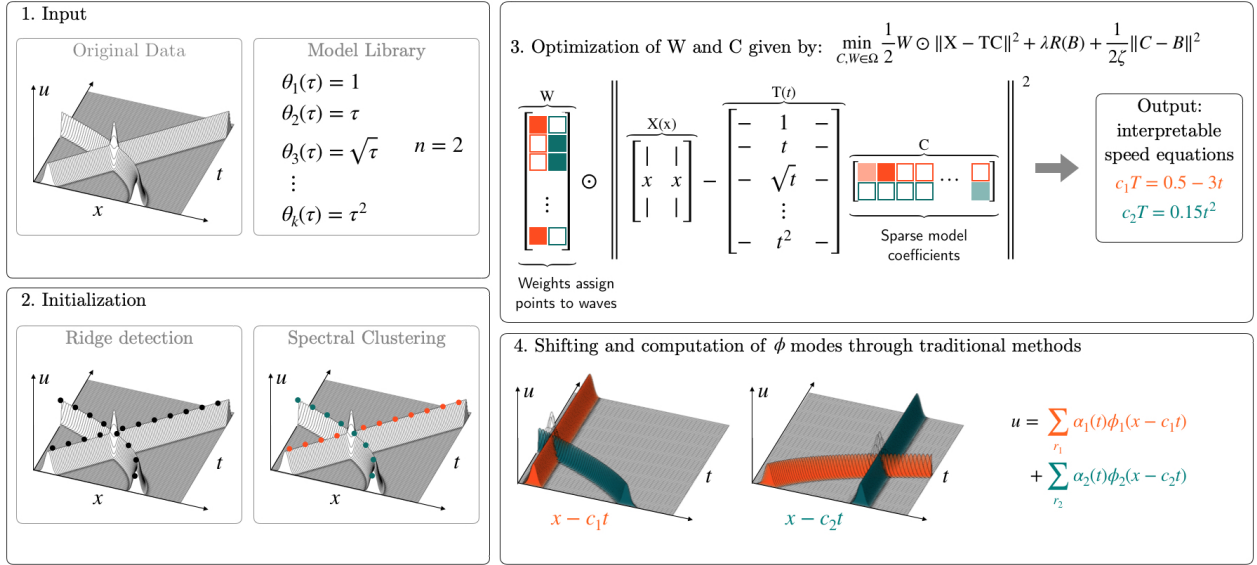


Figure 2.1: Schematic of the optimization-based shift mechanism. (1) The input data is passed in as a 2D snapshot matrix \mathbf{U} . A model library is constructed by choosing simple functions of time to describe candidate wave speeds, and a number of waves is chosen. (2) To initialize, a ridge detection is performed and finds the locations of the wave ridges as (\mathbf{x}, \mathbf{t}) pairs. Then, spectral clustering is performed to preliminarily assign the (\mathbf{x}, \mathbf{t}) pairs to each wave, shown separated into orange and blue waves. (3) An optimization is performed over the \mathbf{W} and \mathbf{C} matrices, with \mathbf{B} acting as an auxiliary variable. \mathbf{W} governs the assignment of the (\mathbf{x}, \mathbf{t}) pairs to each wave. \mathbf{C} gives the coefficients of the library terms that account for each wave speed, orange and blue. This matrix is penalized to be sparse so that only a few library terms are chosen out of many candidates. The speeds for each wave are interpreted as each row of the \mathbf{C} matrix. (4) The original data is shifted into each wave’s reference frame, making each a stationary wave. Traditional model reduction techniques can then be used to find ϕ modes for each.

2.1 Formulating the Optimization Framework

For a traveling wave with time-varying speed $c(t)$, the function $u(x, t) = f(x - c(t))$ describes the waveform over space and time. The first task is to determine a set of \mathbf{x} and \mathbf{t} pairs of

wave peak points that follow the path of the traveling wave over time. Ridge detection is a ubiquitous and well-studied task in computer vision [42], by which we can determine these points. The number of wave peak points N is arbitrary, but a more refined discretization of space and time will yield higher N and more accurate peak point detection.

Simple Case. In the case of a single traveling wave with constant speed c , we can formulate the optimization as $\mathbf{x} - c\mathbf{t} = 0$. Determining the scalar value of c that best describes these points is a simple least-squares regression. In most cases of interest, however, wave speeds are not constant. Rather, they may be non-constant, nonlinear, or have some prescribed time dynamics in the form of $\mathbf{x} - c(\mathbf{t}) = 0$, as seen in the spatial modes ϕ of (1.9).

Non-Constant Speeds. In order to describe waves with speed functions $c(\mathbf{t})$, we form a matrix \mathbf{T} of k candidate wave functions of time, for example, $\mathbf{1}$, \mathbf{t}^2 or $\sin(\mathbf{t})$, that may describe the wave’s non-constant speed over time. The choice of library functions allows the user enormous flexibility to customize the method to their application. The functions can be thoughtfully chosen based on visualization or similarity to other known dynamical systems, though this is not necessary. Using a rich library of simple terms such as polynomial and trigonometric functions usually yield strong candidate models. The input variable n gives the number of waves present in the wave field. We repeat our \mathbf{x} vector n times to form the \mathbf{X} matrix. Searching for a wave’s speed is a least-squares problem which yields *coefficients* $\mathbf{C} \in \mathbb{R}^{k \times n}$ to potential wave speed functions:

$$\min_{\mathbf{C}} \frac{1}{2} \|\mathbf{X} - \mathbf{TC}\|_2^2. \tag{2.1}$$

The matrix \mathbf{C} is composed of rows corresponding to each term in the candidate function library \mathbf{T} . Columns of \mathbf{C} correspond to each wave. In general, we will seek the simplest model for the wave speed, given by the fewest terms or the sparsest matrix \mathbf{C} , similar to the sparse identification of nonlinear dynamics (SINDy) [43].

Multiple Waves. When more than one wave is present in a wave field, however, a mask must be used to differentiate between the wave speed model for each wave. The weighting matrix \mathbf{W} is implemented to isolate rows of \mathbf{X} and \mathbf{T} , corresponding to wave peak points, which are active for each wave. Here, $\mathbf{W} \in \mathbb{R}^{N \times n}$ assigns N points to n waves, with each row of \mathbf{W} encoding a normalized probability distribution for each wave peak point. Entries are forced toward binary values of 0 or 1, corresponding to a wave peak point belonging to only one wave with high probability. We formulate the optimization to include multiple waves and non-constant wave speeds as

$$\min_{\mathbf{C}, \mathbf{W} \in \Omega} \frac{1}{2} \mathbf{W} \odot \|\mathbf{X} - \mathbf{TC}\|_2^2, \quad (2.2)$$

where the symbol \odot denotes the Hadamard element-wise product.

Parsimonious Models. Ultimately, the success of this method relies on accommodating a wide variety and arbitrary number of candidate functions to express the proper wave speed. However, only a small subset of these candidate functions will be relevant for characterizing a real system; even the most complex physically-based systems, such as Navier-Stokes, contain relatively few terms. With this in mind, we introduce a *sparsity* promoting constraint, usually in the form of a ℓ_0 - or ℓ_1 -norm [43]. This additional constraint penalizes a dense matrix and forces many terms to zero. We apply this sparsity constraint on the \mathbf{C} matrix so that many of the coefficients of the candidate model functions become zero, i.e. the wave speeds have few active terms in the library. If this sparsity constraint were not imposed, the \mathbf{C} matrix generically contains many high-order terms since this minimizes a least-square ℓ_2 -norm. This is seldom, if ever, physically meaningful. Further, it generally leads to instability for future-time prediction, and is prone to over-fitting.

2.1.1 The UnTWIST Minimization

Enforcing a sparsity constraint on the wave speed models makes the optimization challenging, and in some cases, non-convex. While many methods have been developed to address the

pursuit of sparsity, such as the ubiquitous LASSO [44], *Sparse Relaxed Regularized Regression* (SR3) [45] was chosen because of its adaptability to different non-convex sparsity constraints and its computational efficiency. The SR3 formulation of the optimization problem to fit speed models to multiple waves with non-constant candidate functions is given by

$$\min_{\mathbf{C}, \mathbf{B}, \mathbf{W} \in \Omega} \frac{1}{2} \mathbf{W} \odot \|\mathbf{X} - \mathbf{T}\mathbf{C}\|_2^2 + \lambda R(\mathbf{B}) + \frac{1}{2\zeta} \|\mathbf{C} - \mathbf{B}\|_2^2. \quad (2.3)$$

The regularizing function $R(\mathbf{B})$ is chosen to be a sparsity constraint, for example the ℓ_0 - or ℓ_1 -norm is used so that $R(\mathbf{B}) = \|\mathbf{B}\|_0$, for instance. The number of nonzero terms is controlled by the value of the λ parameter. The matrix \mathbf{B} is an auxiliary variable which relaxes the sparsity constraints by being enforced here rather than on \mathbf{C} directly. \mathbf{B} is forced to be close to \mathbf{C} , with tuning from the relaxation parameter ζ . Tuning the hyperparameters λ and ζ carefully is critical to reaching a meaningful minimum, although there is no *a priori* knowledge of the optimal values. Generally, λ is considered the sparsity parameter and should be chosen such that only handful of terms in the library are allowed to remain active. The ζ parameter can be considered a step size and should be chosen such that the optimization changes significantly enough in error for each iteration, but that it does not take too many iterations to converge. In the examples presented, a judicious choice of hyperparameters was made for the sake of simplicity. However, this could be done in a principled fashion by using a parameter sweep, and choosing optimal values. Hyperparameter tuning is standard in most machine learning algorithms [46]. The SR3 optimization architecture has proven accurate, robust, and efficient in non-convex problems with sparsity constraints such as the one presented.

2.1.2 Implementation of UnTWIST

Algorithm 1 details the UnTWIST algorithm. Algorithm 2 details the corresponding sparsity promoting SR3 algorithm used for selecting a wave speed model.

It is important to note that initialization of our algorithm plays a key role in the success

Algorithm 1 UnTWIST

Input: \mathbf{U} : wave field, x and t : space-time discretization, n : number of waves, $\boldsymbol{\theta}_{1:k}$: library of candidate speed models, and λ and ζ : optimization hyperparameters

Output: \mathbf{A}

Initialize:

- 1: $(\mathbf{x}, \mathbf{t}) \leftarrow \text{RidgeDetection}(\mathbf{U})$ ▷ Find wave peak points, see Figure 2.1.2
 - 2: $\mathbf{T} \leftarrow \boldsymbol{\theta}_{1:k}(\mathbf{t})$ ▷ Form library from candidate speed functions, from Figure 2.1.1
 - 3: $\mathbf{X} \leftarrow \text{repmat}(\mathbf{x}, k)$
 - 4: $\mathbf{W}_0 \leftarrow \text{SpectralClustering}(\mathbf{x}, \mathbf{t}, n)$ ▷ Find probable wave groups, see Figure 2.1.2
 - 5: $\mathbf{C}_0 \leftarrow \text{ThresholdedLeastSquares}(\mathbf{W}_0, \mathbf{x}, \mathbf{t})$ ▷ Initialize models for each wave speed
 - 6: $(\mathbf{B}_{1:\nu}, \mathbf{C}_{1:\nu}, \mathbf{W}_{1:\nu}) \leftarrow \text{SR3}(\mathbf{C}_0, \mathbf{W}_0)$ ▷ Optimize using SR3, detailed in Algorithm 2, see Figure 2.1.3
 - 7: $\mathbf{U}_{1:n} \leftarrow \text{Shift}(n, \mathbf{U}, \mathbf{C}_\nu, \mathbf{T})$ ▷ Shift in each reference frame, see Figure 2.1.4
 - 8: $\sum_{i=1}^r \alpha_i \phi(x)_i \leftarrow \text{POD}(\mathbf{U}_{1:n})$ ▷ Any dimensionality reduction method may be used, see Figure 2.1.4
-

Algorithm 2 SR3 for UnTWIST

- 1: **Input:** $\mathbf{C}_0, \mathbf{W}_0, \lambda, \zeta$
 - 2: **Initialize:** $\nu = 0, \mathbf{B}_0 = \mathbf{C}_0$
 - 3: **while** not converged **do**
 - 4: $\nu \leftarrow \nu + 1$
 - 5: $\mathbf{C}_{\nu+1} \leftarrow \text{argmin}_{\mathbf{C}} \left(\frac{1}{2} \mathbf{W} \odot \|\mathbf{X} - \mathbf{T} \mathbf{C}\|^2 + \frac{1}{2\zeta} \|\mathbf{C} - \mathbf{B}\|^2 \right)$
 - 6: $\mathbf{B}_{\nu+1} \leftarrow \text{prox}_{\lambda \eta R}(\mathbf{C}_{\nu+1})$
 - 7: $\mathbf{W}_{\nu+1} \leftarrow \text{proj}_{\Omega}((\mathbf{X} - \mathbf{T} \mathbf{C}_{\nu+1}) \odot (\mathbf{X} - \mathbf{T} \mathbf{C}_{\nu+1}))$
 - 8: **Output:** $\nu, \mathbf{C}_{1:\nu}, \mathbf{B}_{1:\nu}, \mathbf{W}_{1:\nu}$
-

of this method over a very broad optimization space. To initialize the weighting matrix \mathbf{W} , as seen in a clustering performed on the initial \mathbf{x} and \mathbf{t} data, assigning initial probabilities

that each wave peak point belongs to a given wave. Spectral clustering [47] is employed to obtain these initial probabilities, relying on the eigenvalue spectrum of a similarity matrix of the points. This algorithm was selected for clustering because of its ability to separate continuous groups of data, as in a wave’s traveling crest, rather than points that are close in Euclidean distance. Points in a given cluster are assigned a weight of unity in a given column of \mathbf{W}_0 and zero elsewhere, with each column of \mathbf{W}_0 corresponding to a different cluster, or wave.

The \mathbf{C} matrix of speed function coefficients was initialized using a sequential thresholded least-squares algorithm, as in [43]. In this scheme, the least-squares solve is performed for each row of the \mathbf{C} matrix, corresponding to each wave. The smallest terms will be thresholded out and “deactivated” to enforce sparsity, another least-squares is performed, and the process is repeated until the sparsity constraint is satisfied. This yields a sparse initial model for each wave as coefficients in \mathbf{C}_0 .

The SR3 optimization is comparable in computational cost to the ADMM algorithm, a popular and ubiquitous optimization scheme, with a one-time overhead of $\mathcal{O}(Nn^2 + n^3)$ and iteration cost of $\mathcal{O}(n^2)$. It is important to note that here, this cost does not scale with the size of the original data. Rather, it scales with the number of points selected by the ridge detection, i.e. the size of the \mathbf{X} and \mathbf{T} matrices, which is typically much smaller than the input. In future use for higher dimensional systems, this cost would not scale proportional to the additional dimensions but with the same points denoting the translated quantity, which would be higher but on the order of magnitude of the 1D examples. One-time computations in the initialization include 10 total operations on the original data for the ridge detection [42], a k -means clustering, which scales with $\mathcal{O}(Nn)$.

The optimization given by (2.3) yields the \mathbf{C} matrix, which is used with the library \mathbf{T} to determine the interpretable speed models for each of the waves. These speed models are used to shift each time slice of the input data matrix. This effectively places the wave field into a traveling reference frame in which one wave appears stationary. Traditional decompositions such as POD are then used and result in meaningful modes for ROM reconstruction.

2.2 Results on Example Systems: Model Discovery

The results of the UnTWIST method on example systems, from physically meaningful PDEs and example traveling wave systems, are described here. First, the results of the sparse wave speed discovery are given to highlight the interpretable physics that are learned from each data set. Then, reduced-order models are presented to showcase the effectiveness of the pre-shifting based on the learned wave speeds as compared to the traditional, unshifted approach.

2.2.1 Korteweg-de Vries Equation

The first example is one of the simplest examples of a traveling wave: a single soliton of constant speed propagating over a domain with periodic boundary conditions. The evolution of the wave field governed by the Korteweg-de Vries (KdV) equation is given by

$$u_t + uu_x + u_{xxx} = 0 \quad (2.4)$$

with initial condition

$$u_0(x) = 1 + 24 \operatorname{sech}^2(\sqrt{2}x). \quad (2.5)$$

The KdV equation is widely used to describe gravity waves propagating through shallow water [48]. The initial condition (2.5) is a known soliton solution of the KdV equation, which maintains a constant wave shape, with a wave propagation speed of 9. This initial condition was discretized into 512 grid points in $[-\pi, \pi)$ and solved for $t \in [0, 0.6]$ using an exponential time-differencing fourth-order Runge-Kutta scheme [49], where results were saved at steps of $\Delta t = 0.001$ in time. The initial data can be seen in the surface plot in Figure 2.2 (left panel).

A library of potential wave speeds, as given in Figure 2.2, was chosen to be polynomials of the time variable up to third order, with a constant included for centering the data in x . The number of waves was chosen to be $n = 2$ in order to allow the algorithm to converge to a meaningful and interpretable wave speed model. A single wave model could have been

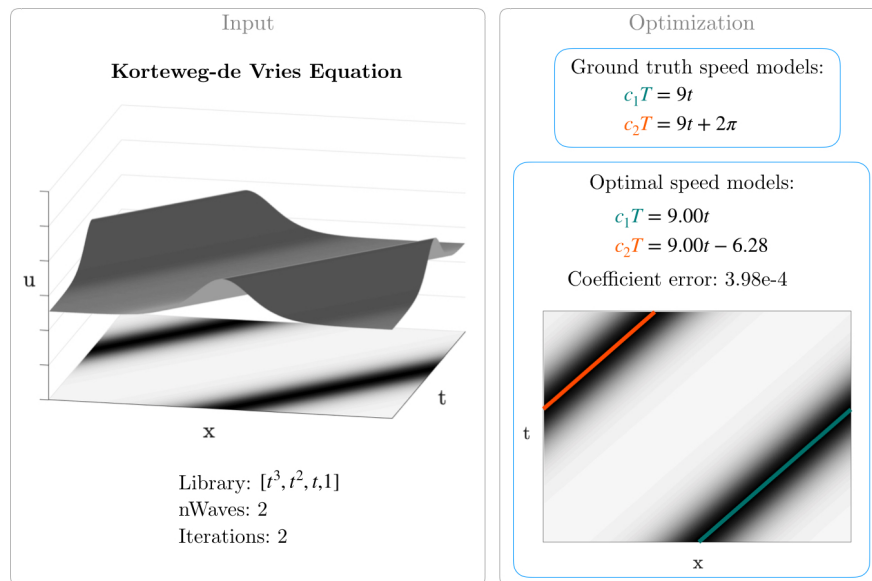


Figure 2.2: Results of the UnTWIST algorithm on the simulated KdV equation input data. The algorithm converged to a coefficient error of $3.98e-4$ within 2 iterations, and is close to the ground truth speed used in simulation.

used, and in this case, found to be a parallel wave directly in between the periodic extension of the right-traveling wave. Hyperparameters were tuned to be $\lambda = 10$ and $\zeta = 1e-4$.

For this simple example, where the wave speeds are constant in time, the UnTWIST algorithm performed with high accuracy, yielding coefficients within $3.98e-4$ of the ground truth wave speed models, written explicitly in Figure 2.2. The two wave crests were identified clearly with the ridge detection step. The algorithm converged in only two iterations, indicating that the initialization via spectral clustering and thresholded least squares solved the optimization nearly to tolerance, and only one optimization step was taken. Figure 2.2 also shows that the wave speed models, given in orange and green lines, closely follow the wave crests.

2.2.2 Two Crossing Solitons: Non-periodic and Non-Constant Speed

The second example we present is a fabricated non-periodic wave field of two crossing Gaussian solitons of non-constant speed. The wave field was constructed using the governing equation

$$u(x, t) = \exp(-0.1(x + 3t - 80)^2) + \exp(-0.1(x - 0.15t^2 - 20)^2) \quad (2.6)$$

on $x, t \in [0, 100] \times [0, 20]$ in a uniform (256×512) grid. The library is taken to be polynomial terms of t up to second-order. Hyperparameters were tuned to be $\lambda = 10,000$ and $\zeta = 1e-4$. The number of waves n was chosen to be 3 as an initial value. The UnTWIST algorithm allows the number of waves to change if the number of points in a cluster shrink below a threshold in the optimization step. In this example, giving a starting value of $n = 3$ allowed a subset of points to be in a temporary third wave group until a more probable wave belonging is found. It can be seen in Figure 2.3 that the addition of a third wave group allows the yellow points, corresponding to the complicated wave intersection area, to belong to a wave group of their own. After the optimization converged in 2,743 iterations, the third wave group has fewer points than the threshold allows, so this model is removed. These points are assigned to the right-traveling wave, and the models closely fit the wave crests. The speed models also have close agreement, with the coefficient error at $1.15e-3$ compared to the ground truth speeds.

2.2.3 Burgers' Equation

The third example we present uses viscous Burgers' Equation, (2.7). Because of its characteristic of forming sharp discontinuities, viscous Burgers' equation is used to describe shock waves common in gas dynamics [48]. This provides an example of a changing wave profile alongside a non-constant wave propagation speed, as seen in Figure 2.4. This example does not provide a known speed, but it is well-known that the wave front propagates at a speed proportional to the wave height.

$$u_t + uu_x - 0.1u_{xx} = 0. \quad (2.7)$$

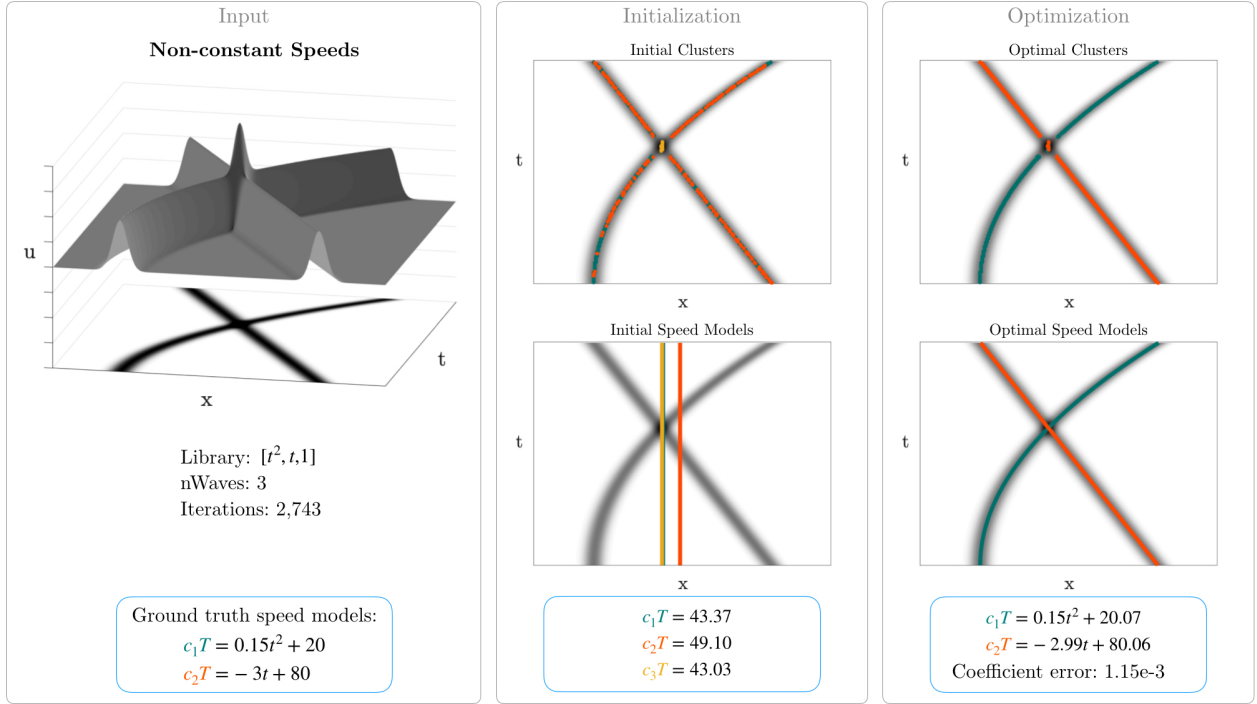


Figure 2.3: Results of the optimization on artificially generated data of non-constant velocity linear soliton waves and a library as given.

This equation was solved using fourth-order Runge-Kutta scheme for $t \in [0, 20]$ in steps of $\Delta t = 0.1$ in a periodic domain of $x \in [-8, 8)$ with 256 grid points. The initial condition is given in (2.8), representing a Gaussian wave.

$$u_0(x) = \exp(-(x + 2)^2). \quad (2.8)$$

The library was chosen to be a linear, square-root, and constant term. The square-root term was chosen in order to accommodate the nonlinear speed simply. The number of waves was chosen to be $n = 1$. Hyperparameters were tuned to be $\lambda = 10$ and $\zeta = 1e-4$. As seen in Figure 2.4, the wave speed model determined by the optimization fits closely to the data, converging in only 15 iterations. Since there is no known ground truth, the wave speed model was also compared to the wave amplitude over time, as seen in Figure 2.4. The close

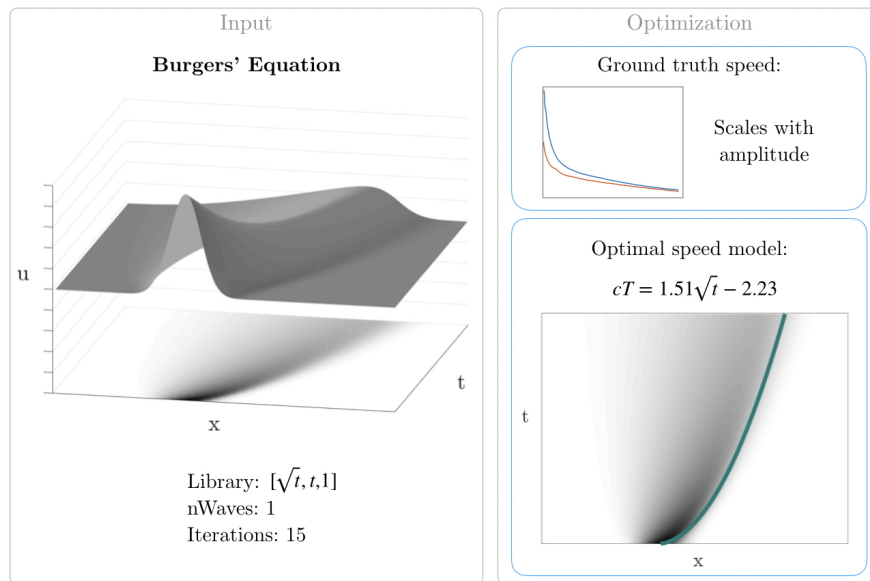


Figure 2.4: Results of the UnTWIST algorithm on data from Burgers' equation. Speeds were determined and compared to the height of the wave as determined by ridge detection, and follow the same trend.

agreement indicates that the wave speed was well-identified. This example also illustrates the success of the ridge detection algorithm for this application, as wake peak points are captured well, even with changing wave height.

2.2.4 *Nonlinear Schrödinger Equation*

The final example we present showcases many of the complications we wish to address with the UnTWIST algorithm— nonlinear wave interaction, non-periodicity, and rapidly-changing wave profiles. The complex nature of the wave field can be seen in Figure 2.5. The Nonlinear Schrödinger (NLS) Equation is used across many fields. In fluid mechanics, it describes gravity-driven surface water waves in the deep water regime [48], and is typically solved

using expensive and finely-tuned numerical methods. NLS is given as

$$iu_t = -\frac{1}{2}u_{xx} + \kappa|u|^2u. \quad (2.9)$$

This was solved on the domain $x, t \in [-15, 15) \times [0, 2\pi]$ in a 1024×501 uniform grid. The initial condition

$$u_0 = 2 \operatorname{sech}(x + 7) \exp(2ix) + 2 \operatorname{sech}(x - 7) \exp(-2ix) \quad (2.10)$$

was propagated forward in time using a pseudo-spectral, fourth-order Runge-Kutta scheme [3].

For this example, the library of candidate functions was made up of polynomials up to third order including a constant term. Similar to the fabricated data in Example 3, this wave field has complicated wave crossings, so an extra wave group was initialized with $n = 3$ waves. Hyperparameters were tuned to be $\lambda = 10$ and $\zeta = 1e-4$. As seen in Figure 2.5, the initialization did not cluster the points into a centered wave-crossing section as expected, rather, they are scattered. The initial models also show very poor correlation to the wave crests. After the optimization with 2,748 iterations, however, the models show a very close agreement to the wave crests and the extraneous wave cluster is eliminated. The discovered wave speed models also show close agreement with a coefficient error of 0.028.

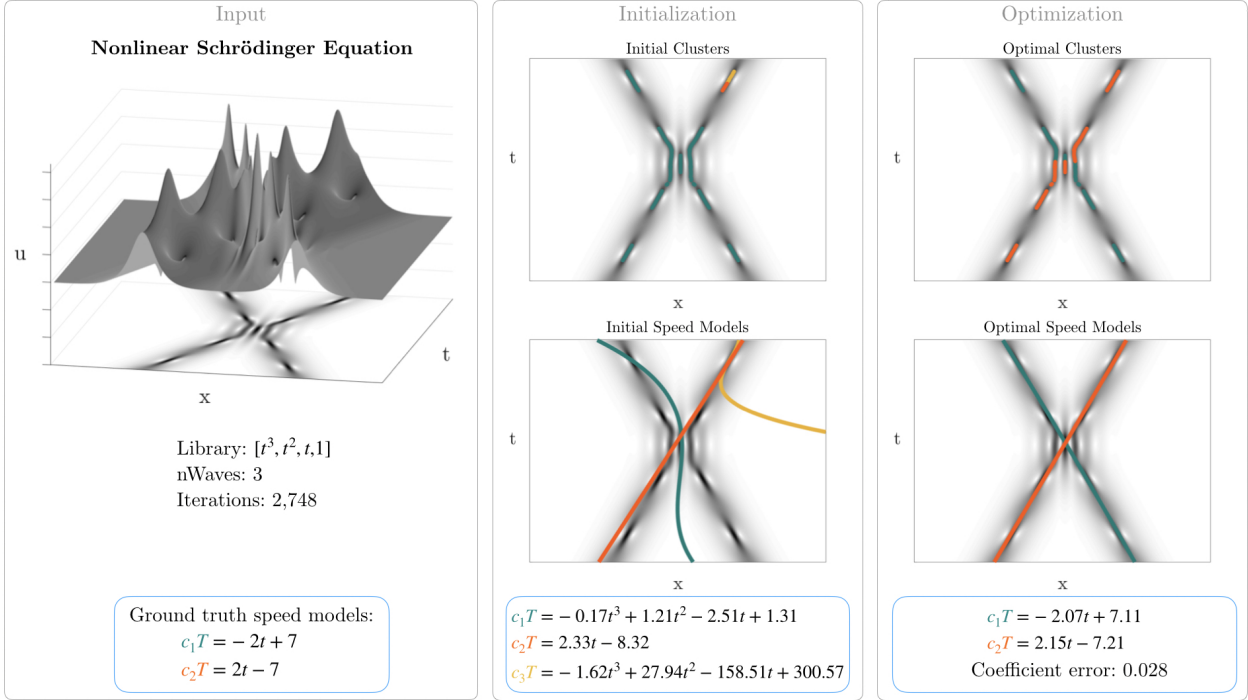


Figure 2.5: Results of the UnTWIST algorithm to find the shifts in simulated NLS equation data. Three clusters are used to allow the algorithm to threshold out complicated areas at first. The initial clustering is poor, as are corresponding initial models. However, after the optimization, the clustering has improved, and the models are visually an excellent fit. The extraneous model is ignored and coefficient error is 0.028 after 2,748 iterations.

2.3 Results on Example Systems: Dimensionality Reduction

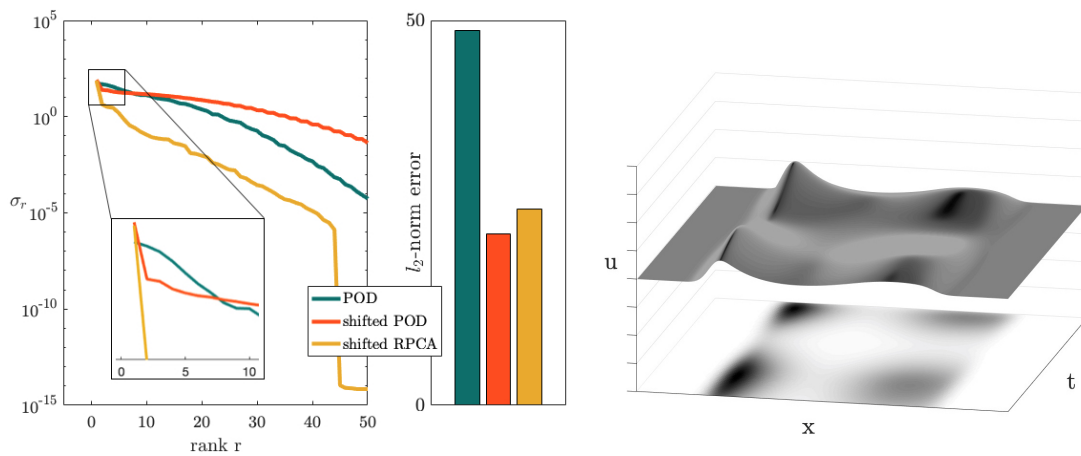
Low-rank reconstructions of two of the example wave fields using UnTWIST are now discussed. UnTWIST is coupled with two common dimensionality reduction methods, POD and Robust Principal Component Analysis (RPCA) [34]. These rank-reduced wave field reconstructions are compared with those of the traditional (unshifted) POD approach, each using only 2 modes. The singular value spectrum is also given for each example, showing the amount of information contained by each successive mode in the decomposition. It is

important to note that when shifting is applied, n wave fields result. For the shifted reconstructions, dimensionality reduction is performed for each aligned wave field separately, one mode is retained, and all wave fields are then superimposed.


The first example wave field is described in Example 2, presenting complications of non-constant wave speeds, non-periodicity, and multiple crossing waves. As seen above, the wave speeds are well-identified using the UnTWIST method. The data was pre-shifted into two frames using the two identified speed functions and both POD and RPCA were performed. These are presented in Figure 2.6 and compared to traditional POD with all methods truncating to rank-2.

The singular value spectra for each of these decompositions is show in Figure 2.6. The shifted RPCA captures the most energy out of the three methods out to around 75 modes, or about 30% of the original size. In this case, traditional POD outperforms the shifted POD, when viewing the total spectrum. For 2-8 modes of retention, shifted methods both outperform traditional POD. In the reconstructions, the two-mode UnTWIST more faithfully represents the dynamics of the system. The shifted RPCA performs the best, with two clear solitons interacting and maintaining smoothness outside this domain. The shifted POD somewhat recreates the wave pattern, but with artifacts obscuring the smooth soliton behavior. In the traditional POD example, however, the wave field is unrecognizable, with no continuity between the two solitons and no clear traveling behavior.


The second example we present is of Example 4, the NLS equation. The results of the decompositions are shown in Figure 2.7. The singular value spectrum shows that for this case, the shifted RPCA decomposition outperforms the others for all ranks of truncation. The shifted POD again outperforms the traditional POD for few modes, from 2-4 modes retained. The traditional POD quickly takes over the shifted, however, and captures more information for 5 and higher modes than the shifted version. Overall, shifted RPCA quickly captures the most possibly information in only about 35, or about 7% of all possible, modes. The rank-2 reconstructions illustrate this concept similarly. While the shifted RPCA does not have the characteristic sharp peaks of the original data, the area outside the breathers




(a) Singular values and error at rank-2



(b) Traditional POD



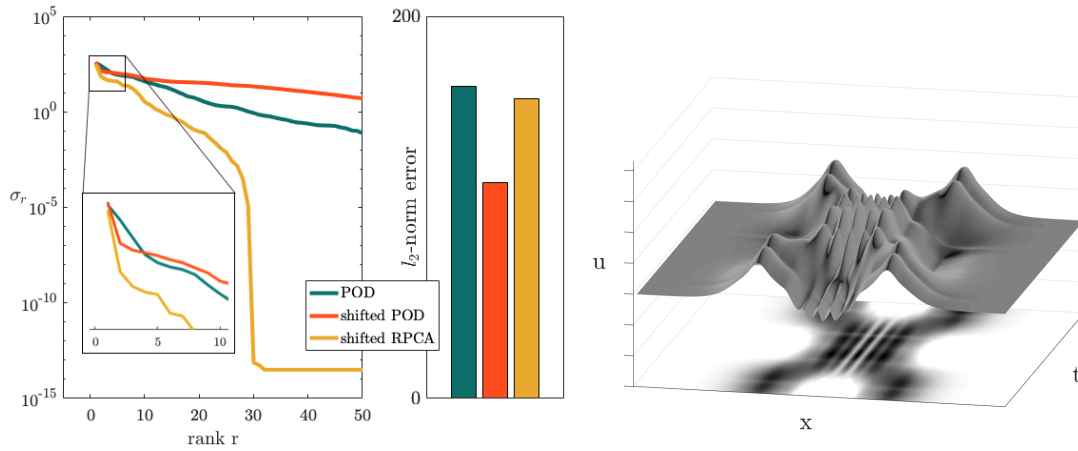
(c) Shifted POD




(d) Shifted RPCA

Figure 2.6: Comparison of different dimensionality reduction methods for generated example data, each using 2 modes. 2.6a shows the singular value spectra for each dimensionality reduction method for various ranks, showcasing the energy captured in each successive rank. Error plots at right show the l_2 -norm error for each method with a rank-2 reconstruction, whose surface plots are shown in 2.6b-2.6d.


is relatively smooth. In contrast, both the shifted POD and the traditional POD maintain oscillations near the edges of the breathers. These are a product of the many high frequency



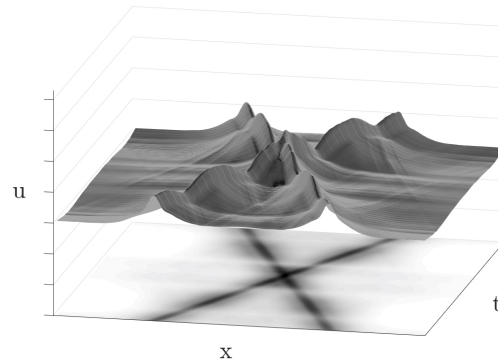
(a) Singular values and error at rank-2



(b) Traditional POD



(c) Shifted POD



(d) Shifted RPCA

Figure 2.7: Comparison of different dimensionality reduction methods for NLS data, each using 2 modes. 2.7a shows the singular value spectra for each dimensionality reduction method for various ranks, showcasing the energy captured in each successive rank. Error plots at right show the l_2 -norm error for each method with a rank-2 reconstruction, whose surface plots are shown in 2.7b-2.7d.

oscillatory modes needed to represent the breathing waves.

The success of these shifted reconstructions may hinge on the ability of the dimensionality

reduction algorithm to filter background from higher frequency phenomenon. When implementing a shift into a moving reference frame, a single traveling wave is viewed as stationary in time. When more than one wave is present in the field, however, shifting simply distorts the wave's speed. This problem is akin to a background separation— one low-frequency, stationary wave is the background of wave field, and the other is the foreground. When performing a dimensionality reduction, this mode is often the most energetic. Many dimensionality reduction methods, including RPCA, have focused on isolating these background modes, explaining the success of RPCA in this example.

Chapter 3

UNTWIST ON EXPERIMENTAL DATA: ROTATING DETONATION ENGINE

In Chapter 2, the UnTWIST algorithm was developed, and shows promising results in noiseless simulated data. Often, ROMs are necessary for experimental data, which may be extremely high-dimensional, noisy, and uncertain. The use of the UnTWIST algorithm is now explored on such experimental data, using the example of a rotating detonation engine.

3.1 Experimental Data: Rotating Detonation Engine

A *rotating detonation engine* (RDE) is a novel combustion engine that uses detonative heat release as the dominant mechanism of energy addition to the reactive, compressible fluid flow, contrasting deflagration-based, constant-pressure heat addition typical of aerospace engines. The RDE offers a number of advantages for application in propulsion or land-based power generation, including mechanical simplification, broad operability limits [50, 51], the potential for increased thermal efficiency [52, 53], and the reduction of propellant pumping requirements [54, 55]. The operating dynamics of the RDE include co- and counter-rotating coherent combustion wave fronts of varying number which interact to produce a rich set of nonlinear dynamics and instabilities. Recent modeling efforts have focused on phenomenological models [56, 57] that are capable of reproducing and characterizing the RDE dynamics and bifurcations observed in experiments. This includes models that characterize the nucleation and formation of combustion pulses, the soliton-like interactions between these combustion fronts, and the fundamental, underlying Hopf bifurcation to periodic modulation of the waves [57]. The goal of this chapter is to characterize the dynamics of the combustion wave front interactions directly from experimental data, specifically with the

goal of developing ROMs for characterizing the origins of dynamic instabilities in RDEs. We will explore several leading techniques in data-driven optimization (i.e., machine learning) of varying complexity.

RDE hardware is designed to amplify thermoacoustic instabilities associated with reacting flows in circular and/or periodic geometries. For thrust-producing RDEs, the typical design is an annular combustion chamber, see Figure 3.1a. Fuel and oxidizer are supplied through independent feeds into the head-end of the annulus, where they promptly mix to form a combustible medium. An ignition source (spark plug) initiates a chemical reaction that quickly and locally releases energy into the fluid. Supposing the geometry of the engine and the rate of heat release allow for a local accumulation of energy (Rayleigh’s criterion), sharp gradients in pressure and density (and therefore temperature) form. This creates a feedback loop where chemical kinetics are further accelerated by the increase in temperature, which in turn releases more energy into the fluid. This process saturates once all propellant is locally consumed and combustion halts. However, in the RDE, the sharp gradients in pressure and density form traveling shock waves strong enough to auto-ignite propellant. These shock-reaction structures, or detonation waves, move supersonically about the periodic chamber of the RDE, consuming the newly injected and mixed propellant in its path. The detonations continuously propagate so long as a sufficient amount of mixed propellant exists in its path to overcome dissipative effects (exhaust, for example). A number of experimental RDE programs have detailed the effects of geometry, injection schemes, and fueling conditions [51, 58–60] on the RDE dynamics.

The detonations follow attractor-like dynamics that are the manifestation of underlying multi-scale balance physics of the driven-dissipative RDE [61]. The RDE is similar in nature to mode-locked lasers [62, 63], where global gain and loss dynamics produce a similar cascading bifurcation diagram of mode-locked states [64]. In this context, the mode-locked structures of the RDE are classified as autosolitons, or stably-propagating nonlinear waves where the local physics of nonlinearity, dispersion, gain, and dissipation exactly balance. These physics are multi-scale in nature: the local fast scale of combustion provides the energy

input to generate the mode-locked state, while the slow scales of dissipation and propellant regeneration shape the waveform and dictate the total number of detonation waves. Thus, the global multi-scale balance physics give the detonations their mode-locking properties - not exclusively the frontal dynamics prescribed by classical detonation theory.

These properties have been experimentally observed at the University of Washington High Enthalpy Flow Laboratory using a gaseous methane-oxygen 76-mm flowpath outside diameter RDE, as described in previous works [56, 57, 65]. This experimental apparatus is unique in that the RDE tested is fully modular and that the apparatus exists in a closed system. The modularity of the RDE allows for parametric testing of engine geometries (flowpath lengths and annular gaps) and injectors (varying injection scheme, orifice count, and total injection area) with respect to varied propellant feed rates and stoichiometry. Because the entire apparatus is closed, implied is both the inlet and outlet boundary conditions of the combustor are able to be set, and are controlled here to give rise to stable traveling wave dynamics. In a successful experiment, the self-organization of traveling waves occurs and persists so long as propellant is flowing into the combustor. For each experiment, a high-speed camera records the duration of the ‘hot’ portion of the run, including the ignition event, the transient mode-locking phase, and steady operation of the combustor. The experiments exhibited in this manuscript are representative of modes of operation and transients observed in this experimental apparatus. These experimental spatiotemporal dynamics are taken from Koch et al. [56, 57, 65].

Computational fluid dynamic (CFD) simulations have been heavily relied upon to diagnose the RDE flowfields. These simulations vary from periodic 2-D ‘unwrapped’ rectangular domains [66–68] to full detailed 3-D engine geometries [69–72]. From these simulations, the canonical RDE flowfield is obtained (a cartoon of which is shown in Figure 3.1a) and relevant metrics can be extracted, such as thrust, specific impulse, available mechanical work, and thermodynamic efficiency [52, 73]. However, long-time parametric simulations of RDE dynamics is prohibitively expensive since the fastest physics (the detonation front) and the slowest physics (mixing and/or exhaustion) both need to be adequately resolved for proper

system behavior. Thus, simulations need to be run for several - if not dozens or hundreds - of cycles (or until the physics of the slowest scales are fully developed). The computational cost of simulations can quickly become prohibitive and it typically requires high-performance computing architectures for even moderate lengths of simulation time. Consequentially, ROMs have been developed, with varying degrees of success, for recreating the RDE canonical flowfield [74, 75], predicting thermodynamic trends [76], predicting application-based propulsive performance [77], or reproducing the dynamics of the waves [56, 57, 78, 79]. However, because of the multi-scale nature of the RDE and the intricate interactions of its fundamental physical processes, these modeling efforts are often constrained to geometry, propellant, or mode-specific operating regimes, with the imposition of wave topology or detonation structure. In an alternative approach, experiments allow us to build ROMs directly from data. To further ease computational burden in the study of RDE flowfields, recent work [56] has indicated that the relevant flow physics can be fully captured by in a single dimension, eliminating the need to compute over the full three-dimensional flow domain.

In order to construct ROMs of the combustion-front dynamics, one must first move to a frame of reference of the mode-locked states. ROMs exploit the intrinsic, low-rank structure of the simulation data in order to create more tractable models for the spatiotemporal evolution dynamics. Typically, ROMs leverage the *singular value decomposition* (SVD) to produce a linear dimensionality reduction [3, 14], whereby a dominant set of correlated modes provide a subspace in which to project the PDE dynamics [2, 15]. Low-energy modes are then truncated, and the governing equations are projected onto the remaining high-energy modes to create an approximate and low-dimensional model. Dimensionality reduction and modal decomposition approaches have been well-studied and are extremely efficient [2, 15, 27–29]. However, SVD-based ROMs are typically compromised by traveling wave physics, which represent an underlying translational invariance. Thus a growing body of literature is aimed at producing mathematical architectures that are capable of determining the traveling wave frame of reference of the underlying wave [30–33]. While these works are critical to addressing the shortcomings of traditional methods, they are limited to applications with constant wave

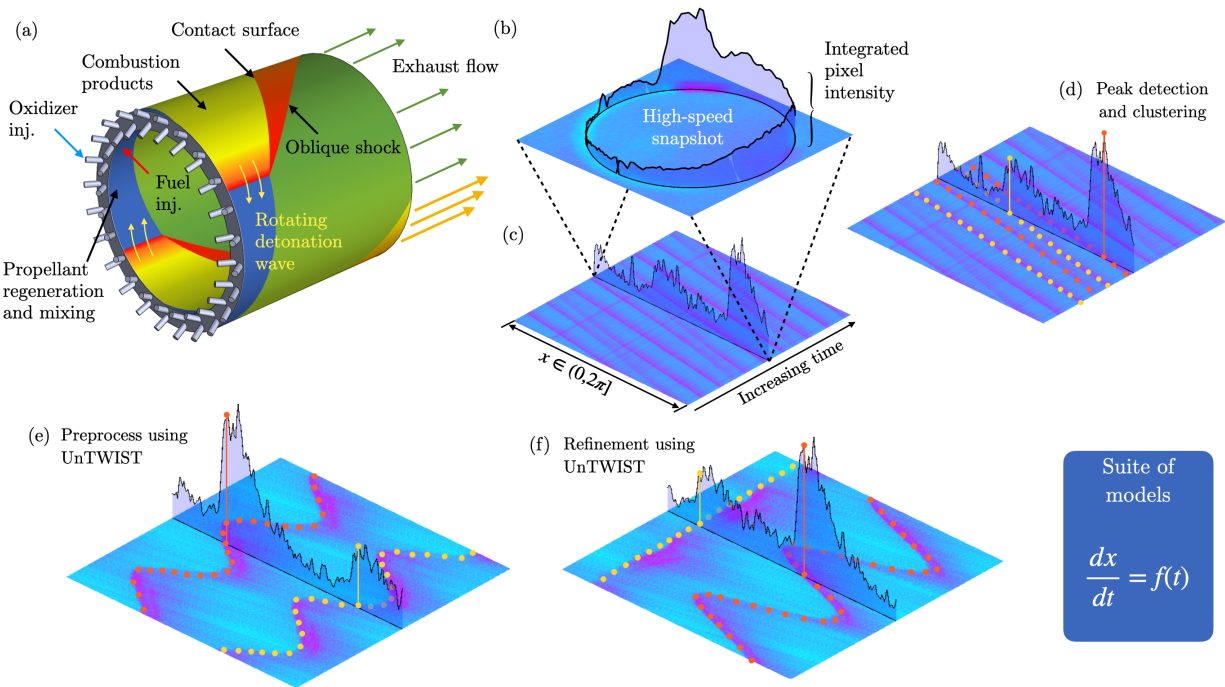


Figure 3.1: (a) RDE schematic, (b) Schematic of one time slice of video data, viewing down the axis of rotation of the RDE, (c) The same time slice viewed in an (x, t) plot, with each column of the data in time constructed by integrating the pixel intensity along the annulus, (d) demonstration of the peak detection and clustering necessary to model the wave speeds with UnTWIST, (e) a preliminary processing of (c) using the UnTWIST algorithm, (f) a refining processing of (e) with the UnTWIST algorithm, which then becomes the basis for a suite of data-driven models.

speeds or knowledge of the underlying physics. Chapter 2 introduces an unsupervised machine learning procedure for transport-dominated systems characterized by traveling waves. This method can be applied with or without knowledge of the governing equations, providing an interpretable mathematical architecture for ROMs exhibiting traveling wave phenomenon. This algorithmic infrastructure can be used to extract the intrinsic features associated with the RDE front evolution, uncovering a coordinate system where it is possible to obtain low-

order models. We then leverage a selection of machine learning algorithms to explore the dynamics prescribing the ubiquitous RDE front interactions. Importantly, the methodology is data-driven in that the ROMs are constructed entirely from detailed experimental observations. This work is part of a growing body of literature that is bringing emerging technology in machine learning to bear on problems in fluid mechanics [1, 9, 10].

3.2 Detonation Wave Tracking with UnTWIST

It is widely known that transport phenomenon such as traveling waves impair the effectiveness of traditional dimensionality reduction methods, mainly due to an issue of separation of space and time variables [14, 30–33]. While an approach like the method of characteristics can be used when governing equations are known, this experimental framework necessitates a system-agnostic method. One numerical approach to resolve this issue is to shift the frame of reference from the laboratory frame to a moving coordinate frame that matches the speed of the traveling waves. Once the traveling quantities have been made stationary in this way, efficient traditional methods such as proper orthogonal decomposition (POD) can be utilized for dimensionality reduction.

In order to build ROMs on the RDE data, rife with traveling shock fronts, it is necessary to preprocess it by aligning these traveling waves in time. Here, we employ the unsupervised traveling wave identification with shifting and truncation (UnTWIST) [80] algorithm to perform this preprocessing step. This method allows for a data-driven and interpretable model for the speeds of the traveling shock fronts, as well as separable low-rank modes, providing an intuitive insight into the physics of the system. A basic overview of UnTWIST and its implementation is described here. For further details and a complete algorithm, please refer to [80].

3.2.1 UnTWIST Method

As described in Chapter 2, UnTWIST learns a moving coordinate frame, given by the speed of a traveling wave. This holds the wave of interest stationary within the coordinate frame,

allowing for models to be built for that particular wave. Unique to UnTWIST is the ability to learn physically-relevant wave speeds directly from the data with little knowledge of underlying dynamics. It allows for a wide variety of physics, including linear, nonlinear, non-constant, and non-smooth wave speeds, to be considered. To do so, UnTWIST relies on an optimization over a user-input library of potential wave speed functions to learn this coordinate frame. This library may include any number of linear or nonlinear functions. If expert knowledge of the system is available, a judicious choice of functions is possible. In many cases, simple functions such as sinusoids and polynomials suffice. Inclusion of many candidate functions will increase the computational cost of the algorithm, therefore, a balance between completeness and the size of the library must be found.

To execute the optimization on wave profile data $u(x, t)$, two main steps are first performed: (1) ridge detection to learn the location in (x, t) space of the traveling wave fronts or peaks, and (2) spectral clustering to divide the points (x_i, t_i) into groupings for each wave. For example, these two steps are shown in Figure 3.2.




Figure 3.2: (a) Example of a traveling wave data set, (b) wave peak points (x_i, t_i) are identified using a ridge detection, shown overlaid with the waves, (c) wave peak points are clustered into wave groups using spectral clustering. Once these points are identified and clustered, a model is fit to each based off of a user-provided library of candidate linear or nonlinear functions.

Once the wave fronts are identified and separated, the data is assembled into the optimization. We construct matrices \mathbf{X} and \mathbf{T} using the (x, t) locations of the wave fronts in $u(x, t)$, where \mathbf{T} contains the values of t evaluated for each function in the user-defined library. The cost function is given by Equation (3.1)

$$\min_{\mathbf{C}, \mathbf{B}, \mathbf{W} \in \Omega} \frac{1}{2} \mathbf{W} \odot \|\mathbf{X} - \mathbf{T}\mathbf{C}\|_2^2 + \lambda R(\mathbf{B}) + \frac{1}{2\zeta} \|\mathbf{C} - \mathbf{B}\|_2^2, \quad (3.1)$$

where \mathbf{W} is the weighting matrix that serves to mask wave peak points for clustering into wave groups. With values of 0 or 1, each row of \mathbf{W} corresponds to each wave peak point (x_i, t_i) , and each column corresponds to a given wave. Values in \mathbf{C} are the coefficients of the speed models that are discovered for each wave. A row of \mathbf{C} corresponds to a wave, and columns give the coefficients of each term in the model library \mathbf{T} , which multiply together to generate the wave speed models. It is desirable for \mathbf{C} to be sparse, i.e. to have few nonzero terms, to glean an interpretable, physically realistic model for the wave speeds. Rather than placing a sparsity constraint on \mathbf{C} directly, the constraint can be relaxed by introducing an auxiliary matrix \mathbf{B} , which is close to \mathbf{C} . \mathbf{B} is directly forced to be sparse via a regularizing function $R(\cdot)$, relieving the burden on \mathbf{C} to meet both sparsity and accuracy goals. The hyperparameter λ is chosen to calibrate the sparseness of auxiliary matrix. The hyperparameter ζ is chosen to enforce the closeness of \mathbf{C} and \mathbf{B} , ensuring that the solution \mathbf{C} itself is also sparse. These two hyperparameters are tuned in tandem in order to meet sparsity and accuracy requirements of the model. This optimization presents a large search space over multiple parameters, and is not guaranteed to be convex. Sparse relaxed regularized regression (SR3) [81] is used to minimize the cost function because of its ability to handle non-convexity and its computational efficiency compared to similar sparsity-promoting optimization schemes.

Once the model coefficients \mathbf{C} are learned, they can be used to shift each time segment of data in order to align the data into one wave group’s moving coordinate frame. The mask matrix \mathbf{W} allows for easy separation of the wave fronts for this alignment. Once the data is aligned into the new coordinate frame, traditional dimensionality reduction methods can

easily be applied and can be expected to reveal extremely low-rank modes for the ‘straightened’ wave or wave group.

3.2.2 *UnTWIST Applied to RDE Data*

An example of the UnTWIST algorithm applied to snapshots of RDE data can be seen in Figure 3.3. UnTWIST was applied in two steps for each data set presented. The models we build are based on time series that are long relative to the spatial dimension— with between 1,000 and 10,000 time steps relative to 180 or 360 spatial points. The snapshots also contain wave fronts that travel on a fast time scale relative to the slow time scale of the relevant dynamics, see Figure 3.3a, necessitating an extreme shift in order to shift into a straightened wave coordinate frame. Because of the fast-moving fronts and long time series, the UnTWIST algorithm was applied in two steps— a preprocessing step, and a refining step, with different inputs for each.

For the preprocessing step, only 10 time steps of the data are considered, as seen in Figure 3.3b. Using the UnTWIST algorithm and the identified wave peak points as shown in Figure 3.3c, we obtain the best linear speed model for each wave, Figure 3.3d. A single linear-speed shift, the average of the speed models, is applied to original data for the entire time series, and gives Figure 3.3e. This first shift reveals critical underlying dynamics of the shock wave interactions.

After the first shift is performed, a second refining shift can be used with a more diverse library of potential wave speed models to completely align the data for building low-rank models. The refining shift is performed similarly to the first shift, but we now include potential wave speed functions such as sinusoids, polynomials, exponentials, and nonlinear combinations of these terms. For the example shown in Figure 3.4, sinusoids and exponentials were included in the candidate function library. Once the models were computed, the data was shifted into n refined coordinate frames, one for each shock wave.

The outcome of the second shift can be seen in Figure 3.4. Each coordinate frame allows

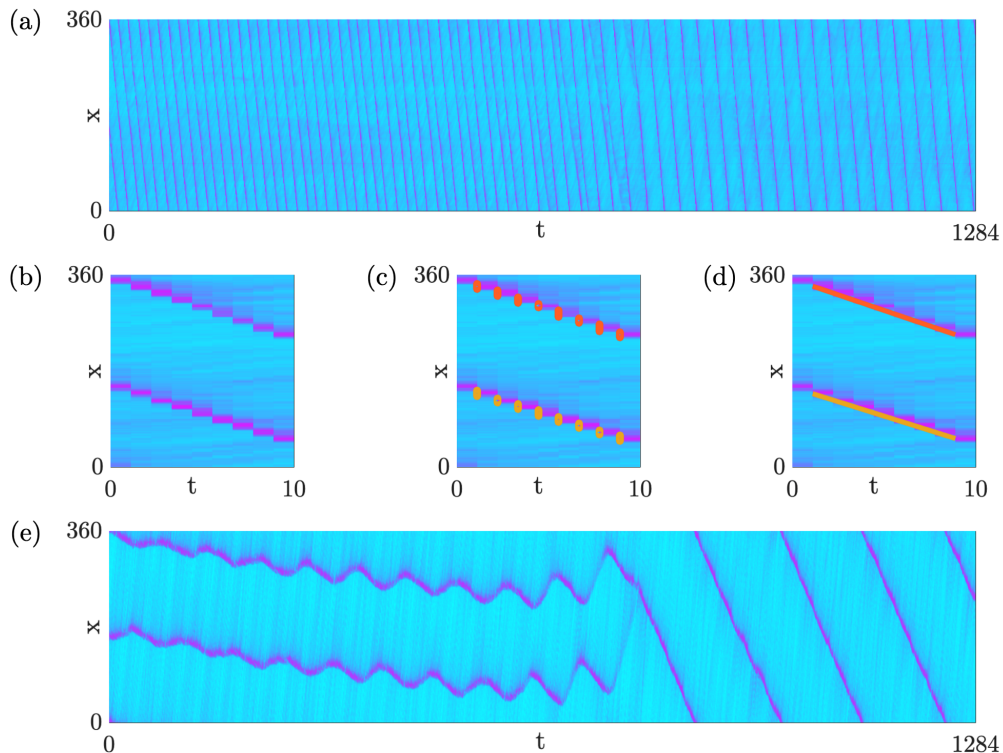


Figure 3.3: Preprocessing step using UnTWIST: (a) Example of an original data set, presenting fast-moving shock fronts and a long time series relative to the spatial dimension, (b) A 10 time step segment of (a) showing approximately linear-speed shock front propagation, (c) wave peak points identified and separated, (d) linear models of the shock wave speeds, (e) Data from (a) shifted into the average wave speed, dictated by the models identified in (d), which reveals the relevant interactions between the shock waves.

one shock wave to appear stationary at a time. The shifted data is now aligned in a manner that is amenable to traditional dimensionality reduction methods like POD. An example of the first mode of RPCA [34] of the shifted data is shown in Figures 3.4b and 3.4d, compared to the first mode of the POD of the original (laboratory frame) data. Here, we explore the first segment of the time series, before the bifurcation point.

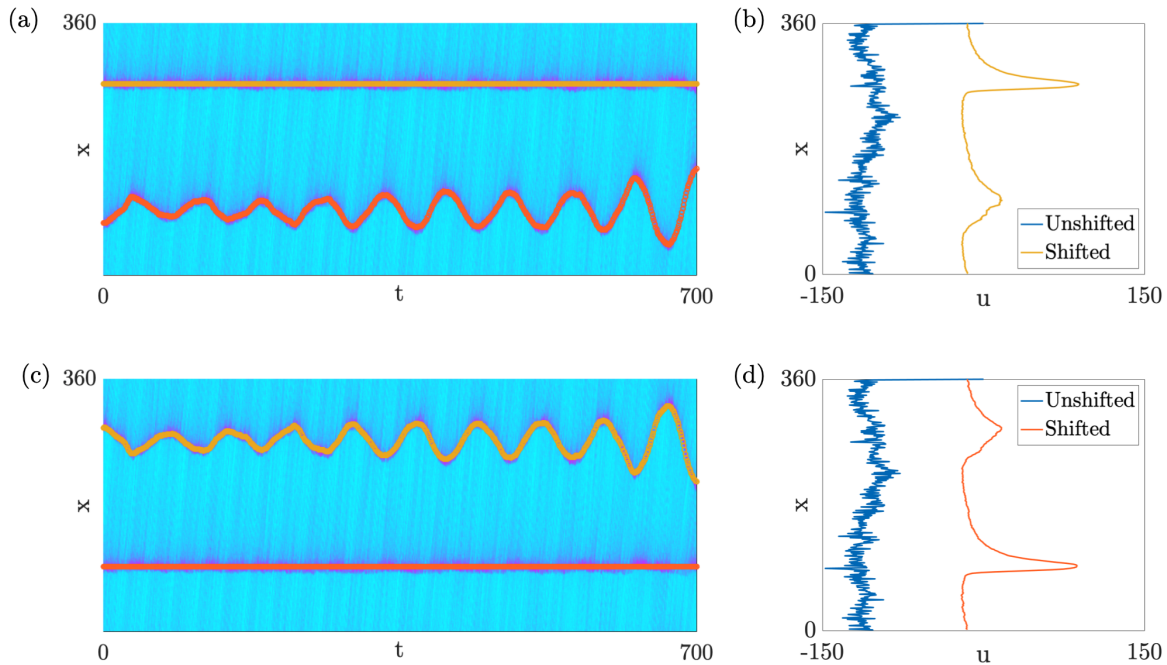


Figure 3.4: Refinement step using UnTWIST: (a) Data shifted into coordinate frame of the first (yellow) wave. (b) Corresponding first mode of a robust dimensionality reduction of the shifted data in (a) compared to the first mode of the original data. The shifted mode shows a clear shock front where the first wave is straightened, and a smaller artifact where the second wave exists, whereas the original data mode reveals no interpretable structures within the data. (c) Data shifted into the coordinate frame of the second (orange) wave. (d) First mode from dimensionality reduction of the second wave frame, showing the same shock front shape in the correct position, and artifact of the first wave. Higher modes of the unshifted dimensionality reduction are not shown, but similarly do not indicate soliton structure.

This provides an example of how UnTWIST is used on a particular data set in order to align the traveling waves to uncover low-rank representations of their wave fronts. The same

steps have been used to process various data sets. While UnTWIST can align these wave fronts into more amenable coordinate frames for dimensionality reduction of the wave field as a whole, it is also of great interest to study the linear and nonlinear interactions between wave fronts in RDEs. Using similarly aligned data and the wave speeds and locations throughout the time series, we explore models of the shock wave interactions in the following section.

3.3 *Data-Driven Models of Rotating Detonation Front Dynamics*

The ability to automate the discovery of a moving coordinate system pinned to a shock front allows for a wide range of reduced-order modeling possibilities. In what follows, we utilize data-driven modeling strategies that reduce the dynamics to simple models that characterize the observed interactions of the rotating detonation waves. Experiments show that these interactions can range from simple linear dynamics to more complicated nonlinear dynamical interactions. Our choice of methods allow us to characterize the full range of observed data.

3.3.1 *Linear Dynamics: Dynamic Mode Decomposition*

Given the linear-appearing nature of the data from Figure 3.4, a linear model for these dynamics is sought. One good candidate is the Dynamic Mode Decomposition, or DMD, as described in Section 1.2. A simple DMD model is computed by stacking one time shift as follows in Equation (3.2) and utilizing a variable projection method to compute an optimized DMD [82]. This simplified DMD formulation circumvents many of the biases introduced by standard DMD algorithms [23–26] by directly fitting an exponential solution.

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}(t_1) & \mathbf{x}(t_2) & \cdots & \mathbf{x}(t_{m-2}) \\ \mathbf{x}(t_2) & \mathbf{x}(t_3) & \cdots & \mathbf{x}(t_{m-1}) \end{bmatrix} \quad \text{and} \quad \mathbf{X}' = \begin{bmatrix} \mathbf{x}(t_2) & \mathbf{x}(t_3) & \cdots & \mathbf{x}(t_{m-1}) \\ \mathbf{x}(t_3) & \mathbf{x}(t_4) & \cdots & \mathbf{x}(t_m) \end{bmatrix} \quad (3.2)$$

Figure 3.5 shows the true distance between the shock fronts compared to discovered DMD model, using one pair of modes to approximate the dynamics, with an interpretable function to describe them, given by Equation (3.3).




Figure 3.5: True distance between the shock wave fronts compared to the linear DMD model.

$$x(t) = 18 \exp(0.001t) \cos(0.0837t) \quad (3.3)$$

3.3.2 Nonlinear Dynamics: Lotka-Volterra Model

Although there are many interacting wave dynamics that appear to be described by simple harmonic motion, i.e. linear oscillators well-captured by DMD, the RDE also produces dynamics that are strongly nonlinear in nature. Figure 3.6a presents three RDE shock fronts interacting in an oscillatory manner. When shifted into the coordinate frame of the top wave, shown in 3.6b, it is clear that the middle and lower wave exhibit sharp, periodic changes in wave speed with respect to the top wave. Such oscillations are beyond a simple linear description.

The Lotka-Volterra equations, also known as the predator-prey equations, are a coupled pair of nonlinear equations often used to model population changes in two species and are given by:

$$\frac{dy}{dt} = \alpha y - \beta yz \quad (3.4)$$

$$\frac{dz}{dt} = \delta yz - \gamma z, \quad (3.5)$$

where α , β , δ , and γ are positive real parameters controlling the growth and decay of y (prey), and growth and decay of z (predators), respectively.

A Lotka-Volterra model may be a good candidate to describe the dynamics in this system,

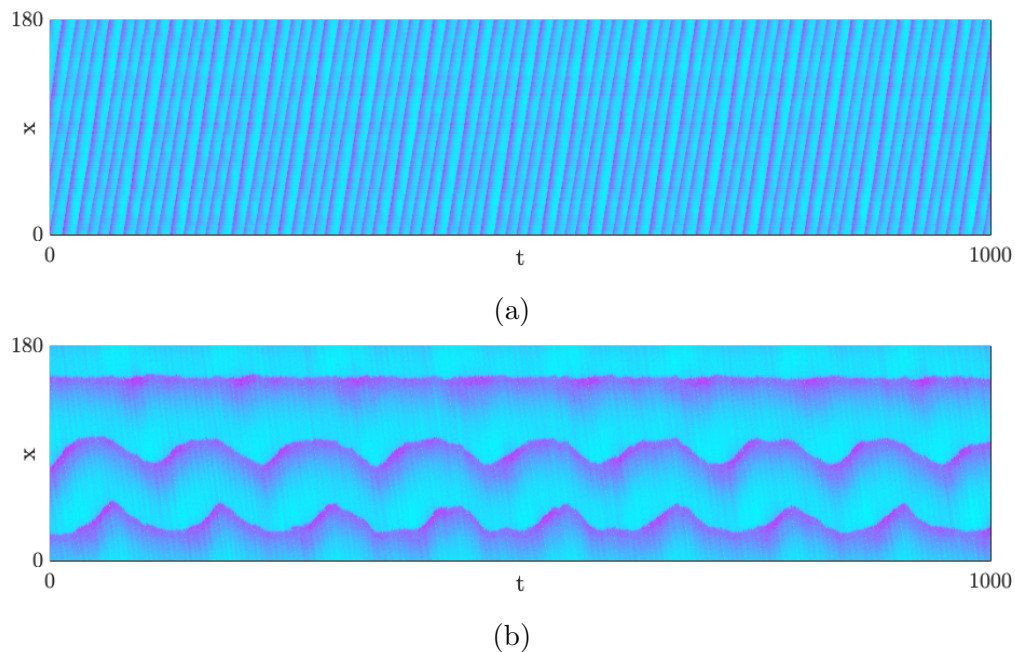


Figure 3.6: Figure (a) an example data set in the laboratory frame. (b) Same data set shifted into the coordinate frame learned by UnTWIST using linear models for wave speed for both the preprocessing and refinement steps.

not least because of its capture the periodic changes with sharp peaks. Preliminary analytical models for traveling waves within RDEs [56] indicate a similar form: by eliminating the spatial dependence in these equations, which is achieved in practice by processing with the UnTWIST algorithm, this model is phenomenologically equivalent to Lotka-Volterra. Additionally, the Lotka-Volterra model intuitively fits the nature of the data. The competition between combustion and regeneration of the propellant gives the analogous dynamics, where the prey is the reactant and the predator is combustion. One wave's acceleration combusts more fuel along the annulus, and therefore leads to the other waves' deceleration due to lack of fuel regeneration, similar to how an increase in predator population results in a decrease in prey.

We explore these nonlinear dynamics by fitting a Lotka-Volterra model to the peak loca-




Figure 3.7: Lotka-Volterra model for the two traveling waves in Figure 3.6. Yellow and orange indicate the true peak locations y and z of the two wave fronts of the data set in the shifted frame, and black shows the Lotka-Volterra model. Parameters of $\alpha = 0.07$, $\beta = 0.13$, $\delta = 0.10$ and $\gamma = 0.05$ yield a well-fitting model for the two waves which compete for resources.

tions of the middle and lower waves. We take y and z to be the locations of the two traveling waves from Figure 3.6b, using the negative of the middle wave to orient the sharp peaks in the positive direction. Parameters for the best-fit model were determined via a Nelder-Mead simplex optimization, and were found to be $\alpha = 0.07$, $\beta = 0.13$, $\delta = 0.10$ and $\gamma = 0.05$. The first 500 time steps were used as a training set. The error was computed over these time steps as the Frobenius norm of the difference between model and true $[y, z]$. The resulting model, with forward-time prediction for testing data from time 501-1000 can be seen in Figure 3.7. This model proves to be a good fit for the periodic nonlinear dynamics of wave interactions, with the frequency matching through to the end of the test data set.

The Lotka-Volterra model is only one of a potential wealth of interpretable nonlinear

models that may describe and give insight into the physics governing the RDE. An interesting avenue of future work is to automate the identification of nonlinear dynamics, for example using the sparse identification of nonlinear dynamics (SINDy) [43] algorithm. SINDy has been widely applied to identify reduced-order models for fluid systems [8, 83], including those with predator-prey dynamics [84], and is a promising candidate for obtaining low-order models of RDE dynamics.

3.4 Discussion

In this chapter, the UnTWIST method was used on observational data of a rotating detonation engine to find a coordinate system that is pinned to any desired detonation wave. The transformation gives a rotating coordinate system which is amenable to constructing ROMs that characterize the detonation front interactions. The ROMs are constructed from observational data, requiring no previous physics knowledge of the complex, multi-scale physics driving the combustion dynamics themselves. Chiefly, the most advantageous aspect of this approach is that it is able to separate the wave groups cleanly, and provide a clear representation of the waves traveling in each direction. This is particularly useful in the RDE: the traveling wave shapes and velocities give direct, though qualitative, indication of wave strength, chemical reaction rate, and relative strengths of dissipative effects.

Moreover, UnTWIST allows a diversity of model reduction techniques to be applied. Two modeling approaches were demonstrated: (i) the DMD for building the best-fit linear dynamics model and (ii) a Lotka-Volterra model for constructing nonlinear dynamical systems models for the detonation wave interactions. Both modeling paradigms are relevant as the RDE data and detonation front interactions exhibit dynamics that range from approximately linear to strongly nonlinear. Such models provide reductions that enable exploration of the complex and multi-scale dynamics of the reactive, compressible fluid dynamics of RDEs. Other modeling approaches such as deep Koopman neural networks may also provide functionally equivalent results but lack interpretability compared to UnTWIST [85].

The architecture presented here emphasizes the critical components for data-driven physics

discovery, specifically the joint discovery of coordinates and parsimonious models that represent interpretable and extrapolatory models of the physics. Given the recent emergence of RDE data, and the lack of theory characterizing detonation wave interactions, our data-driven method gives the beginnings of physical insights of the dynamics within RDEs. This proves to be a promising first step for exploring data-driven models for similar transport-dominated experimental data. Discovered models allow for engineering design and suggest control strategies that can be imposed in order to manipulate the output of RDE. They can also better inform engineering of the thermodynamic work loop [61] in order to optimize engine performance. This work shows that these engineering challenges can be approached even if a detailed physics model is not available, or if the computations are intractable.

Chapter 4

HIGHER-DIMENSIONAL SYMMETRIES

In the previous chapters, the issues posed by translational symmetry have been tackled in one spatial dimension. While these examples have expounded that the shifting approach presented in UnTWIST has promise, many applications we wish to consider are higher-dimensional and may contain more complex symmetries, such as scaling and rotation. The focus of this chapter will be to alter the UnTWIST algorithm to accommodate higher dimensions as well as a scaling symmetry. First, an example problem is presented to showcase the sparsity, interpretability, and accuracy of the higher-dimensional models. Then, a case study of high-altitude nuclear explosions, approximated with a 2D Euler equation simulation, is presented to test UnTWIST in realistic higher-dimensional ROMs applications with complex physics. Examples of one and two initial detonation waves are explored.

4.1 Extensions to UnTWIST: Arbitrary Dimensions and Symmetries

4.1.1 Modifications to the Initialization: Symmetry Detection

To detect the location of the wave structure, as well as the symmetries we seek to account for, the initialization approach for UnTWIST must be modified. In general, there are many methods that could be utilized to detect these parameters with computer vision and machine learning providing a suite of potential detection improvements. Here, a simple threshold-based detection of locations of each wave in the x - and z -dimensions and a scaling parameter will be determined. These approaches can be generalized to a third spacial dimension, or additional rotational or other symmetry if it were present. An example of the problems that will be considered here is given in Figure 4.1.

The input data is given as $\mathbf{U} \in \mathbb{R}^{m \times n \times q}$, where $\mathbf{x} \in i = 1 : m$, $\mathbf{z} \in j = 1 : n$, and




Figure 4.1: Time series showing the wave height of an example Gaussian traveling in two dimensions and increasing in size, exemplifying the translational and scaling symmetries.

$\mathbf{t} \in \ell = 1 : q$. The first step in the wave detection is setting a simple threshold. In the examples that follow, a threshold τ is determined to be a certain percentile of wave height over the full data set.

$$\mathbf{U}_\tau = \mathbf{U} > \tau \quad (4.1)$$

For example, the 98th percentile of values in \mathbf{U} yields τ , and all points above this τ will be considered within a quantity of interest. Time slices of those points within the \mathbf{U}_τ map can be seen as disks in Figure 4.2.

Indices of the points that meet the threshold are then used to find the area of the quantity of interest, as well as its (x, z, t) location. The centers of the waves in the two dimensions, \mathbf{x}_c and \mathbf{z}_c , shown as the pink lines projected on to the respective planes in 4.2, are given in cm by

$$\begin{aligned} \mathbf{x}_c(t_\ell) &= \frac{\sum_{i,j} \mathbf{U}_\tau(x_i, z_j, t_\ell) \mathbf{x}}{\sum_{i,j} \mathbf{U}_\tau(x_i, z_j, t_\ell)} \\ \mathbf{z}_c(t_\ell) &= \frac{\sum_{i,j} \mathbf{U}_\tau(x_i, z_j, t_\ell) \mathbf{z}}{\sum_{i,j} \mathbf{U}_\tau(x_i, z_j, t_\ell)} \end{aligned} \quad (4.2)$$

The shifts for these quantities are then computed as the distance from the center of the wave of interest $\mathbf{x}_c, \mathbf{z}_c$ to the center of the data set. For an input data set of size $\mathbf{U}(t_\ell) \in \mathbb{R}^{m \times n}$,

the following two quantities, in pixels, are computed

$$\begin{aligned}\text{shift}_x(t_\ell) &= \lfloor (i(t_\ell) - m/2) \rfloor \\ \text{shift}_z(t_\ell) &= \lfloor (j(t_\ell) - n/2) \rfloor\end{aligned}\tag{4.3}$$

where $\lfloor \cdot \rfloor$ indicates rounding to the nearest integer, and (i, j) represent the indices within (\mathbf{x}, \mathbf{z}) that correspond to centers $(\mathbf{x}_c, \mathbf{z}_c)$. With these shift values, all waves will be centered within the frame at each time step, and will expand radially from this center point.

The area in cm, shown by the cone encompassing the disks in Figure 4.2, is computed by

$$\mathbf{a}(t_\ell) = dx \cdot dz \cdot \sum_{i,j} \mathbf{U}_\tau(x_i, z_j, t_\ell).\tag{4.4}$$

The overall growth of the wave quantities must be counteracted in order to view the waves in stationary coordinate frames. To achieve this, an overall scaling parameter is determined based on the radius of the area of interest as approximated by a circle. The radius of each area of interest, shown by the gray area the $z - t$ and $x - t$ planes in Figure 4.2 is easily computed as

$$\mathbf{r}(t_\ell) = \sqrt{\mathbf{a}(t_\ell)/\pi}.\tag{4.5}$$

To retain as much detail in the data as possible, the scaling parameter is determined such that the largest original radius of the quantity of interest is scaled to 1, and all smaller radii are scaled up proportionally. The scaling parameter for each time step is then computed as

$$\mathbf{s}(t_\ell) = \frac{\max(\mathbf{r})}{\mathbf{r}(t_\ell)}\tag{4.6}$$

and is used alongside a up-sampling function, here using `imresize` in MATLAB. Other up-sampling or resizing approaches can be used; the field of superresolution is abundant with approaches for this step. Three quantities, the shifts $\text{shift}_x(t_\ell)$ and $\text{shift}_z(t_\ell)$, as well as the scaling parameter $\mathbf{s}(t_\ell)$ are used as inputs to UnTWIST. The algorithm is modified to find sparse, interpretable models for these quantities.




Figure 4.2: Schematic of the shifting and scaling captured in the initialization step, with the same data shown in Figure 4.1. Slices of the wave above the threshold, τ , are plotted in time (disks). The center locations of the wave are shown piercing through the data slices in pink. The center locations are projected in pink to show the \mathbf{x}_c and \mathbf{z}_c in time. Radius of the wave in time \mathbf{r} is also shown projected in gray, with the maximum radius in this example at the end of the time series.

4.1.2 Modifications to the UnTWIST Algorithm

In the original formulation of the UnTWIST algorithm, only one symmetry is accounted for— a shift in one spatial dimension. Moving into systems with higher dimensions and more complex symmetries necessitates an inclusion of more models while retaining the capability of separating multiple waves. In the previous presentation of UnTWIST, time was accommodated in the rows of matrix elements. Columns corresponded to individual translational elements. For example, rows of \mathbf{W} corresponding to each wave peak point in time, and columns correspond to each wave. While higher-dimensional data can be considered in this formulation, all models and shifts are restricted to the one spatial dimension under consideration. Ideally, including models for shifts in more than one dimension simultaneously is desired. To accommodate this, we add a dimension in the UnTWIST structure. Reiterating

the original cost function in equation (2.3), given as

$$\min_{\mathbf{C}, \mathbf{B}, \mathbf{W} \in \Omega} \frac{1}{2} \mathbf{W} \odot \|\mathbf{X} - \mathbf{T}\mathbf{C}\|_2^2 + \lambda R(\mathbf{B}) + \frac{1}{2\zeta} \|\mathbf{C} - \mathbf{B}\|_2^2, \quad (4.7)$$

we now augment the structure of each term as follows, retaining the original parameters, where N is the number wave peak points, k is the number of potential models in the library, n is the number of waves, and a newly introduced d is the total number of symmetries to consider:

Term	1D	d-D
\mathbf{W}	$N \times n$	$N \times n$
\mathbf{X}	$N \times n$	$N \times n \times d$
\mathbf{T}	$N \times k$	$N \times k$
\mathbf{C}, \mathbf{B}	$k \times n$	$k \times n \times d$

It is important to note here that \mathbf{X} and \mathbf{C} , corresponding to the \mathbf{x} points and models in the fitting, respectively, have changed sizes to accommodate a higher number of symmetries. However, \mathbf{T} remains the same dimension because the candidate models apply equally across each symmetry. In addition, \mathbf{W} remains identical to the one-dimensional case. This is because \mathbf{W} separates the wave groups and by construction must remain the same across each symmetry. When considering an x and z shift for each wave group, if the x shift applies for a given wave, the corresponding z shift and scaling should also apply for that wave, rather than being reassigned to another wave group. The new dimensions of the data are also incorporated to the spectral clustering used in the initialization of \mathbf{W} , making it more robust. To initialize the \mathbf{C} matrix, initializing the model functions, the same sequentially thresholded least squares approach is used, this time being performed on each of the d symmetries separately.

To account for these new dimensions in the optimization, we can represent the new formulation as

$$\min_{\mathbf{C}, \mathbf{B}, \mathbf{W} \in \Omega} \sum_{j=1}^d \frac{1}{2} \mathbf{W} \odot \|\mathbf{X}_j - \mathbf{T}\mathbf{C}_j\|_2^2 + \lambda R(\mathbf{B}_j) + \frac{1}{2\zeta} \|\mathbf{C}_j - \mathbf{B}_j\|_2^2, \quad (4.8)$$

where sums in j correspond to the third dimension of \mathbf{X} , \mathbf{C} , and \mathbf{B} . In this formulation, the UnTWIST algorithm applies in any number of spatial dimensions. One iteration of the algorithm will fit, for example, models in x , y , and z separately, with (x_i, y_i, z_i, t_i) all corresponding to the wave or wave group.

4.2 Example Problem: Growing Two-Dimensional Gaussian

To verify that the UnTWIST algorithm can indeed discover true, interpretable underlying physics, an example problem with prescribed shifting and scaling is created.

Example Data In this example, a two-dimensional Gaussian solution is constructed as

$$u(x, z, t) = \exp\left(\frac{-(x - c_x(t))^2 - (z - c_z(t))^2}{2t^2}\right) \quad (4.9)$$

with speeds in the x and z directions given by

$$\begin{aligned} c_x(t) &= 4t^2 - 2 \\ c_z(t) &= 6t^3 - 4 \end{aligned} \quad (4.10)$$

respectively, in the domain $x = [-5 : 0.1 : 5]$, $z = [-6 : 0.1 : 6]$, and $t = [0.1 : 0.1 : 1]$. This example is shown in Figure 4.1. These prescribed speeds, written in sparse, identifiable terms will be used to compare to the models resulting from UnTWIST.

In order to find the centers of the Gaussian, a threshold τ is assigned to be the 98th percentile of the total data. Here, $\tau = 0.32$. Solving for the radius of the intersection between the Gaussian and this threshold, we find

$$r = t\sqrt{-2\ln(\tau)} \approx 1.5t. \quad (4.11)$$

This is the true physical radius that bounds the data we wish to consider. It will be used to compare to the outcome of the UnTWIST model.

UnTWIST Results Using the initialization detailed in Section 4.1.1, shifts shift_x and shift_z , and scaling parameter based on the radius \mathbf{r} were determined for this data set. The UnTWIST algorithm was initialized with parameters $\lambda = 1$ and $\zeta = 1e - 2$, and a candidate speed library of

$$\mathbf{T}(\mathbf{t}) = \begin{bmatrix} \frac{1}{\mathbf{t}} & \mathbf{t}^4 & \mathbf{t}^3 & \mathbf{t}^2 & \sqrt{|\mathbf{t}|} & \mathbf{t} & \mathbf{1} \end{bmatrix}. \quad (4.12)$$

After 2 iterations, UnTWIST yields the following models for the invariances depicted in Figure 4.3.

$$\begin{aligned} \text{shift}_x(t) &= 4.0008t^2 - 2.0004 \\ \text{shift}_z(t) &= 6.0007t^3 - 4.0003 \\ s(t) &= 1.5095t. \end{aligned} \quad (4.13)$$

Compared to the known speed models given by (4.10) and the true radial expansion given by (4.11), the resulting UnTWIST models perform well qualitatively and quantitatively, with a ℓ_2 -norm fit error of $6.7e-04$. Of unique importance is that the models are sparse, match the correct form of the model terms, and are accurate in coefficient values, with a ℓ_2 -norm coefficient error of $1.3e-03$.

Now that the models have been determined, the shifting and scaling operators can be applied. The results of these manipulations can be seen in Figure 4.3. It is clear that applying the shifting and scaling operators makes the wave appear as a stationary quantity, which is more amenable to low-rank representation. Once the data has been repositioned and scaled so as to make the wave stationary, the singular value spectrum of each level of manipulation will disclose the role that these shifts play. The POD/SVD was computed for each data set and the decay of the singular values can be seen in Figure 4.3. The shifted & scaled data's singular value decay shows a large share of the system's energy contained by the first mode, as well as the rapid drop off of the energy contained in the remaining modes. Relative to the original and shifted data, this indicates that very low rank representations may be suitable for this data. Reconstructions were generated for each

data set. For the original data set, the first mode of the POD/SVD was retained. For the shifted data set, the first mode was retained, and the shift operation was inverted on this low-rank reconstruction. Similarly, for the shifted & scaled data set, the scaling was also inverted on the low-rank reconstruction, with zeros padding the matrix to account for down sampling. The comparison of these 1-mode reconstructions can be seen in Figure 4.4. It is clear from this comparison that for equivalently low-rank approximations, the shifting and scaling operators yield more physically relevant and interpretable reconstructions. Early in the original data's approximation, wave forms are not retained, with the spherical nature of the true wave being distorted. Conversely, the spherical wave shape, complete with traveling and scaling phenomena, is accurately captured when shifting and scaling are accounted for.

Overall, the models discovered by UnTWIST prove to be physically relevant, accurate to the original system, and useful to improve the dimensionality reduction in both efficiency and interpretability. These promising results on this example problem indicate that the UnTWIST algorithm can be expected to perform similarly on systems with two-dimensional shifting and scaling phenomena.

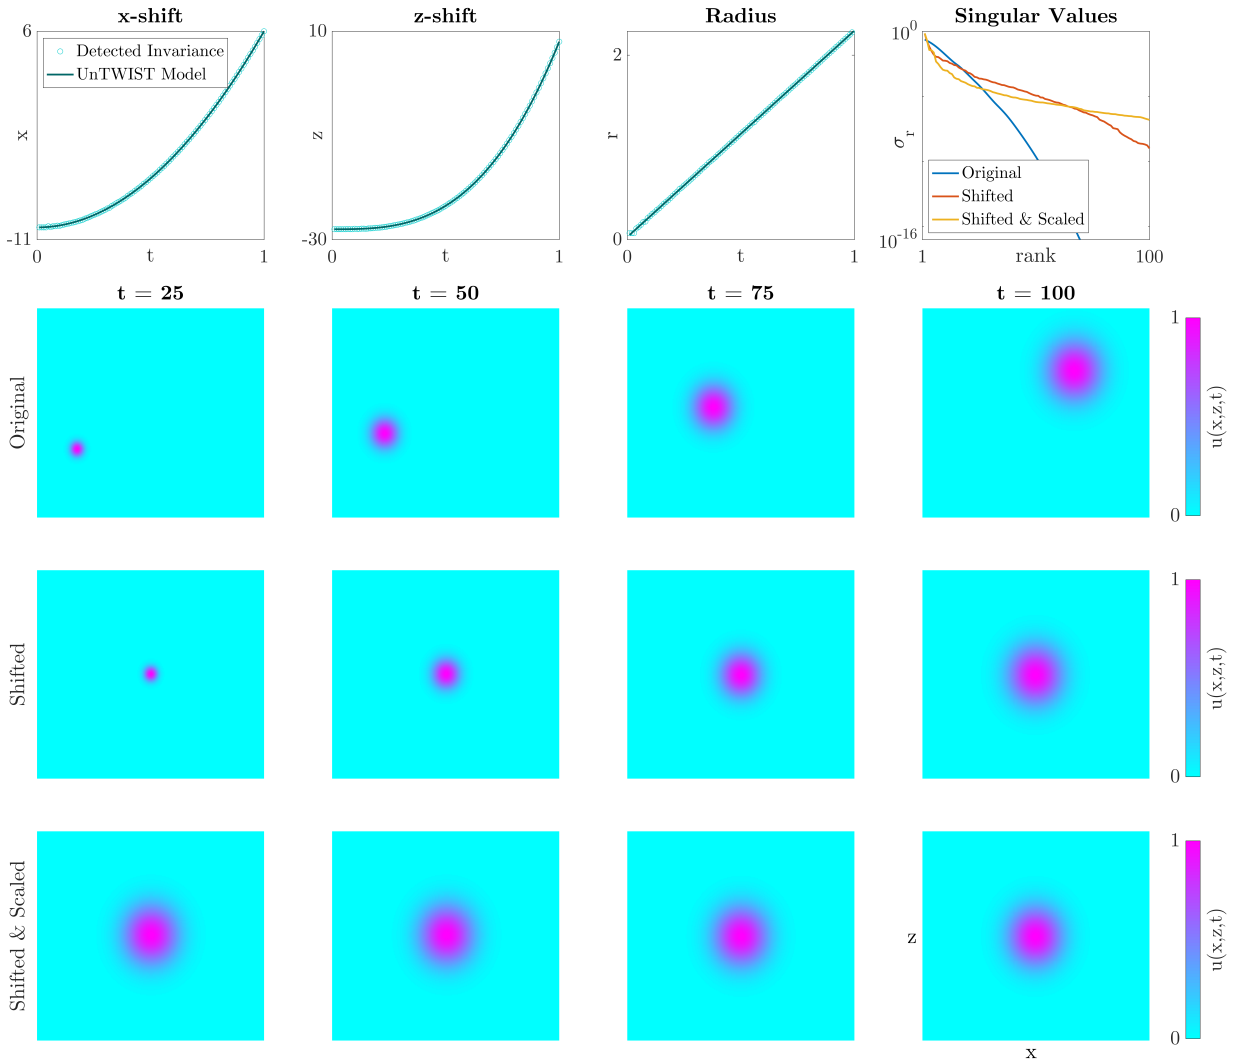


Figure 4.3: (top) Models discovered by the UnTWIST algorithm plotted alongside the detected invariances and singular value spectra of each level of alignment. (second) Original data, (third) shifted data aligning the quantity of interest in the center of the frame. (bottom) Effect of scaling based on the discovered radius model, yielding a perfectly stationary Gaussian.




Figure 4.4: Comparison of 1-mode reconstructions for each of the 3 stages of invariance reversal— unadulterated original data, shifted data, and shifted & scaled data. It is clear that 1 mode is insufficient for reconstructing the wave shapes for the original data, yields some information for the shifted data without capturing the scaling, and clearly describes the shifted & scaled data.

4.3 High-Altitude Nuclear Explosions

High-altitude nuclear explosions (HANEs) are nuclear explosions that detonate within Earth’s atmosphere above an altitude of about 30 kilometers. Historically, several HANE tests have been conducted between 1958 and 1962 to determine the effect of such explosions. A HANE differs from a typical nuclear explosion on Earth’s surface in that it creates a strong electromagnetic pulse that spreads high-energy particles over a much wider radius, degrading radio frequency communications and damaging or destroying satellites and power infrastructure [86, 87]. The detonation cloud emitted by such an explosion begins as a spherical cloud, contrasting the well-known mushroom cloud, until distorting under Earth’s magnetic field.

There is some information known about the wide-reaching multiphysics effects of HANEs from these historical experimental observations [88]. However, because of the obvious drawbacks of experimental testing and the immense threat posed by such weapons, it is critical to model and predict the behavior and effects of HANEs via computation. Similar to the challenges of time scale posed by the rotating detonation engine, HANE simulations must be resolved on the sub-nanosecond scale to study the initial explosion, but must cover hours or days to capture the full ionization and fallout behavior. Further physical considerations such as magnetohydrodynamics and buoyancy and chemistry of the atmosphere must also be taken into consideration. Dimensionality reduction and reduced-order modeling are therefore critical tools to studying HANEs in a wide parameter space.

4.4 Three-Dimensional Euler Equations

While detailed HANE simulations are complex and include the modeling of many types of physics, an idealized test problem can capture some of the dominant aspects of high-altitude phenomena. For example, solving the Euler equations on a domain with a spherical perturbation of high temperature in a vertically stratified background atmosphere. This idealized problem setup can provide enough complexity to test ROM algorithms while capturing important phenomenology such as the generation of radially expanding waves and gravitational

buoyancy effects.

The three-dimensional system of equations is given as

$$\frac{\partial}{\partial t} \begin{bmatrix} \rho \\ \rho \mathbf{v} \\ E \end{bmatrix} + \nabla \cdot \begin{bmatrix} \rho \mathbf{v} \\ \rho \mathbf{v} \mathbf{v} + p \mathbf{I} \\ (E + p) \mathbf{v} \end{bmatrix} = \begin{bmatrix} 0 \\ -\rho \mathbf{g} \\ 0 \end{bmatrix} \quad (4.14)$$

where ρ is the mass density, \mathbf{v} is the velocity vector, p is the pressure, \mathbf{I} is the identity matrix, E is the total energy defined by $E = \frac{p}{\gamma-1} + \frac{1}{2}\rho|\mathbf{v}|^2$, γ is the ratio of specific heats, and \mathbf{g} is the vector of acceleration due to gravity. This system was simulated for time 0 to 160 seconds using the PyClaw software package [89–91], using the Euler with gravity Riemann solver. A $CFL = 0.5$ is used with data outputs every 1 second. The simulations use a spherical coordinate system domain

$$\theta \in [\pi/4 - 0.15, \pi/4 + 0.15] \text{ rad}$$

$$\phi \in [0, 0.30] \text{ rad}$$

$$r \in [80e5, 1559e5] \text{ cm}$$

which corresponds to about a 1000 km by 1000 km latitude and longitude extent and 80 to 1559 km altitude range. A background mass density and pressure are initialized from tables to yield a vertically stratified background atmosphere. The pressure is slightly modified to maintain hydrostatic equilibrium using the f-wave [92] method.

Extrapolation boundary conditions are used in the $\pm\theta$, $\pm\phi$ and the top ($+r$) boundaries. At the bottom ($-r$) boundary the initial background atmosphere conditions are held fixed

$$\begin{aligned} \rho_{k=0} &= \rho_0 \\ \rho \mathbf{v}_{k=0} &= 0 \\ E_{k=0} &= \frac{p_0}{\gamma - 1} \end{aligned}$$

where ρ_0 and p_0 are the initial background mass density and pressure and k is the index for the radial r dimension.




Figure 4.5: Velocity magnitude of each type of data set: (top) one initial energy perturbation of varying magnitude, (bottom) two initial energy perturbations of a given magnitude at varying distances, each at four time steps as viewed from the central ϕ plane of the three-dimensional view.

Two main data sets are simulated using these equations. In the first set of simulations, a single spherical temperature perturbation was constructed with various initial energies, 1.0, 1.6 and 2.0 eV, and propagated through the time domain. In the second set of simulations two spherical temperature perturbations of 1.0 eV and 1.6 eV initial energy were simulated with various distances separating the two. The perturbations are initialized at an altitude of 200 km with a peak temperature at the center and falling off linearly until a radius of 100 km.

4.5 One Wave in Two Dimensions

Data A 100 km radius hot sphere with no initial velocity is initialized in the center of the domain on the $\theta - \phi$ plane and at 200 km altitude. Three test cases with initial hot sphere energies of 1.0 eV, 1.6 eV and 2.0 eV were simulated. The resulting velocity magnitudes for




Figure 4.6: Velocity magnitude of 1.0, 1.6 and 2.0 eV initial energy perturbation simulations at four example time slices as viewed from a central plane of the three-dimensional view.

three initial detonation energies can be seen in Figure 4.6, depicting the center ϕ slice of the front at example times. Throughout the following examples, though three dimensions were computed, two-dimensional planar slices were considered.

UnTWIST Results This data, similar to the 2D Gaussian example, depicts three main symmetries, the first being the translational symmetry of the sphere rising due to buoyancy. The same steps shown in the previous example from Section 4.2 are used to process the Euler equation data, yielding θ and r centers of the waves and radii over time. First, the UnTWIST algorithm was used to discover models for the θ and r shifts and a scaling parameter for each data set. In this application, all values, including t , shift_θ , shift_r , and $\mathbf{s}(t)$

were first normalized. In doing so, each model could be fit agnostic to the scale relative to the other quantities, i.e. $\mathbf{s}(t)$ fitting on the order of 10^1 and shift_r fitting on the order of 10^2 . After the model discovery step, the normalizations were accounted for. The identified wave centers, scaling parameter, as well as the discovered UnTWIST models for each quantity can be seen compared in Figure 4.7 for each initial energy value. The input hyperparameters, as well as the total iterations, model fit errors, and resulting model equations are given in Appendix A, Table A.1.

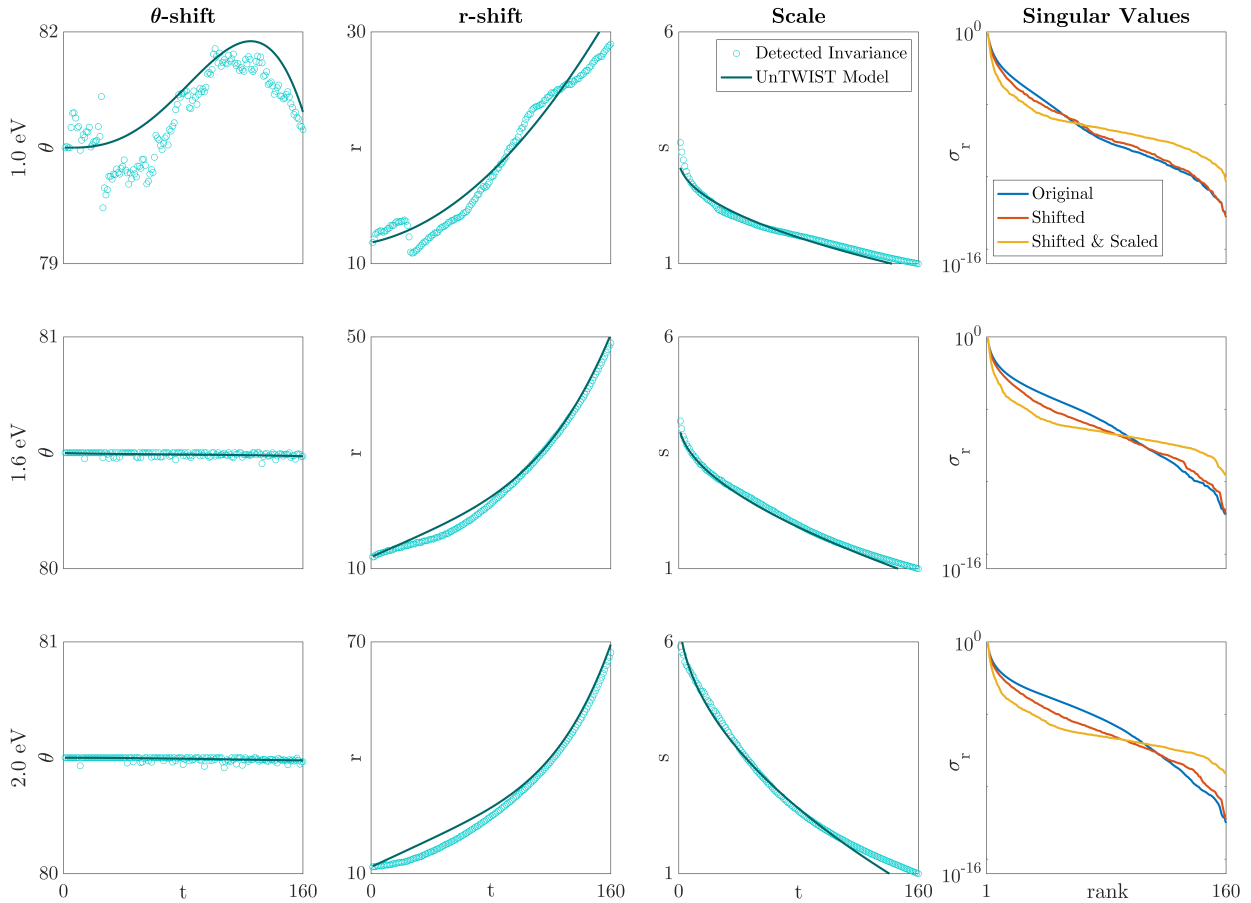


Figure 4.7: UnTWIST models for shift_θ , shift_r and $\mathbf{s}(t)$ each initial perturbation energy, 1.0, 1.6, and 2.0 eV, compared to the data collected in the initialization step. Singular value spectra for each data set is also shown before and after accommodating for each invariance.

Figure 4.7 depicts the models for each data set for the three invariance we wish to account for– the θ -shift, r -shift, and scaling. While the data taken from initialization is somewhat noisy, especially as seen in the θ - and r -shifts for the 1.0 eV detonation, it is clear that the UnTWIST models fit the data well while smoothing inconsistencies in the wave detection. This will be useful when accounting for these invariances in the data, yielding smooth transitions between time steps rather than jumpy shifts.

These three invariances were accounted for in each data set, and a standard POD/SVD was computed for each data set. The resulting singular value decompositions can also be seen in Figure 4.7. These singular value decompositions show that the shifted & scaled data is most amenable to POD. The energy contained in the first mode is higher in the shifted & scaled data sets than the shifted or original data for each data set. The significance of the first few modes is most prominent in the 2.0 eV data set. In this case, the higher-energy wave front travels faster than the other initial energies. The improved initialization and model fit of UnTWIST on this data set indicates that performance of the UnTWIST method increases with wave speed. In addition, the singular value spectra indicate that very few modes can be used to reconstruct all three data sets.

For each data set, 1-POD-mode reconstructions were created, and invariance modifications were reversed as in the previous example. Figure 4.8 shows these low-rank reconstructions. The results are effective and interpretable, and show that UnTWIST has aided in the dimensionality reduction. For all data sets, the low-rank reconstruction of the original data set is unable to capture the invariances– neither the translating nor the scaling is captured. For the 1.0 and 1.6 eV data sets, the UnTWIST-modified reconstruction is extremely close to the original data. Some of the intensity of the shock front is not captured, but the overall dynamics, as well as the shifting and scaling, are well-captured. Similarly, the 2.0 eV data set captures these invariances, but presents difficulties as the wave front exceeds the boundaries. In this case, the periodic shift adds the wave front to the bottom of the r -domain. While these results are promising, one capability of UnTWIST has not yet been explored in this higher-dimensional example– accommodating and separating multiple waves.

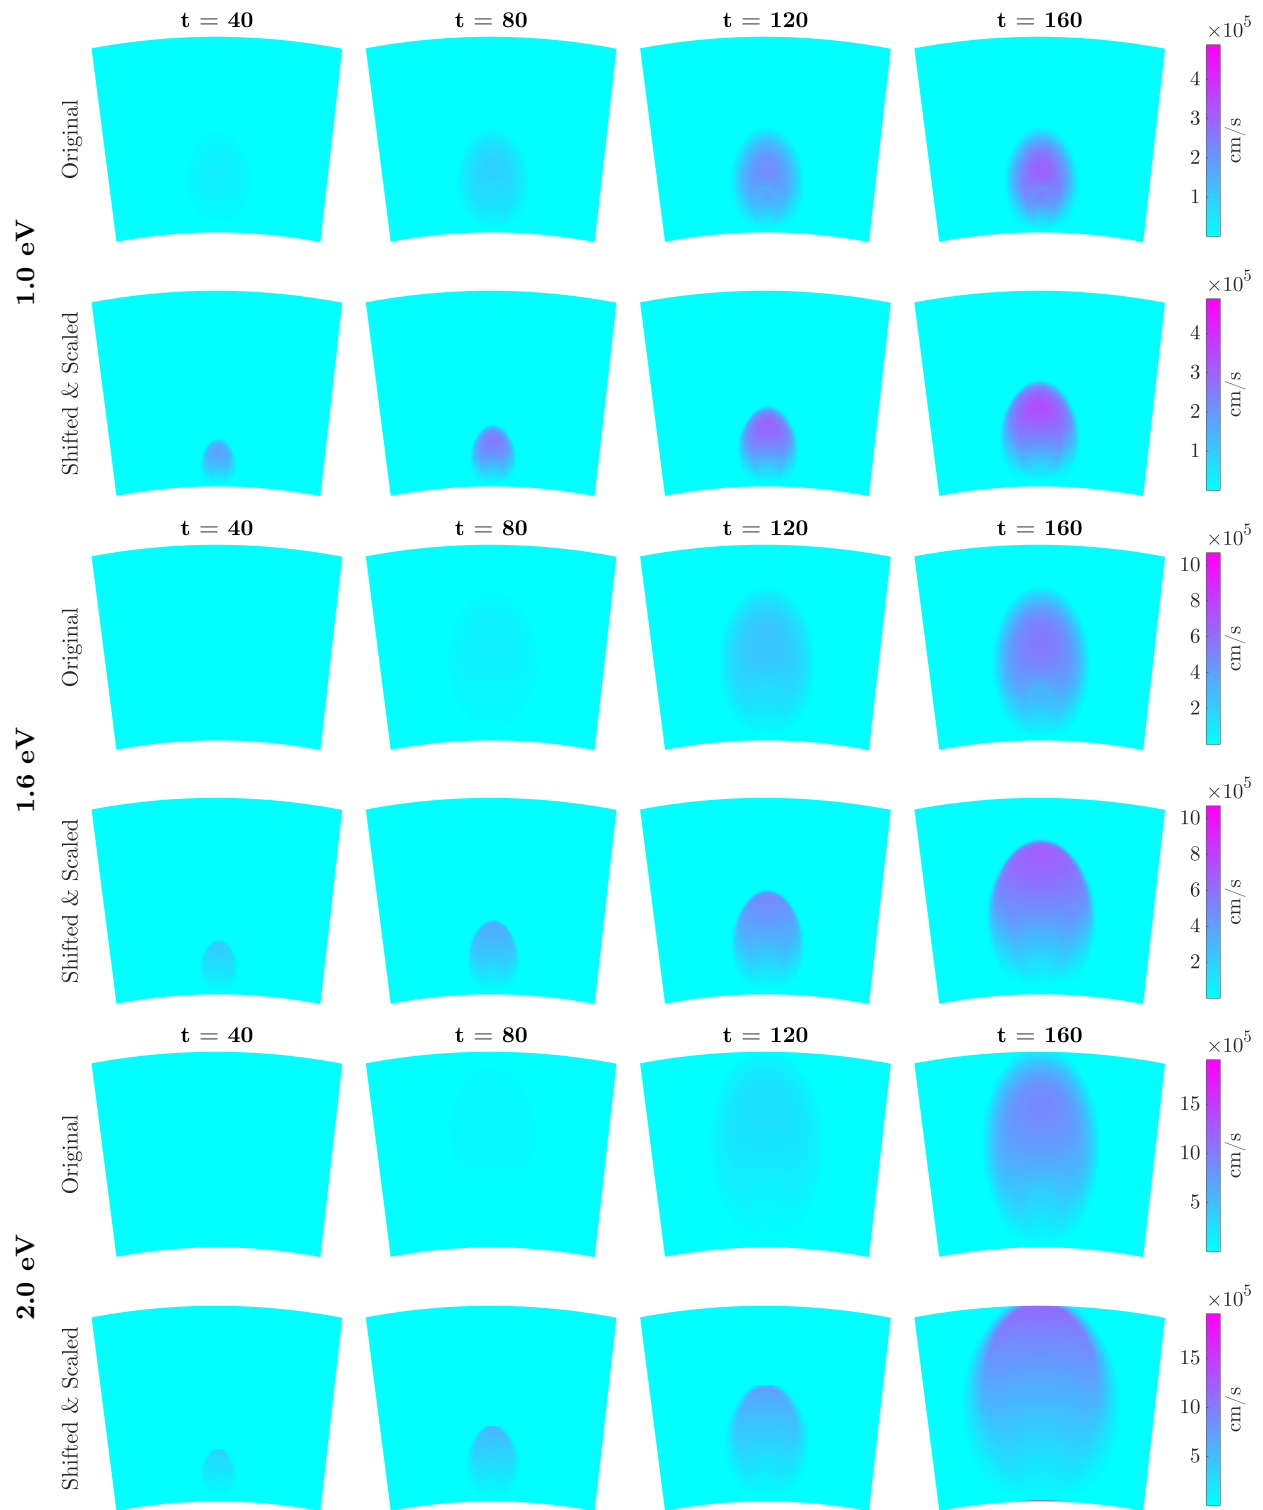


Figure 4.8: 1-mode reconstructions of original vs UnTWIST-modified single-wave data.

4.6 Two Waves in Two Dimensions

In the previous section, the use of UnTWIST on 2D single-detonation phenomenon was explored. In this section, the same Euler equations will be used to determine the limits of UnTWIST in the presence of more than one wave quantity in 2D.

Here, under the same domain, one detonation of 1.0 eV (left) and one of 1.6 eV (right) each with a 100 km radius are initialized. 6 total data sets are generated, with the spacing between the two detonations varying between 0.5 to 1.5 initial diameters apart. Time slices of this initial data can be seen in Figure 4.9. As the initial separation distance between the detonations increases, the two quantities remain separate for a longer time before interfering. This interference will pose a challenge to UnTWIST, as the distinction between the wave becomes unclear, so does the necessary symmetry accommodation.

Similar to the previous one-wave example, the same steps were performed for finding values for the wave centers and scalings. Here, the wave centers were naively computed by separating the domain into a left and right half. The same shift_θ , shift_r and $\mathbf{s}(t)$ parameters were computed for each data set. In this case, the separation function of UnTWIST is used, yielding models for wave 1 and wave 2 separately. The resulting models can be seen in Figure 4.10. The details of the form of the models as well as the hyperparameters used in the optimization can be found in Appendix A, Tables A.2 and A.3.

It is clear that from the increase in spacing come more clear models for the symmetries. Ideally, the θ -shift for each data set would be captured as a constant value, since there is no underlying translation in the θ direction. However, with close detonation wave fronts, such as the 0.5, 0.7, and 0.9 diameter spacings, the θ -shift becomes muddled and veers toward the opposite wave. The intermediately spaced θ -shifts also show that while the initial model may be accurate, when the waves interact, the center of one individual wave is difficult to capture accurately. The 1.1 diameter spacing data set shows the constant θ -shift captured well until a certain point in time, where the models become skewed. Figure 4.10 shows an improved picture of the θ -shift with 1.3 and 1.5 diameter separation, with values for each

wave fairly constant throughout time.

The r -shifts are fairly accurately captured for both waves, though there is some noise in the two most closely-spaced examples. The scaling is well-captured, with all models matching closely to the data. This is in part due to the nature of the scaling computation, where each wave is assumed to be circular, yielding smooth radius and therefore scaling models. Nevertheless, the UnTWIST models capture well-fitting scaling parameters.

In order to see the effect of these various symmetry modifications, the shifting and scaling parameters were accounted for as in the previous examples. Whereas there are two desired wave frames, each stationary with respect to one wave, each wave frame was filtered using the predetermined threshold. Values outside the area of interest for that frame were set to zero. The entire data set was shifted and scaled to make each wave stationary separately, and the singular value decomposition was taken for each wave frame. The results, compared to the singular value spectra of each original data set can be seen in Figure 4.11.

The compared singular value spectra in Figure 4.11 indicate that dimensionality reduction may be slightly improved. For many of the individual wave frames, the fraction of information captured by the first mode is slightly improved from the original data set. For some, the first mode performs noticeably better than the original data, but a closer look at the reconstructions is needed. To reconstruct the data using these low-rank modes, each wave frame was considered at a time. 1 mode was retained for each wave frame, and each wave frame was inversely shifted and scaled. The sum of the two low-rank reconstructions are used to represent the original data. The resulting 2-mode reconstructions for each data set can be seen in time slices in Figures 4.12 and 4.13, comparing the original data to the UnTWIST-modified data.

From these reconstructions, UnTWIST shows a benefit over traditional dimensionality reduction approaches in cases where the waves are separable. For the farther-spaced data sets, such as the 1.5 diameter spacing, UnTWIST provides a clear advantage. The wave structures are accurately represented, retaining their separate scales and shifts. Compared to the reconstruction of the original data set, the fronts are clearly captured, and contain a larger

share of the original energy. For more closely spaced wave fronts, the advantage of UnTWIST is less clear. Whereas the traditional POD approach does not capture the scaling accurately, and leaves out a large share of the energy in the center of the wave, the reconstructions using UnTWIST are also not physically realistic. In general, closely-spaced wave data sets may not be amenable to the current formulation due to the naive approach to separating the wave structures. Reconstructions exacerbate the hard boundaries used to detect invariance values. Alternatively, methods such as RPCA or other background subtraction approach [93] may aid in separating these wave structures in their stationary wave frames. An improvement in the separation of the wave structures would allow UnTWIST to more viably represent these waves in low-rank.




Figure 4.9: Euler simulation data with two initial perturbations 0.5-1.5 diameters apart.

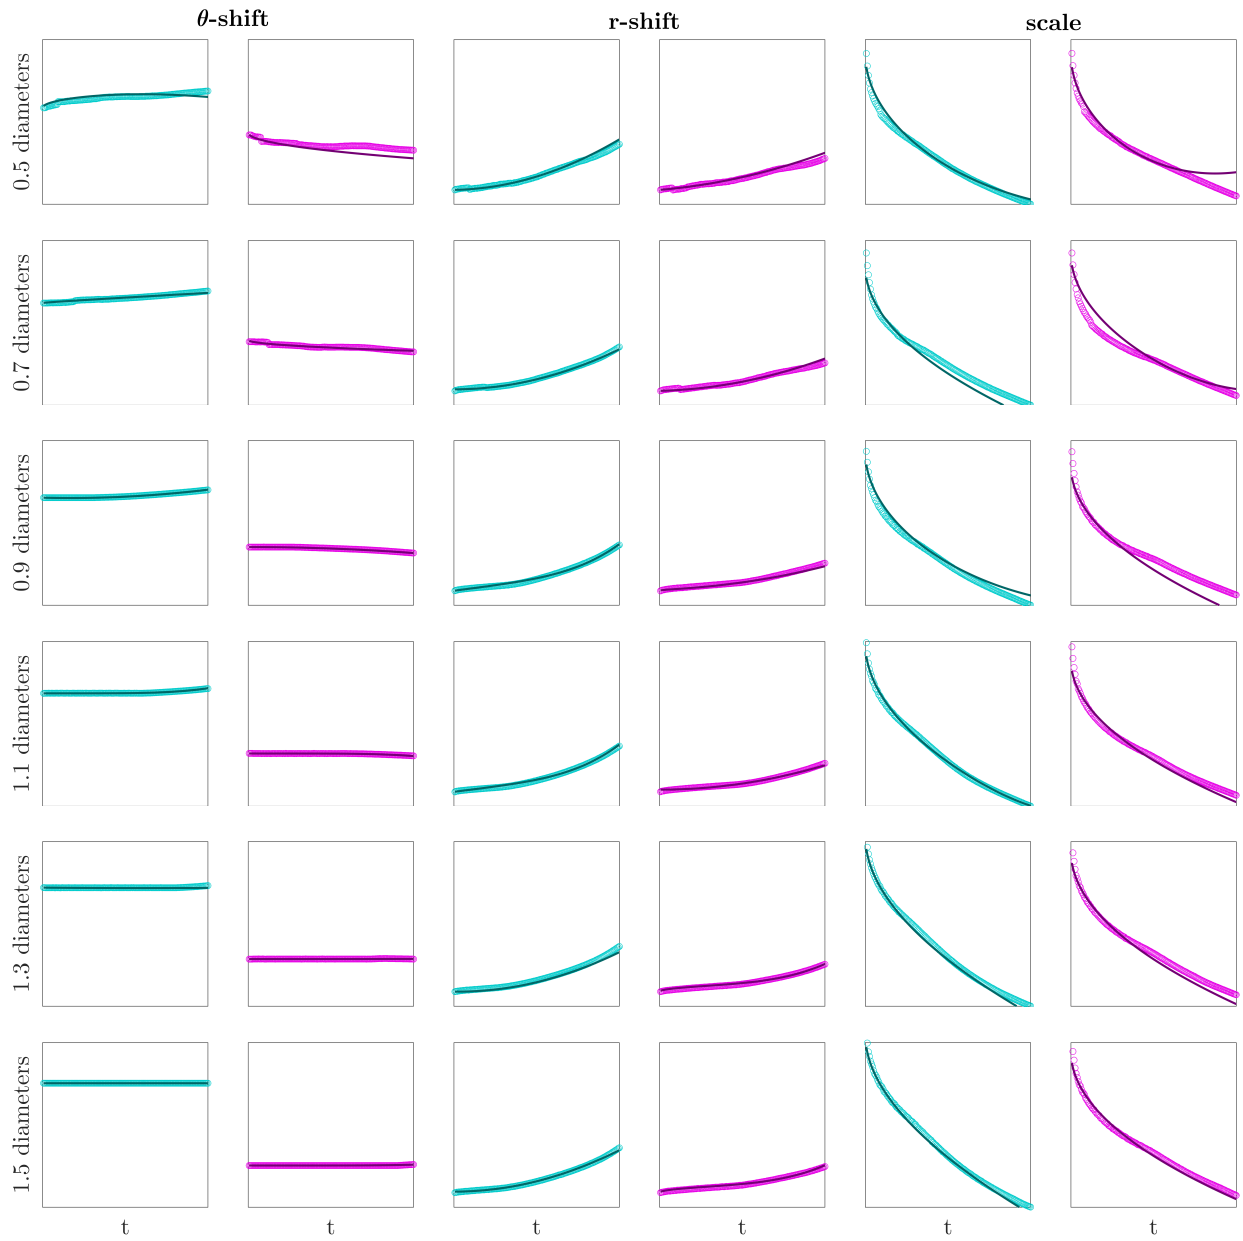


Figure 4.10: UnTWIST models for shift_θ , shift_r and $\mathbf{s}(t)$ for wave 1 (teal) and wave 2 (pink) for the each two-wave detonation data set. Limits on θ and r are the same as the domain, and scale is plotted in $[1,4]$.

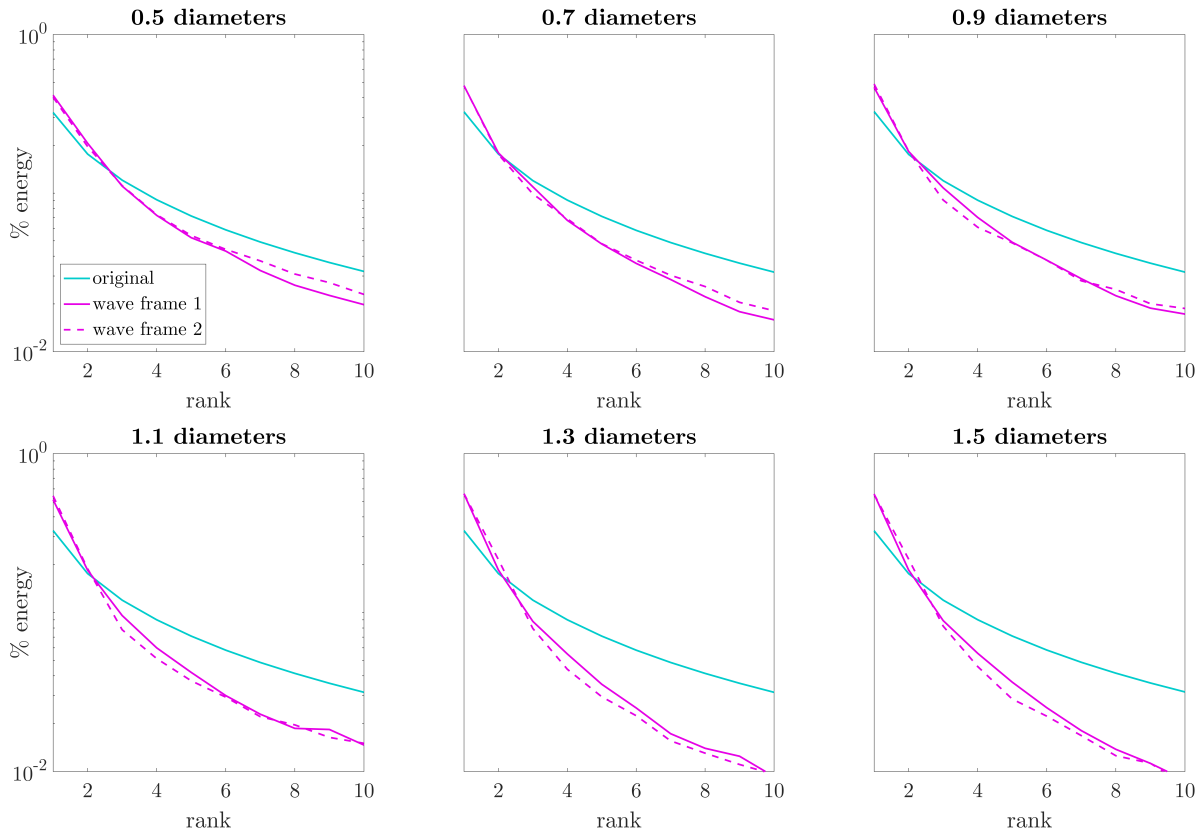


Figure 4.11: Singular value decompositions of each two-wave data set, 0.5-1.5 diameter detonation spacings. Blue lines shows the singular value decay of the original data set, where pink solid and dashed lines show the singular value decays after shifting & scaling for each wave frame.

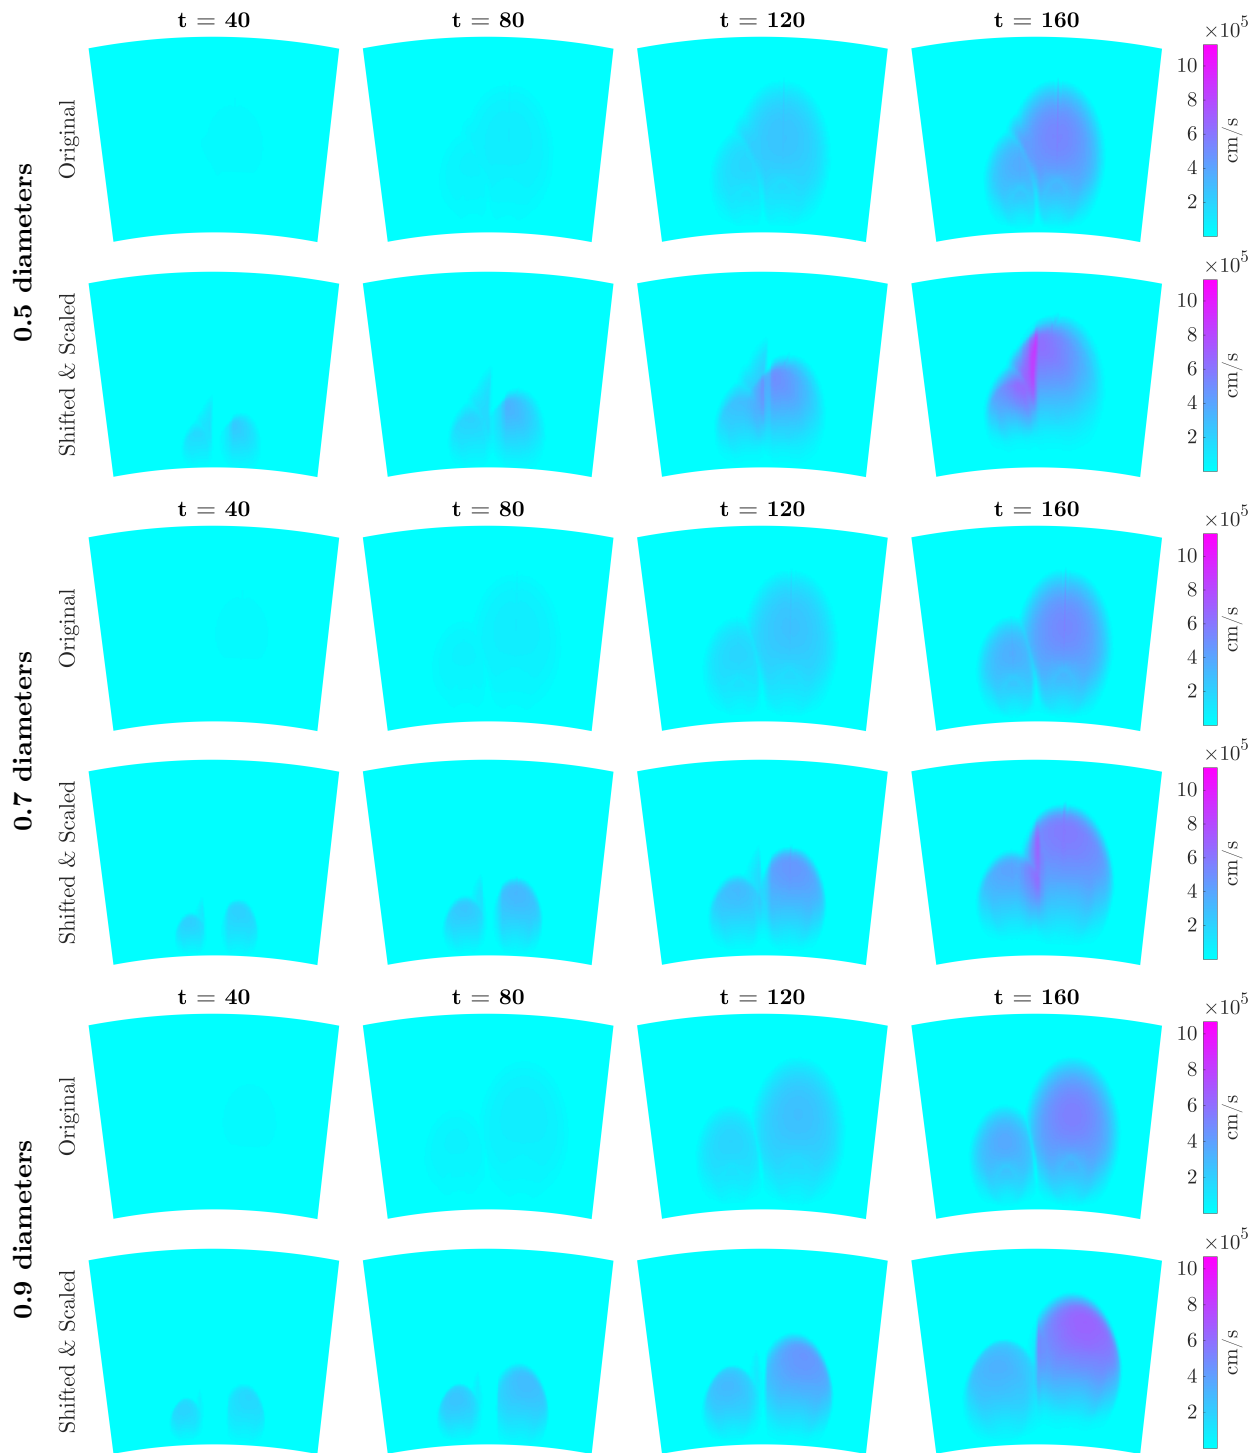


Figure 4.12: 2-mode reconstructions of original vs UnTWIST-modified two-wave data, 0.5-0.9 diameter spacings.




Figure 4.13: 2-mode reconstructions of original vs UnTWIST-modified two-wave data, 1.1-1.5 diameter spacings.

4.7 Discussion

In this chapter, the UnTWIST algorithm was modified to accept inputs of a higher dimension with more invariances. Two-dimensional data with scaling invariance was explored, with examples including a single-wave, noiseless system with known parameters, a single-wave system with more complex physics, and a two-wave system with various amounts of interaction between the two waves. In general, the UnTWIST algorithm was able to find well-fitting models for these invariances, and improved the dimensionality reduction. In cases with close wave spacing, where individual wave quantities were difficult to parse, the UnTWIST algorithm was not able to accurately determine the invariances, and the reconstructions were poor. However, in farther-spaced cases, low-rank reconstructions were improved, yielding more physically-relevant wave shapes while retaining the correct invariance phenomenon.

There are still, however, clear improvements that can be made to the UnTWIST approach for two dimensional data with scaling. Primarily, the naive approach to invariance detection can be improved upon. In these examples, all wave centers were determined with a simple threshold. Accounting for quantities with different heights or intensities may be necessary for accurate wave detection. For example, identifying coherent vortical structures of various vorticities would necessitate a more nuanced detection approach. Another suggested improvement is to the assumption that the wave quantities are approximately circular. With the scaling parameter dependent on approximate radial symmetry, severely non-symmetrical waves would be scaled incorrectly. Detection of the *shape* of the wave quantity could be incorporated using more sophisticated image processing approaches or deep learning. Then, scaling in two dimensions separately could be utilized to account for the shape variations. Additionally, the scaling approach could be improved by utilizing more sophisticated method for resizing data. Recent advances in superresolution, specifically with deep learning [94], have made strides in accuracy and efficiency in superresolving, or rescaling, single images. Overall, advanced object detection algorithms could be employed to better identify irregular wave quantities, with the UnTWIST algorithm applicable no matter the detection method.

Chapter 5

CONCLUSIONS

Data-driven ROMs are of growing importance across the engineering, physical and biological sciences given our increasing ability to exploit emerging sensor technologies to observe and quantify complex dynamical systems. Building models directly from observational data is at the forefront of data-driven science and engineering [14] and addresses the increasing need for interpreting and utilizing big data [95]. Importantly, good ROMs require that an appropriate coordinate system be used in order for a low-rank representation of the dynamics to be achieved [15]. Invariances, particularly rotational and translational, present significant challenges in making ROM models useful for spatiotemporal systems. Simple traveling waves compromise standard ROM architectures, thus requiring additional methods to handle the translational invariance [14, 30–33].

In this work, an automated methods are developed to handle traveling waves [80]. The so-called UnTWIST method uses spectral clustering and machine learning techniques to provide a reference frame pinned to a traveling wave. In Chapter 2, the UnTWIST method is developed, allowing for multiple waves within a one-dimensional data set to be detected, separated, and shifted in moving coordinate frames. The resulting coordinate frames prove to be amenable to traditional dimensionality reduction methods. To demonstrate the use of the method, four physically-motivated example problems were explored, with varying difficulties presented in multiple waves, non-constant wave speeds, nonlinear wave interactions, and complex breathing phenomenon. UnTWIST successfully learned models that improved the dimensionality reduction and matched known wave speeds for each system. In addition, the resulting low-rank reconstructions were better able to capture the essential wave phenomenon when UnTWIST was employed.

In Chapter 3, the use of the UnTWIST method was tested under more challenging experimental data taken from a rotating detonation engine. This system, which has yet-to-be-fully known physics, exhibits rich nonlinear dynamics rife with bifurcations and instabilities. In specific configurations, the RDE maintains stable traveling detonation wave fronts. UnTWIST was used in a preprocessing and refinement step for this data, allowing a given wave in the detonation annulus to be made stationary at a time. The UnTWIST method was successful in finding coordinate frames to center a given wave, allowing for linear and nonlinear models to be built on the wave fronts. Dynamic mode decomposition was employed, and found a low-rank, interpretable model for the interaction between two waves. A nonlinear Lotka-Volterra model was fit for a three-wave data set, capturing the, sharp periodic nonlinearities. Notably, Lotka-Volterra is phenomenologically equivalent to the physical model for this system with spatial dependence removed, as is accomplished with UnTWIST's shifting. These physically relevant models may prove useful in understanding, designing, and controlling RDEs.

In Chapter 4, UnTWIST is modified to accept higher-dimensional data. While this chapter considers only two-dimensional phenomenon with an additional scaling symmetry, the changes to UnTWIST allow for an arbitrary number of symmetries and dimensions of the data to be considered, such as three-dimensional and rotational data. A simple example is demonstrated, to compare the performance of the UnTWIST models to a known system. The two-dimensional shifts and a scaling parameter are found for a growing Gaussian, and prove to match the analytical values. Next, an analog model to a high-altitude nuclear explosion is considered using a single temperature perturbation wave propagated using the Euler equations. The UnTWIST method uncovers sparse models for the wave's speed and growth, and the shifted and scaled data proves to be more amenable to traditional dimensionality reduction methods. A second data set is considered with two detonation waves, each with different initial separation distances. While closely-spaced wave quantities hamper results, low-rank models show that UnTWIST aids the POD in the case of two growing and traveling waves.

Overall, the UnTWIST algorithm has been able to fill a critical gap left in similar approaches. UnTWIST has presented an interpretable framework for the discovery of underlying symmetries, allowing for the description of unknown dynamics. This allows for the use of robust reduced order modeling for these systems, opening tractable computational exploration for systems with no analytical representation. Allowing for translations in more than one dimension has not been approached computationally and improves the dimensionality reduction further. In addition, symmetries such as scaling have not yet been considered by comparable methods. This critical augmentation now allows for complex physics within expanding quantities, such as detonation waves and coherent vortex flow structures, to be accounted for. In doing so, the evolution of the dynamics within the wave can be considered independent of its scale. This acts as a nondimensionalization, unlocking potential relevant phenomenon to be explored and represented with dimensionality reduction.

While the work presented here develops the field of dimensionality reduction for data with symmetries, there are many extensions and improvements that can make UnTWIST robust and applicable to a broader range of systems. The architecture can accommodate any type of symmetry as long as there is an operator to reverse the symmetry. In the future, symmetries like rotations can be considered. With the proper rotation being determined in an initialization step, and a rotation operator to “unrotate” the data into a stationary frame, such models can be learned with UnTWIST. Similarly, three-dimensional data can already be explored with the current architecture with little increase in overall computational cost—only significantly increasing the one-time initialization step. Primarily, incorporating deep learning-based techniques for learning symmetries hold potential to greatly improve UnTWIST. Rather than using naive thresholding, more sophisticated image processing methods could be employed to learn more nuanced symmetry information. Detecting two-dimensional waves of drastically different scales would be improved using local rather than global thresholding. Determining directional scaling would improve scaling for non-radially-symmetric shapes. The work presented here provides a foundational and adaptable architecture for accounting for symmetries for use with traditional dimensionality reduction methods.

Appendix A

	1.0 eV	1.6 eV	2.0 eV
λ	5.0e-2	3e-1	1e-1
ζ	1e1	1e1	1e1
iterations	8,694	3,026	2,753
shift _x	$-11.1025t^4 + 11.5660t^3$	$-0.0044t^5 - 0.0105\sqrt{t}$	$0.009t^3 - 0.0206t^2$
shift _z	$15.3825t^2 + 4.4434t$	$20.1387t^4 + 18.2669t$	$28.6947t^5 + 28.8754t$
$\mathbf{s}(t)$	$-2.3937\sqrt{t} - 0.3567$	$-2.9889\sqrt{t} - 0.3535t$	$-6.1005\sqrt{t} + 0.8180$

Table A.1: Hyperparameters, iterations, errors, and resulting models for Euler equation simulations of various initial energies.

sep. dist.	λ	ζ	iterations	error
0.03	50e1	1e1	4619	3.430e-1
0.05	50e-1	1e1	4883	1.949e-1
0.07	13e-1	1e1	2638	2.718e-1
0.09	9e-1	1e1	4579	1.968e-1
0.11	1.2e-1	1e1	6455	2.780e-1
0.13	5e-1	1e1	1725	1.992e-1
0.15	1e-1	1e1	1355	3.929e-1

Table A.2: Hyperparameters, iterations, and errors for Euler equation simulations of various spacings between two waves.

sep. dist.	shift _{x1}	shift _{x2}	shift _{z1}	shift _{z2}	$\mathbf{s}_1(t)$	$\mathbf{s}_2(t)$
0.03	$-18.5t^2 + 31.7\sqrt{t}$	$-31.9\sqrt{t} - 17.5$	$35.7t^2 + 18.5t$	$13.6t^2 + 20.2t$	$0.607t^2 - 3.56\sqrt{t}$	$0.965t^2 - 3.47\sqrt{t}$
0.05	$-12.2t^2 + 22.7\sqrt{t}$	$-25\sqrt{t} - 24.1$	$39.4t^2 + 2.32t$	$19.3t^2 + 11.3t$	$0.386t^2 - 3.05\sqrt{t}$	$0.959t^2 - 3.12\sqrt{t}$
0.07	$1.61\sqrt{t} + 8.28t$	$-10.6\sqrt{t} - 35.9$	$33.1t^2 + 1.25$	$23.2t^2 + 3.48t$	$-2.78\sqrt{t} - 0.229$	$0.327t^3 - 2.81\sqrt{t}$
0.09	$11.6t^2 - 3.85t$	$-5.83t^2 - 48$	$20.4t^4 + 18.1t$	$16.2t^2 + 4.04\sqrt{t}$	$0.379t^2 - 3.01\sqrt{t}$	$-2.69\sqrt{t} - 0.259$
0.11	$4.34t^5 + 0.676t^4$	$-2.51t^5 - 58.5$	$21.7t^4 + 17.4t$	$20.2t^2 + 1.8$	$0.109t^5 - 3.08\sqrt{t}$	$-2.6\sqrt{t} - 0.308$
0.13	$1.26t^2 - 1.53t$	$0.164t^4 - 69.5$	$34t^2 - 1.5t$	$14.1t^4 + 8.72\sqrt{t}$	$-3.07\sqrt{t} - 0.193t$	$-2.79\sqrt{t} - 0.281$
0.15	$-0.00432t^5 - 0.00599\sqrt{t}$	$0.791t^5 - 80$	$34.2t^2 + 0.807$	$13.7t^4 + 8.86\sqrt{t}$	$-2.95\sqrt{t} - 0.333t$	$-2.7\sqrt{t} - 0.313$

Table A.3: Resulting models for Euler equation simulations of various spacings between two waves.

BIBLIOGRAPHY

- [1] Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. *Annual Review of Fluid Mechanics*, 51:357–377, 2019.
- [2] Kunihiko Taira, Steven L Brunton, Scott Dawson, Clarence W Rowley, Tim Colonius, Beverley J McKeon, Oliver T Schmidt, Stanislav Gordeyev, Vassilios Theofilis, and Lawrence S Ukeiley. Modal analysis of fluid flows: An overview. *AIAA Journal*, 55(12):4013–4041, 2017.
- [3] J. N. Kutz. *Data-Driven Modeling & Scientific Computation: Methods for Complex Systems & Big Data*. Oxford University Press, 2013.
- [4] P. J. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge Monographs in Mechanics. Cambridge University Press, Cambridge, England, 2nd edition, 2012.
- [5] Michele Milano and Petros Koumoutsakos. Neural network modeling for near wall turbulent flow. *Journal of Computational Physics*, 182(1):1–26, 2002.
- [6] Julia Ling, Andrew Kurzawski, and Jeremy Templeton. Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *Journal of Fluid Mechanics*, 807:155–166, 2016.
- [7] Romit Maulik and Omer San. A neural network approach for the blind deconvolution of turbulent flows. *Journal of Fluid Mechanics*, 831:151–181, 2017.
- [8] J.-C. Loiseau, B. R. Noack, and S. L. Brunton. Sparse reduced-order modeling: sensor-based dynamics to full-state estimation. *Journal of Fluid Mechanics*, 844:459–490, 2018.

- [9] MP Brenner, JD Eldredge, and JB Freund. Perspective on machine learning for advancing fluid mechanics. *Physical Review Fluids*, 4(10):100501, 2019.
- [10] Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52:477–508, 2020.
- [11] L. Sirovich. Turbulence and the dynamics of coherent structures, parts I-III. *Q. Appl. Math.*, XLV(3):561–590, 1987.
- [12] Gal Berkooz, Philip Holmes, and John L Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual review of fluid mechanics*, 25(1):539–575, 1993.
- [13] Aaron Towne, Oliver T Schmidt, and Tim Colonius. Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics*, 847:821–867, 2018.
- [14] S. L. Brunton and J. N. Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.
- [15] Peter Benner, Serkan Gugercin, and Karen Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM review*, 57(4):483–531, 2015.
- [16] Kunihiko Taira, Maziar S Hemati, Steven L Brunton, Yiyang Sun, Karthik Duraisamy, Shervin Bagheri, Scott TM Dawson, and Chi-An Yeh. Modal analysis of fluid flows: Applications and outlook. *AIAA journal*, 58(3):998–1022, 2020.
- [17] Bernd R Noack, Marek Morzynski, and Gilead Tadmor. *Reduced-order modelling for flow control*, volume 528. Springer Science & Business Media, 2011.
- [18] Kevin Carlberg, Charbel Bou-Mosleh, and Charbel Farhat. Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approx-

- imations. *International Journal for Numerical Methods in Engineering*, 86(2):155–181, 2011.
- [19] Kevin Carlberg, Matthew Barone, and Harbir Antil. Galerkin v. least-squares Petrov–Galerkin projection in nonlinear model reduction. *Journal of Computational Physics*, 330:693–734, 2017.
- [20] B. R. Noack, K. Afanasiev, M. Morzynski, G. Tadmor, and F. Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *Journal of Fluid Mechanics*, 497:335–363, 2003.
- [21] Gilead Tadmor, Oliver Lehmann, Bernd R Noack, Laurent Cordier, Joël Delville, Jean-Paul Bonnet, and Marek Morzyński. Reduced-order models for closed-loop wake control. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1940):1513–1524, 2011.
- [22] Marek Morzyński, Witold Stankiewicz, Bernd R Noack, Rudibert King, Frank Thiele, and Gilead Tadmor. Continuous mode interpolation for control-oriented models of fluid flow. In *Active Flow Control*, pages 260–278. Springer, 2007.
- [23] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D.S. Henningson. Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 645:115–127, 2009.
- [24] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, August 2010.
- [25] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz. On dynamic mode decomposition: theory and applications. *J. Comp. Dyn.*, 1(2):391–421, 2014.
- [26] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016.

- [27] Athanasios C Antoulas. *Approximation of large-scale dynamical systems*, volume 6. Siam, 2005.
- [28] Jan S Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*. Springer, 2016.
- [29] Alfio Quarteroni, Andrea Manzoni, and Federico Negri. *Reduced basis methods for partial differential equations: an introduction*, volume 92. Springer, 2015.
- [30] Michael Kirby and Dieter Armbruster. Reconstructing phase space from PDE simulations. *Zeitschrift für angewandte Mathematik und Physik ZAMP*, 43(6):999–1022, 1992.
- [31] Clarence W Rowley and Jerrold E Marsden. Reconstruction equations and the Karhunen–Loève expansion for systems with symmetry. *Physica D: Nonlinear Phenomena*, 142(1-2):1–19, 2000.
- [32] Donsub Rim, Scott Moe, and Randall J LeVeque. Transport reversal for model reduction of hyperbolic partial differential equations. *SIAM/ASA J. Uncertainty Quantification*, 6(1):118–150, 2018.
- [33] Julius Reiss, Philipp Schulze, Jörn Sesterhenn, and Volker Mehrmann. The shifted proper orthogonal decomposition: A mode decomposition for multiple transport phenomena. *SIAM Journal on Scientific Computing*, 40(3):A1322–A1344, 2018.
- [34] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11–1–11–37, 2011.
- [35] Angelo Iollo and Damiano Lombardi. Advection modes by optimal mass transfer. *Physical Review E*, 89(2):022923, 2014.
- [36] David Lucia, Paul King, Mark Oxley, and Philip Beran. Reduced order modeling for

- a one-dimensional nozzle flow with moving shocks. In *15th AIAA Computational Fluid Dynamics Conference*, page 2602, 2001.
- [37] Nicolas Cagniard, Yvon Maday, and Benjamin Stamm. Model order reduction for problems with large convection effects. In *Contributions to Partial Differential Equations and Applications*, pages 131–150. Springer, 2019.
- [38] Rambod Mojjani and Maciej Balajewicz. Lagrangian basis method for dimensionality reduction of convection dominated nonlinear flows. *arXiv preprint arXiv:1701.04343*, 2017.
- [39] Francesco Fedele, Ozeair Abessi, and Philip J Roberts. Symmetry reduction of turbulent pipe flows. *Journal of Fluid Mechanics*, 779:390–410, 2015.
- [40] Jean-Frédéric Gerbeau and Damiano Lombardi. Approximated Lax pairs for the reduced order integration of nonlinear evolution equations. *Journal of Computational Physics*, 265:246–269, 2014.
- [41] Philipp Schulze, Julius Reiss, and Volker Mehrmann. Model reduction for a pulsed detonation combustor via shifted proper orthogonal decomposition. In *Active Flow and Combustion Control 2018*, pages 271–286. Springer, 2019.
- [42] Christoph Rasche. Rapid contour detection for image classification. *IET Image Processing*, 12(4):532–538, 2017.
- [43] Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Nat. Acad. Sci.*, page 201517384, 2016.
- [44] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

- [45] Peng Zheng, Travis Askham, Steven L Brunton, J Nathan Kutz, and Aleksandr Y Aravkin. A unified framework for sparse relaxed regularized regression: SR3. *IEEE Access*, 7:1404–1423, 2018.
- [46] Ian H Witten, Eibe Frank, Mark A Hall, and Christopher J Pal. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2016.
- [47] Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems*, pages 849–856, 2002.
- [48] Mark J Ablowitz. *Nonlinear dispersive waves: asymptotic analysis and solitons*, volume 47. Cambridge University Press, 2011.
- [49] Aly-Khan Kassam and Lloyd N Trefethen. Fourth-order time-stepping for stiff PDEs. *SIAM Journal on Scientific Computing*, 26(4):1214–1233, 2005.
- [50] Vijay Anand, Andrew St. George, Robert Driscoll, and Ephraim Gutmark. Investigation of rotating detonation combustor operation with h₂-air mixtures. *International Journal of Hydrogen Energy*, 41(2):1281–1292, Jan 2016.
- [51] Matthew L. Fotia, Fred Schauer, Tom Kaemming, and John Hoke. Experimental study of the performance of a rotating detonation engine with nozzle. *Journal of Propulsion and Power*, 32(3):674–681, May 2016.
- [52] C. A. Nordeen, D. Schwer, F. Schauer, J. Hoke, Th. Barber, and B. Cetegen. Thermodynamic model of a rotating detonation engine. *Combustion, Explosion, and Shock Waves*, 50(5):568–577, 2014.
- [53] Ye-Tao Shao, Meng Liu, and Jian-Ping Wang. Numerical investigation of rotating detonation engine propulsive performance. *Combustion Science and Technology*, 182(11-12):1586–1597, Oct 2010.

- [54] Jorge Sousa, Guillermo Paniagua, and Elena Collado Morata. Thermodynamic analysis of a gas turbine engine with a rotating detonation combustor. *Applied Energy*, 195:247–256, Jun 2017.
- [55] Brent A. Rankin, Matthew L. Fotia, Andrew G. Naples, Christopher A. Stevens, John L. Hoke, Thomas A. Kaemming, Scott W. Theuerkauf, and Frederick R. Schauer. Overview of performance, application, and analysis of rotating detonation engine technologies. *J. Prop. Power*, 33(1):131–143, Jan 2017.
- [56] James Koch, Mitsuru Kurosaka, Carl Knowlen, and J. Nathan Kutz. Mode-locked rotating detonation waves: Experiments and a model equation. *Physical Review E*, 101(1), Jan 2020.
- [57] James Koch, Mitsuru Kurosaka, Carl Knowlen, and J. Nathan Kutz. Multi-scale physics of rotating detonation engines: Autosolitons and modulational instabilities. *arXiv preprint arXiv:2003.06655*, 2020.
- [58] Richard Dyer, Andrew Naples, Thomas Kaemming, John Hoke, and Fred Schauer. Parametric testing of a unique rotating detonation engine design. In *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*. American Institute of Aeronautics and Astronautics, Jan 2012.
- [59] Matthew L. Fotia, John Hoke, and Fred Schauer. Experimental performance scaling of rotating detonation engines operated on gaseous fuels. *Journal of Propulsion and Power*, 33(5):1187–1196, Sep 2017.
- [60] Ian V. Walters, Christopher L. Journell, Aaron Lemcherfi, Rohan Gejji, Stephen D. Heister, and Carson D. Slabaugh. Parametric survey of a natural gas-air rotating detonation engine at elevated pressure. In *AIAA Scitech 2019 Forum*, Jan 2019.
- [61] James Koch and J Nathan Kutz. Modeling thermodynamic trends of rotating detonation engines. *Physics of Fluids*, 32(12):126102, 2020.

- [62] Herman A Haus. Mode-locking of lasers. *IEEE J. Sel. Top. Quan. Elec.*, 6(6):1173–1185, 2000.
- [63] J. Nathan Kutz. Mode-locked soliton lasers. *SIAM Review*, 48(4):629–678, Jan 2006.
- [64] Feng Li, PKA Wai, and J Nathan Kutz. Geometrical description of the onset of multipulsing in mode-locked laser cavities. *JOSA B*, 27(10):2068–2077, 2010.
- [65] James Koch, Lien Chang, Chinmay Upadhye, Kevin Chau, Mitsuru Kurosaka, and Carl Knowlen. Influence of injector-to-annulus area ratio on rotating detonation engine operability. In *AIAA Propulsion and Energy Forum*, 2019.
- [66] Douglas Schwer and Kailas Kailasanath. Numerical investigation of the physics of rotating-detonation-engines. *Proc. Comb. Inst.*, 33(2):2195–2202, Jan 2011.
- [67] Douglas A. Schwer, Ryan F. Johnson, Andrew Kercher, David Kessler, and Andrew T. Corrigan. Progress in efficient, high-fidelity, rotating detonation engine simulations. In *AIAA Scitech 2019 Forum*. AIAA, Jan 2019.
- [68] Sathyanarayanan Subramanian and Joseph Meadows. Novel approach for computational modeling of a non-premixed rotating detonation engine. *Journal of Propulsion and Power*, 36(4):617–631, Jul 2020.
- [69] T. Gaillard, D. Davidenko, and F. Dupoirieux. Numerical simulation of a rotating detonation with a realistic injector designed for separate supply of gaseous hydrogen and oxygen. *Acta Astronautica*, 141:64–78, Dec 2017.
- [70] Jian Sun, Jin Zhou, Shijie Liu, Zhiyong Lin, and Jianhua Cai. *Acta Astronautica*, November.
- [71] Jian Sun, Jin Zhou, Shijie Liu, and Zhiyong Lin. *Acta Astronautica*, November.

- [72] Christopher Lietz, Nathan L. Mundis, Stephen A. Schumaker, and Venke Sankaran. Numerical investigation of rotating detonation rocket engines. In *2018 AIAA Aerospace Sciences Meeting*. AIAA, Jan 2018.
- [73] Rui Zhou and Jian-Ping Wang. Numerical investigation of flow particle paths and thermodynamic performance of continuously rotating detonation engines. *Combustion and Flame*, 159(12):3632–3645, Dec 2012.
- [74] Robert T. Fievisohn and Kenneth H. Yu. Steady-state analysis of rotating detonation engine flowfields with the method of characteristics. *Journal of Propulsion and Power*, 33(1):89–99, Jan 2017.
- [75] Jorge Sousa, James Braun, and Guillermo Paniagua. Development of a fast evaluation tool for rotating detonation combustors. *Applied Mathematical Modelling*, 52:42–52, Dec 2017.
- [76] Tom Kaemming, Matthew L. Fotia, John Hoke, and Fred Schauer. Thermodynamic modeling of a rotating detonation engine through a reduced-order approach. *Journal of Propulsion and Power*, 33(5):1170–1178, Sep 2017.
- [77] Andrew R. Mizener and Frank K. Lu. *Journal of Propulsion and Power*, November.
- [78] Myles Bohon, Richard Bluemner, Alessandro Orchini, Christian Paschereit, and Ephraim Gutmark. Analysis of rdc operation by dynamic mode decomposition (dmd). In *AIAA Propulsion and Energy 2019 Forum*. American Institute of Aeronautics and Astronautics, Aug. 2019.
- [79] Jenna Humble, Swanand V. Sardeshmukh, and Stephen D. Heister. Reduced order modeling of rotational detonation engines. In *AIAA Scitech 2019 Forum*. American Institute of Aeronautics and Astronautics, Jan 2019.

- [80] Ariana Mendible, Steven L Brunton, Wes Lowrie, Aleksander Y Aravkin, and J Nathan Kutz. Dimensionality reduction and reduced-order modeling for traveling wave physics. *Theor. Comput. Fluid Dyn*, 2020.
- [81] Peng Zheng, Travis Askham, Steven L Brunton, J Nathan Kutz, and Aleksandr Y Aravkin. A unified framework for sparse relaxed regularized regression: Sr3. *IEEE Access*, 7:1404–1423, 2018.
- [82] Travis Askham and J Nathan Kutz. Variable projection methods for an optimized dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 17(1):380–416, 2018.
- [83] J.-C. Loiseau and S. L. Brunton. Constrained sparse Galerkin regression. *Journal of Fluid Mechanics*, 838:42–67, 2018.
- [84] Magnus Dam, Morten Brøns, Jens Juul Rasmussen, Volker Naulin, and Jan S Hesthaven. Sparse identification of a predator-prey system from simulation data of a convection model. *Physics of Plasmas*, 24(2):022310, 2017.
- [85] Ariana Mendible. Data-driven modeling of rotating detonation waves. *Physical Review Fluids*, 6(5), 2021.
- [86] Herman Hoerlin. United states high-altitude test experiences. *Rep. LA-6J05, Los Alamos Natl. Lab., Los Alamos, NM*, 1976.
- [87] PR Arendt and H Soicher. Effects of arctic nuclear explosions on satellite radio communication. *Proceedings of the IEEE*, 52(6):672–676, 1964.
- [88] Wilmot N Hess. *The effects of high altitude explosions*. Citeseer, 1964.
- [89] Clawpack Development Team. Clawpack software, 2020. Version 5.7.0.
- [90] Kyle T Mandli, Aron J Ahmadi, Marsha Berger, Donna Calhoun, David L George, Yiannis Hadjimichael, David I Ketcheson, Grady I Lemoine, and Randall J LeVeque.

- Clawpack: building an open source ecosystem for solving hyperbolic pdes. *PeerJ Computer Science*, 2:e68, 2016.
- [91] David I. Ketcheson, Kyle T. Mandli, Aron J. Ahmadi, Amal Alghamdi, Manuel Quezada de Luna, Matteo Parsani, Matthew G. Knepley, and Matthew Emmett. Py-Claw: Accessible, Extensible, Scalable Tools for Wave Propagation Problems. *SIAM Journal on Scientific Computing*, November.
- [92] D. Bale, R. J. LeVeque, S. Mitran, and J. A. Rossmannith. A wave-propagation method for conservation laws and balance laws with spatially varying flux functions. *SIAM J. Sci. Comput.*, 24:955–978, 2002.
- [93] Namrata Vaswani, Thierry Bouwmans, Sajid Javed, and Praneeth Narayanamurthy. Robust subspace learning: Robust pca, robust subspace tracking, and robust subspace recovery. *IEEE signal processing magazine*, 35(4):32–55, 2018.
- [94] Wenming Yang, Xuechen Zhang, Yapeng Tian, Wei Wang, Jing-Hao Xue, and Qingmin Liao. Deep learning for single image super-resolution: A brief review. *IEEE Transactions on Multimedia*, 21(12):3106–3121, 2019.
- [95] Tony Hey, Stewart Tansley, Kristin Tolle, et al. *The fourth paradigm: data-intensive scientific discovery*, volume 1. Microsoft research Redmond, WA, 2009.