



Data-driven modeling of rotating detonation wavesAriana Mendible ^{1,*}, James Koch,² Henning Lange ³, Steven L. Brunton,¹ and J. Nathan Kutz³¹*Department of Mechanical Engineering, University of Washington, Seattle, Washington 98195, USA*²*Oden Institute for Computational & Engineering Sciences, University of Texas at Austin, Austin, Texas 78712, USA*³*Department of Applied Mathematics, University of Washington, Seattle, Washington 98195, USA*

(Received 10 August 2020; accepted 20 April 2021; published 12 May 2021)

The direct monitoring of a rotating detonation engine (RDE) combustion chamber has enabled the observation of combustion front dynamics that are composed of a number of corotating and/or counterrotating coherent traveling shock waves whose nonlinear mode-locking behavior exhibits bifurcations and instabilities which are not well understood. Computational fluid dynamics simulations are ubiquitous in characterizing the dynamics of the RDE's reactive compressible flow. Such simulations are prohibitively expensive when considering multiple engine geometries, different operating conditions, and the long-time dynamics of the mode-locking interactions. Reduced-order models (ROMs) provide a critically enabling simulation framework because they exploit low-rank structure in the data to minimize computational cost and allow for rapid parametrized studies and long-time simulations. However, ROMs are inherently limited by translational invariances manifest by the combustion waves present in RDEs. In this work, we leverage machine learning algorithms to discover moving coordinate frames into which the data are shifted, thus overcoming limitations imposed by the underlying translational invariance of the RDE and allowing for the application of traditional dimensionality reduction techniques. We explore a diverse suite of data-driven ROM strategies for characterizing the complex shock wave dynamics and interactions in the RDE. Specifically, we employ the dynamic mode decomposition and a deep Koopman embedding to give modeling insights and understanding of combustion wave interactions in RDEs.

DOI: [10.1103/PhysRevFluids.6.050507](https://doi.org/10.1103/PhysRevFluids.6.050507)**I. INTRODUCTION**

A rotating detonation engine (RDE) is a novel combustion engine that uses detonative heat release, a nearly constant-volume process, 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 [1,2], the potential for increased thermal efficiency [3,4], and the reduction of propellant pumping requirements [5,6]. The operating dynamics of the RDE includes corotating and counterrotating 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 [7,8] 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 solitonlike interactions between these combustion fronts, and the fundamental underlying

*Corresponding author: mendible@uw.edu

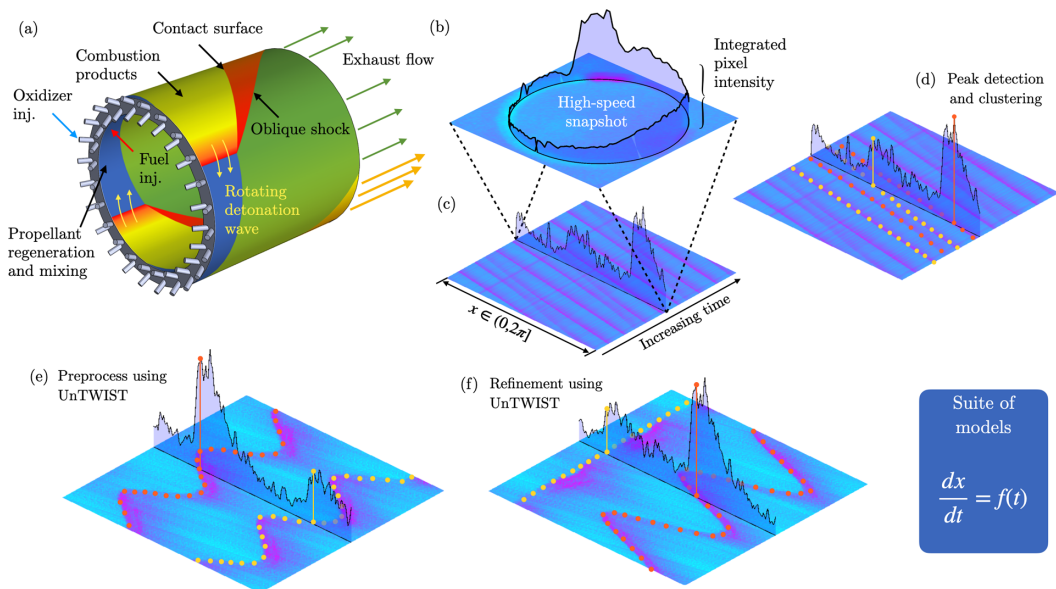


FIG. 1. (a) Schematic of the RDE. (b) Schematic of one time slice of video data, viewing down the axis of rotation of the RDE. (c) Same time slice as in (b) viewed in an (x, t) plot, with each column of the data in time constructed by integrating the pixel intensity along the annulus. (d) Demonstration of the peak detection and clustering necessary to model the wave speeds with UnTWIST. (e) Preliminary processing of (c) using the UnTWIST algorithm. (f) Refining processing of (e) with the UnTWIST algorithm, which then becomes the basis for a suite of data-driven models.

Hopf bifurcation to periodic modulation of the waves [8]. The goal of the present work is to characterize the dynamics of the combustion wave front interactions directly from experimental data, specifically with the goal of developing reduced-order models (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.

Rotating detonation engine 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 Fig. 1(a)]. 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 autoignite 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 [2,9–11] on the RDE dynamics.

The detonations follow attractorlike dynamics that is the manifestation of underlying multiscale balance physics of the driven-dissipative RDE [12]. The RDE is similar in nature to mode-locked

lasers [13,14], where global gain and loss dynamics produces a similar cascading bifurcation diagram of mode-locked states [15]. 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 multiscale 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 multiscale 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 flow-path outside diameter RDE, as described in previous works [7,8,16]. 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 (flow-path 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, it is implied that 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. The inlet boundary condition is implicitly set via a desired flow rate and propellant mixture, thereby constraining the manifold pressures. The outlet boundary condition is set via controlling the back pressure of a large (approximately 4 m³) dump volume. Finally, the exhaust routing of the engine has allowed for the installation of an optical viewport approximately 2 m downstream of the exit plane of the combustor. Each experiment consists of four main phases. First, a prepurge of inert diluent, typically nitrogen, floods the system. Second, the diluent is shut off and propellant begins to flow through the combustor. Third, chemical reactions are triggered, typically via an automotive spark plug or a predetonation tube. In a successful experiment, the self-organization of traveling waves occurs and persists so long as propellant is flowing into the combustor. Finally, the propellant is shut off and diluent is reintroduced 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.* [7,8,16].

Computational fluid dynamic simulations have been heavily relied upon to diagnose the RDE flow fields. These simulations vary from periodic two-dimensional “unwrapped” rectangular domains [17–19] to full detailed three-dimensional engine geometries [20–23]. From these simulations, the canonical RDE flow field is obtained [a cartoon of which is shown in Fig. 1(a)] and relevant metrics can be extracted, such as thrust, specific impulse, available mechanical work, and thermodynamic efficiency [3,24]. However, long-time parametric simulations of RDE dynamics is prohibitively expensive since both the fastest physics (the detonation front) and the slowest physics (mixing and/or exhaustion) 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 flow field [25,26], predicting thermodynamic trends [27], predicting application-based propulsive performance [28], or reproducing the dynamics of the waves [7,8,29,30]. However, because of the multiscale 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 flow fields, recent work [7] 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. Reduced-order models 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 [31,32], whereby a dominant set of correlated modes provide a subspace in which to project the partial differential equation dynamics [33–35]. 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 [33–38]. 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 [39–42]. While these works are critical to addressing the shortcomings of traditional methods, they are limited to applications with constant wave speeds or knowledge of the underlying physics. Mendible *et al.* [43] recently developed 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 [44–46].

II. 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 [32,39–42]. 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 reduced-order models 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) [43] 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, refer to [43].

A. UnTWIST method

Similar to other methods, 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, nonconstant, and nonsmooth 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




FIG. 2. (a) Example of a traveling wave data set. (b) Wave peak points (x_i, t_i) identified using a ridge detection, shown overlaid with the waves. (c) Wave peak points 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.

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: (i) ridge detection to learn the location in (x, t) space of the traveling wave fronts or peaks and (ii) spectral clustering to divide the points (x_i, t_i) into groupings for each wave. For example, these two steps are shown in Fig. 2.

Once the wave fronts are identified and separated, the data are 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

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

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} . The matrix \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 [47] is used to minimize the cost function because of its ability to handle nonconvexity 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 are 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.

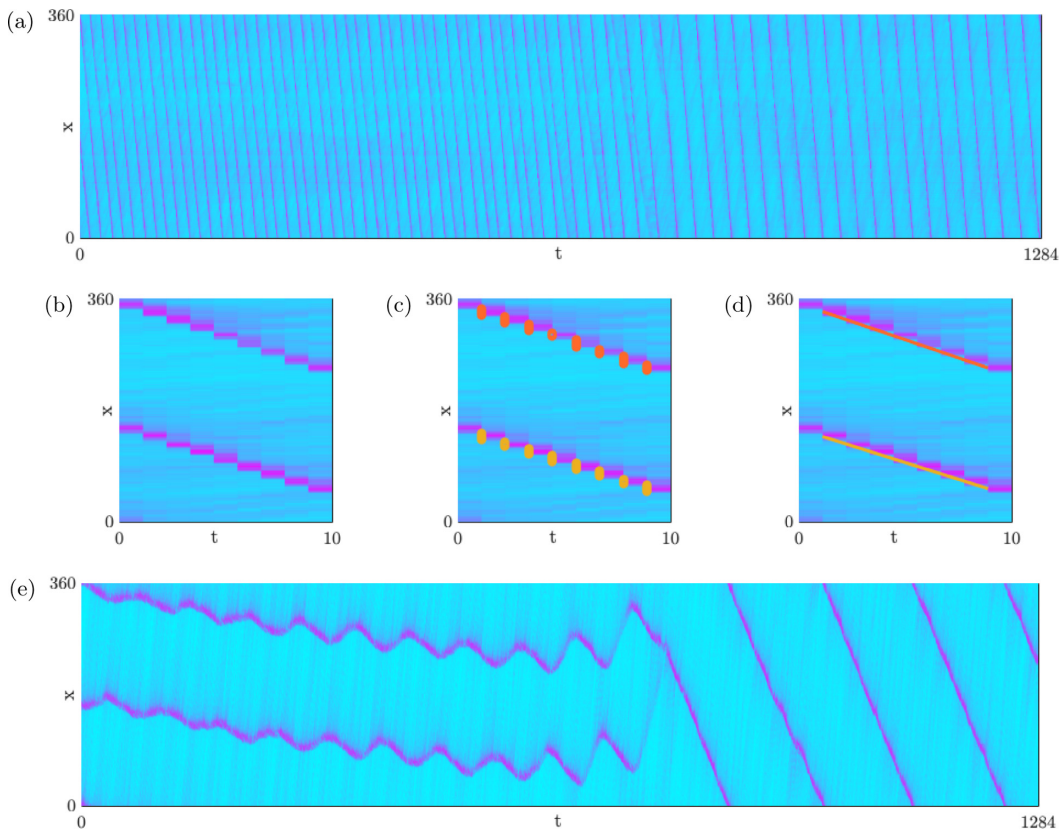


FIG. 3. Preprocessing step using UnTWIST. (a) Example of an original data set, presenting fast-moving shock fronts and a long-time series relative to the spatial dimension. (b) Ten-time-step segment of (a) showing approximately linear-speed shock front propagation. (c) Wave peak points identified and 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.

B. UnTWIST applied to RDE data

An example of the UnTWIST algorithm applied to snapshots of RDE data can be seen in Fig. 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 1000 and 10 000 time steps relative to 180 or 360 spatial points. The snapshots also contain wave fronts that travel on a fast timescale relative to the slow timescale of the relevant dynamics [see Fig. 3(a)], 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 ten time steps of the data are considered, as seen in Fig. 3(b). Using the UnTWIST algorithm and the identified wave peak points as shown in Fig. 3(c), we obtain the best linear-speed model for each wave [Fig. 3(d)]. A single linear-speed shift, the average of the speed models, is applied to original data for the entire time series and gives Fig. 3(e). 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

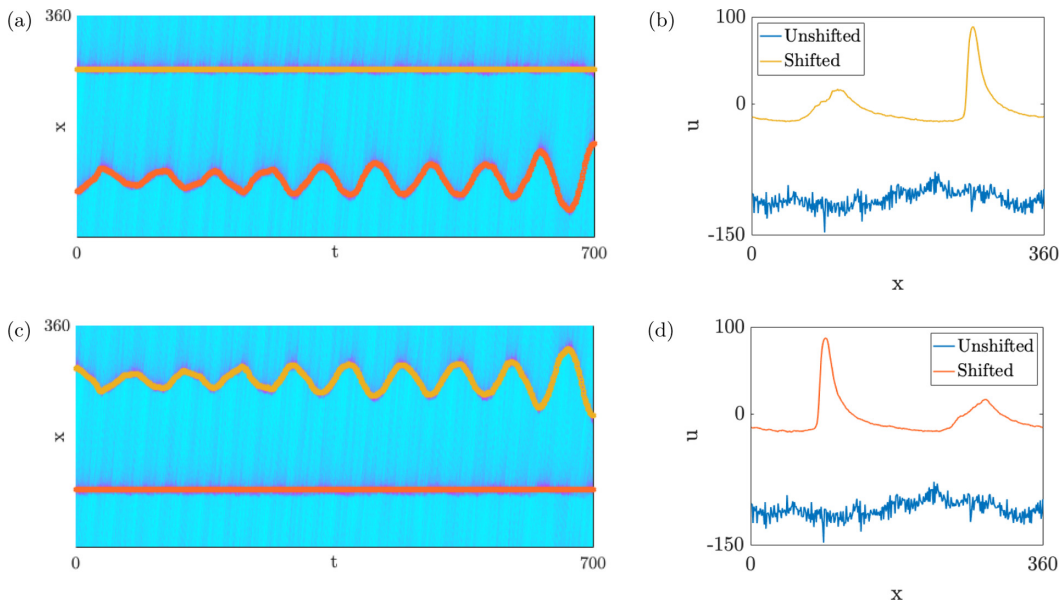


FIG. 4. Refinement step using UnTWIST. (a) Data shifted into the coordinate frame of the first (yellow) wave. (b) Corresponding first mode of a robust dimensionality reduction of the shifted data in (a) compared to the first mode of the original data. The shifted mode shows a clear shock front where 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.

functions such as sinusoids, polynomials, exponentials, and nonlinear combinations of these terms. For the example shown in Fig. 4, sinusoids and exponentials were included in the candidate function library. Once the models were computed, the data were shifted into n refined coordinate frames, one for each shock wave.

The outcome of the second shift can be seen in Fig. 4. Each coordinate frame allows one shock wave to appear stationary at a time. The shifted data are now aligned in a manner that is amenable to traditional dimensionality reduction methods, such as POD. An example of the first mode of a robust dimensionality reduction [48] of the shifted data is shown in Figs. 4(b) and 4(d), compared to the first mode of the POD of the original (laboratory frame) data. Here we only explore the first segment of the time series, before the bifurcation point.

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.

III. 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.

A. Linear dynamics: Dynamic mode decomposition

The dynamic mode decomposition (DMD) [49–52] is an alternative to 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. Thus, in addition to performing a low-rank SVD approximation, it further performs an eigendecomposition on a best-fit linear operator that advances measurements forward in time in the computed subspaces in order to extract critical temporal features. The DMD algorithm is a least-squares regression. In its simplest form [51], one can consider two sets of measurement data \mathbf{u} ,

$$\mathbf{X} = \begin{bmatrix} \vdots & \vdots & \cdots & \vdots \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_{m-1} \\ \vdots & \vdots & \cdots & \vdots \end{bmatrix}, \quad \mathbf{X}' = \begin{bmatrix} \vdots & \vdots & \cdots & \vdots \\ \mathbf{u}'_1 & \mathbf{u}'_2 & \cdots & \mathbf{u}'_{m-1} \\ \vdots & \vdots & \cdots & \vdots \end{bmatrix}, \quad (2)$$

where the primed data are 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

$$\mathbf{A} = \mathbf{X}'\mathbf{X}^\dagger, \quad (3)$$

where the dagger represents the Moore-Penrose pseudoinverse. This gives a least-squares fit to the best linear model fitting the data whose solution is

$$\mathbf{u}_k = \sum_{j=1}^n \phi_j \lambda_j^k b_j = \Phi \Lambda^k \mathbf{b}, \quad (4)$$

where ϕ_j and λ_j are the eigenvectors and eigenvalues of the matrix \mathbf{A} , respectively, 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 $\omega = \ln(\lambda)/\Delta t$, 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.

A simple DMD model is computed by stacking one time shift as

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}(t_1) & \mathbf{x}(t_2) & \cdots & \mathbf{x}(t_{m-2}) \\ \mathbf{x}(t_2) & \mathbf{x}(t_3) & \cdots & \mathbf{x}(t_{m-1}) \end{bmatrix}, \quad \mathbf{X}' = \begin{bmatrix} \mathbf{x}(t_2) & \mathbf{x}(t_3) & \cdots & \mathbf{x}(t_{m-1}) \\ \mathbf{x}(t_3) & \mathbf{x}(t_4) & \cdots & \mathbf{x}(t_m) \end{bmatrix} \quad (5)$$

and utilizing a variable projection method to compute an optimized DMD [53]. This simplified DMD formulation circumvents many of the biases introduced by standard DMD algorithms [49–52] by directly fitting an exponential solution. Figure 5 shows the true distance between the shock fronts compared to discovered DMD model, using one pair of modes to approximate the dynamics.

B. 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 is strongly nonlinear in nature. Figure 6(a) presents three RDE shock fronts interacting in an oscillatory manner. When shifted into the coordinate frame of the top wave, shown in Fig. 6(b), it is




FIG. 5. True distance between the shock wave fronts compared to the linear DMD model.

clear that the middle and lower waves 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

$$\frac{dy}{dt} = \alpha y - \beta yz, \quad (6)$$

$$\frac{dz}{dt} = \delta yz - \gamma z, \quad (7)$$

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, not least because of its capture of the periodic changes with sharp peaks. Preliminary analytical models for traveling waves within RDEs [7] 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 the Lotka-Volterra model. 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 locations of the middle and lower waves. We take y and z to be the locations of the two traveling waves from




FIG. 6. (a) 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.




FIG. 7. Lotka-Volterra model for the two traveling waves in Fig. 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, respectively, and black shows the Lotka-Volterra model. The parameters $\alpha = 0.07$, $\beta = 0.13$, $\delta = 0.10$, and $\gamma = 0.05$ yield a well-fitting model for the two waves which compete for resources.

Fig. 6(b), 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 Fig. 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) [54] algorithm. SINDy has been widely applied to identify reduced-order models for fluid systems [55,56], including those with predator-prey dynamics [57], and is a promising candidate for obtaining low-order models of RDE dynamics.

C. Deep Koopman theory

In the following, ROMs based on Koopman theory are explored that do not rely on an explicit dimensionality reduction technique and can therefore avoid translational invariances. Koopman theory postulates that any nonlinear dynamical system can be lifted by the means of a nonlinear and time-invariant functional, oftentimes referred to as observables, into a space where its time evolution can be described by linear methods. It was first introduced in the seminal paper for Hamiltonian systems [58] but later generalized to continuous-spectrum systems [59]. Even then it was of considerable importance as a building block for advances in ergodic theory [60–64]. Koopman theory has experience renewed interest in the past two decades [65–67].

Letting $u(t)$ be the collected measurements at time t and ψ be the time-invariant observable functional, Koopman theory dictates that a linear operator \mathcal{K} always exists such that

$$\mathcal{K}\psi(u(t)) = \psi(u(t+1)).$$

Note that \mathcal{K} is usually referred to as the Koopman operator and might in practice be infinite dimensional as in, for example, chaotic systems.

Recent applied research has focused on algorithmic approaches to estimate the Koopman operator from measurement data. Early approaches relied on autoencoder structures [68–72] and solved an optimization objective, usually by means of gradient descent, usually composed of terms that encourage linearity in “Koopman space” and reconstruction performance. These approaches were extended in various ways. For example, Bayesian neural networks as encoders were utilized to extend Koopman theory to the probabilistic setting [73]. Furthermore, Champion *et al.* [74] relaxed the linearity requirement and allowed for sparse dynamics in the latent space. Because linearity in Koopman space is part of the optimization objective, these approaches are usually only approximately and locally linear, which in turn impedes their ability to predict far into the future. A more recent approach called Koopman forecast [75] is linear in Koopman space by construction and does not require training an encoder network. In order to overcome its nonlinear and nonconvex objective, the Koopman Forecast algorithm employs gradient descent in conjunction with the fast Fourier transform. In the following, variants of the Koopman Forecast algorithm are introduced and their efficacy in modeling rotating detonation waves is evaluated.

1. Koopman Forecast algorithm

The Koopman Forecast algorithm provides the tools to approximate the Koopman operator from data that is assumed to be quasiperiodic. The quasiperiodicity assumption in turn restricts the Koopman operator to have purely imaginary eigenvalues, i.e., the Koopman operator describes a stable linear dynamical system.

Note that for any linear system $y(t)$ with purely imaginary eigenvalues the following holds:

$$y(t) \propto \begin{bmatrix} \cos(\omega_1 t) \\ \vdots \\ \cos(\omega_n t) \\ \sin(\omega_1 t) \\ \vdots \\ \sin(\omega_n t) \end{bmatrix} := \Omega(\omega t).$$

Because of this, the Koopman Forecast algorithm solves the following optimization problem:

$$E(\vec{\omega}, \Theta) = \sum_{t=1}^T \|u(t) - f_{\Theta}(\Omega(\omega t))\|_2^2.$$

In this case, f_{Θ} is some nonlinear function parametrized by Θ , for example, a neural network. Thus, colloquially speaking, the Koopman Forecast algorithm fits a neural network driven by a linear oscillator to data. Because of the nonlinearity of f , the objective E is notoriously difficult to solve for ω . However, as laid out in [75], by exploiting periodicities in temporally local loss functions in conjunction with coordinate descent, globally optimal values in the direction of ω_i can be obtained. Specifically, for every t , the temporally local loss function $\|u(t) - f_{\Theta}(\Omega(\omega t))\|_2^2$ is periodic in $\frac{2\pi}{t}$. Thus, it is sufficient to sample each temporally local loss function within its first period in order to reconstruct E . The reader is referred to [75] for a more thorough discussion of this technique.

2. Temporal Koopman Forecast with spatial decoder

In the following, we will show how the Koopman Forecast with a spatial decoder can alleviate the problem of translational invariance. For this, a spatial decoder function f_{Θ} is devised that converts

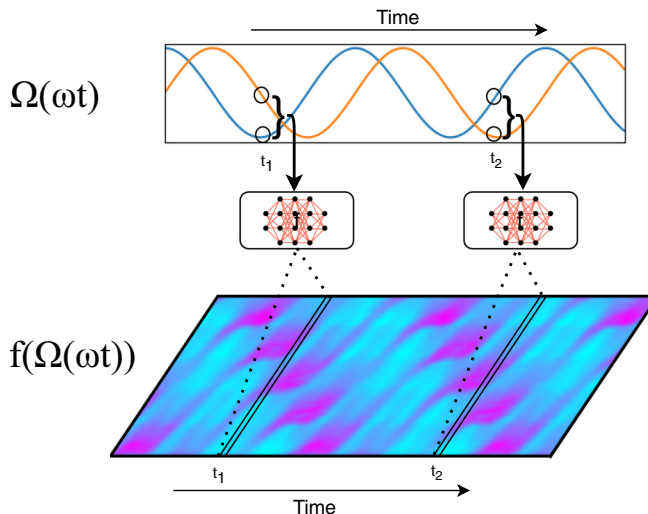


FIG. 8. Graphical depiction of the Koopman Forecast algorithm with one underlying frequency. Colloquially speaking, the Koopman Forecast algorithm resembles a neural network driven by a linear oscillator.

time into space. As f_{Θ} we choose a fully connected feedforward neural network with the topology $2 \rightarrow 32 \rightarrow 32 \rightarrow 180$ and tanh nonlinearity in intermediate layers. The output dimensionality is 180 because space is sampled at 180 locations. The input dimensionality is 2 because we assume the system to be driven by a single frequency. Figure 8 shows graphically the setup for the experiment. Note that f_{Θ} can learn and preserve the nonlinear interactions between waves but because it is a neural network, it is difficult to extract interpretable information about the nature of the nonlinear interactions. Figure 9 shows the spatiotemporal prediction of the system into the future [Fig. 9(b)] against the true data [Fig. 9(a)], with Fig. 9(c) comparing a single time slice. The algorithm correctly explains approximately 75% of the variance. Considering the measurements are obtained by experimentation and therefore exhibit a considerable amount of noise, the Koopman Forecast algorithm seems to perform well.

3. Spatiotemporal Koopman Forecast

The Koopman Forecast model paired with a spatial decoder seems to explain the data reasonably well. However, because f_{Θ} is a neural network, it is hard for practitioners to understand what the algorithm has learned, i.e., the model gives little insight apart from the extracted frequency. In the following a more idealized model will be introduced that is less flexible but also more interpretable. The assumptions of the model are the following: We assume that N modes interact linearly and that both modes travel at a constant speed. Specifically, we model the data as a superposition of N rotating or shifted modes m_i .

Let $u(x, t)$ denote the wave height at position x and time t . Mathematically speaking, we assume

$$u(x, t) = \sum_i^N m_i(f_i(t) + x),$$

where $f_i(t)$ is a time-dependent function that models the offset of mode i at time t . In the next step, we incorporate the knowledge that spatial boundary conditions are periodic. Assume that $u \in [0, K]$. We can find a periodic parametrization of m_i as

$$m_i(x) = g_{\theta_i} \left(\begin{bmatrix} \sin \left(\frac{2\pi}{K} x \right) \\ \cos \left(\frac{2\pi}{K} x \right) \end{bmatrix} \right),$$

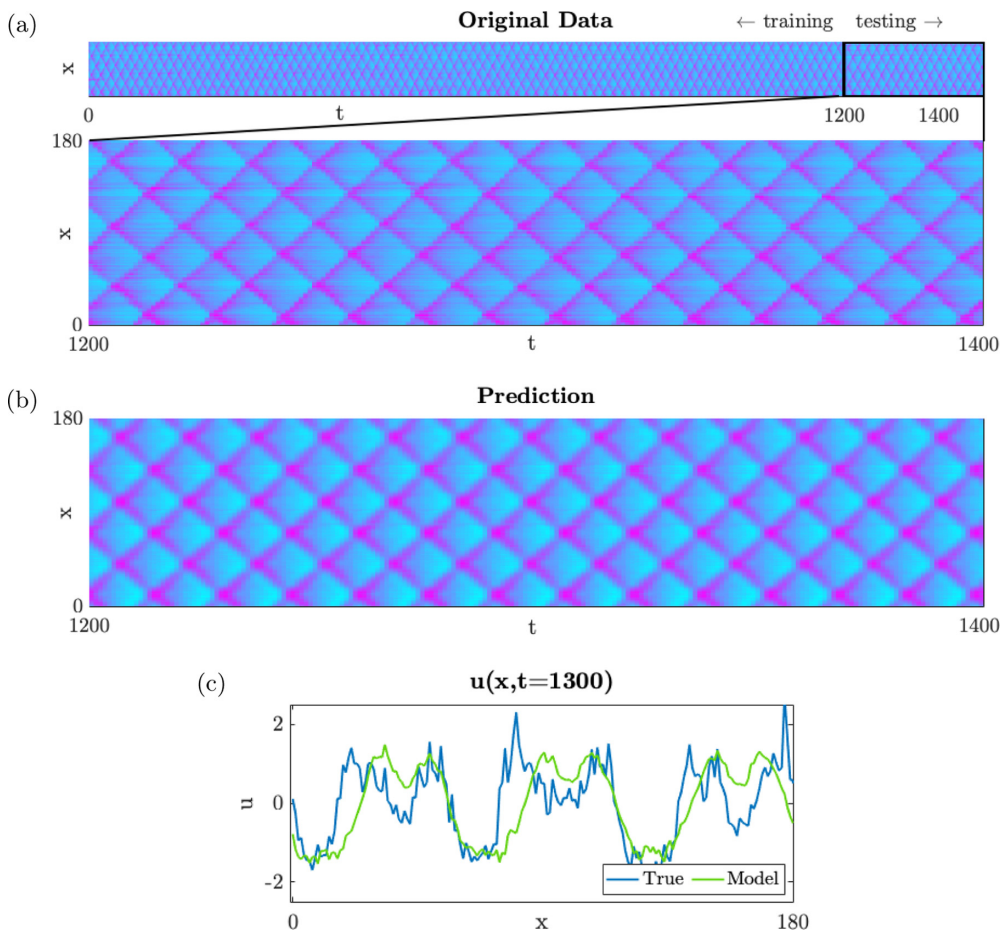


FIG. 9. (a) Original data in the laboratory coordinate frame, showing training data from times $t = [0, 1200]$ and testing data from times $t = [1200, 1400]$, and a broader view of the testing data set. (b) Prediction of the Koopman Forecast algorithm over the same testing time. (c) Comparison of the prediction of the wave shape using the Koopman Forecast algorithm to the ground truth wave shape at an example time step of $t = 1300$.

with θ_i the model parameters of the i th mode (e.g., weights of a neural network). Note that this parametrization is periodic in K because

$$m_i(x + K) = g_{\theta_i} \left(\begin{bmatrix} \sin \left(\frac{2\pi}{K}x + 2\pi \right) \\ \cos \left(\frac{2\pi}{K}x + 2\pi \right) \end{bmatrix} \right) = m_i(x).$$

If we now assume that the offset of the i th mode increases at a constant speed ω_i , we can rewrite $u(x, t)$ as

$$u(x, t) = \sum_i^N g_{\theta_i} \left(\begin{bmatrix} \sin \left(\frac{2\pi}{K}x + \omega_i t \right) \\ \cos \left(\frac{2\pi}{K}x + \omega_i t \right) \end{bmatrix} \right).$$

Let $\Theta = \{\theta_i\}_{i=1}^N$. Fitting $u(x, t)$ to measured data $y(x, t)$ requires solving

$$E(\bar{\omega}, \Theta) = \sum_t \sum_x [y(x, t) - u(x, t)]^2.$$

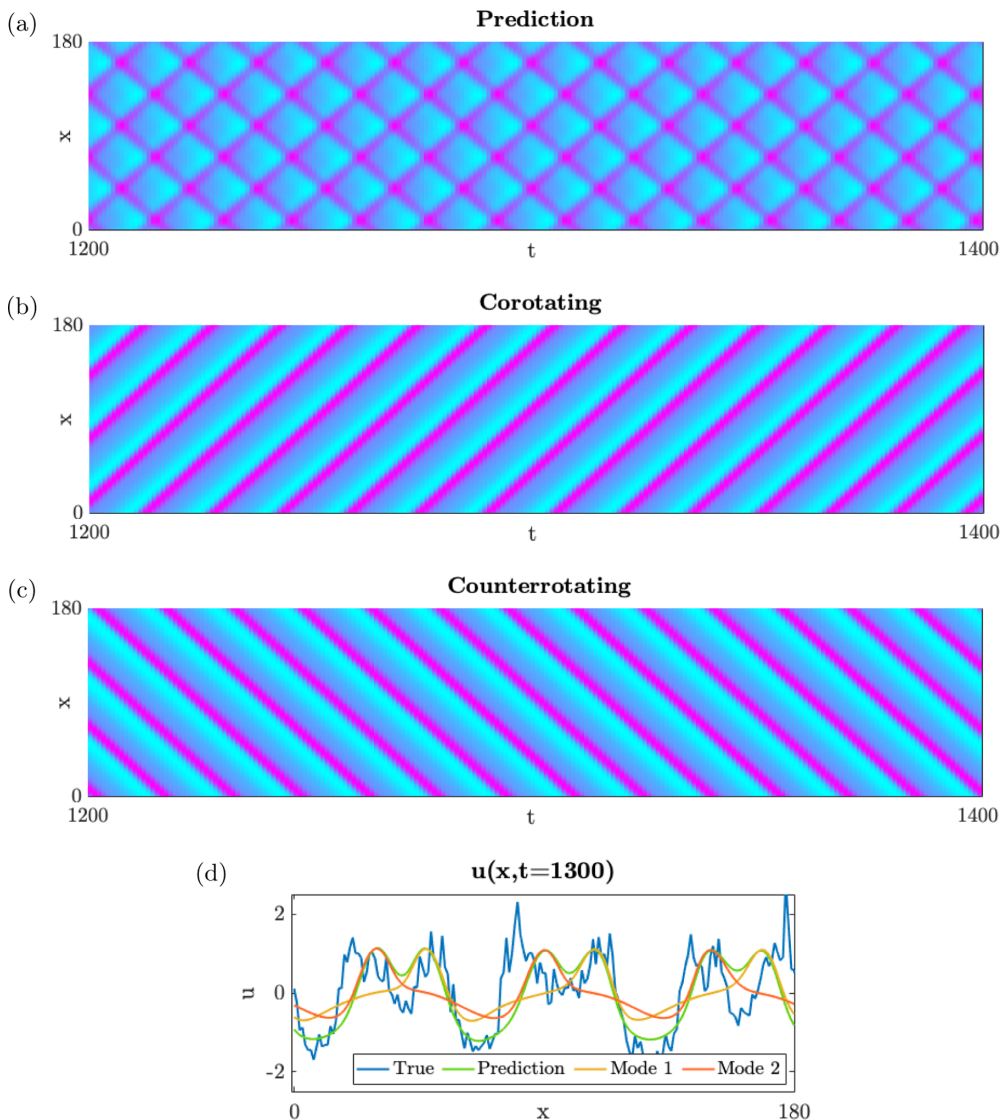


FIG. 10. (a) Prediction of the modal Koopman model on the same data from Fig. 9(a). (b) and (c) Inferred modes. Note that, colloquially, (a) = (b) + (c). (d) Inferred modes at the same time step and the aggregate prediction (mode 1 plus mode 2) in comparison to the ground truth measurement data at an example time $t = 1300$.

Solving this optimization objective for ω is again notoriously difficult as it is not only nonconvex but also nonlinear. However, note that for every x and t , $[y(x, t) - u(x, t)]^2$ is again periodic in $2\pi/t$. Therefore, a strategy analogous to the Koopman Forecast algorithm can be employed to solve for ω .

Figure 10 shows the results when $N = 2$, i.e., when two additive modes are assumed. The modal Koopman model is considerably stiffer as it only explains 65% of the variance in comparison to the Koopman Forecast algorithm, which explains 75%. However, the increase in stiffness also results in an increase of interpretability. The algorithm allows us to decompose the data into the corotating and counterrotating modes [Figs. 10(b) and 10(c)]. This enables practitioners to examine and study

modes individually. Figure 10(d) shows the prediction of the individual modes, the aggregate prediction (mode 1 plus mode 2), and ground truth, respectively.

These two Koopman models prove to be useful tools in reduced-order modeling for RDE data, giving reconstructions and predictions which provide a number of advantages. Primarily, the models are extremely low rank, representing the entire wave field in only two modes for the case of Fig. 10. The models also prove to be applicable on RDE dynamics that behave linearly and nonlinearly, preserving important nonlinear wave interactions. Another advantage of the Koopman approach is that the methods are robust to noise that is rife in the experimental RDE data, effectively acting as a denoising filter as can be seen in Fig. 11(g). However, there is a lack of interpretability in the wave interactions learned in the Koopman algorithms. This limits the use of Koopman and neural network architecture in gaining intuition into the physics at play in the RDE system. While methods such as UnTWIST may not perform as robustly in determining and separating modes in traveling wave systems, they are able to provide more insight into the underlying dynamics.

IV. DISCUSSION AND 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 [32] and addresses the increasing need for interpreting and utilizing a large amount of data [76]. Importantly, good ROMs require that an appropriate coordinate system be used in order for a low-rank representation of the dynamics to be achieved [33]. 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 [32,39–42]. More recently, automated methods have been developed to handle traveling waves [43]. The so-called UnTWIST method uses spectral clustering and machine learning techniques to provide a reference frame pinned to a traveling wave. While the relevant RDE traveling wave dynamics is fully described in one dimension, other open problems exist which demand higher-dimensional translations to be captured. At present, there is a dearth of literature to approach translations in higher than one dimension, and this is an interesting avenue for future work.

We have shown that the UnTWIST method can be 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 multiscale 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 a direct, though qualitative, indication of wave strength, chemical reaction rate, and relative strengths of dissipative effects. Wave strength can be inferred by base-to-peak amplitude of the waves, corresponding to a shock jump condition. The chemical reaction rate can be related to a chemical length scale and is typically observed as the distance from the shock front to the point of greatest luminosity, i.e., the peak of the waveform. Finally, the rate of decay of the expansion side of the waves relates to the timescales associated with expelling the burnt combustion products away from the combustion zone.

Moreover, UnTWIST allows a diversity of model reduction techniques to be applied. We demonstrated three modeling approaches: (i) the DMD for building the best-fit linear dynamics model, (ii) a Lotka-Volterra model for constructing nonlinear dynamical system models for the detonation wave interactions, and (iii) a deep Koopman model that uses a neural network to map the time dynamics to Fourier temporal behavior in order to characterize the dynamics. All three modeling paradigms are relevant as the RDE data and detonation front interactions exhibit dynamics that range from

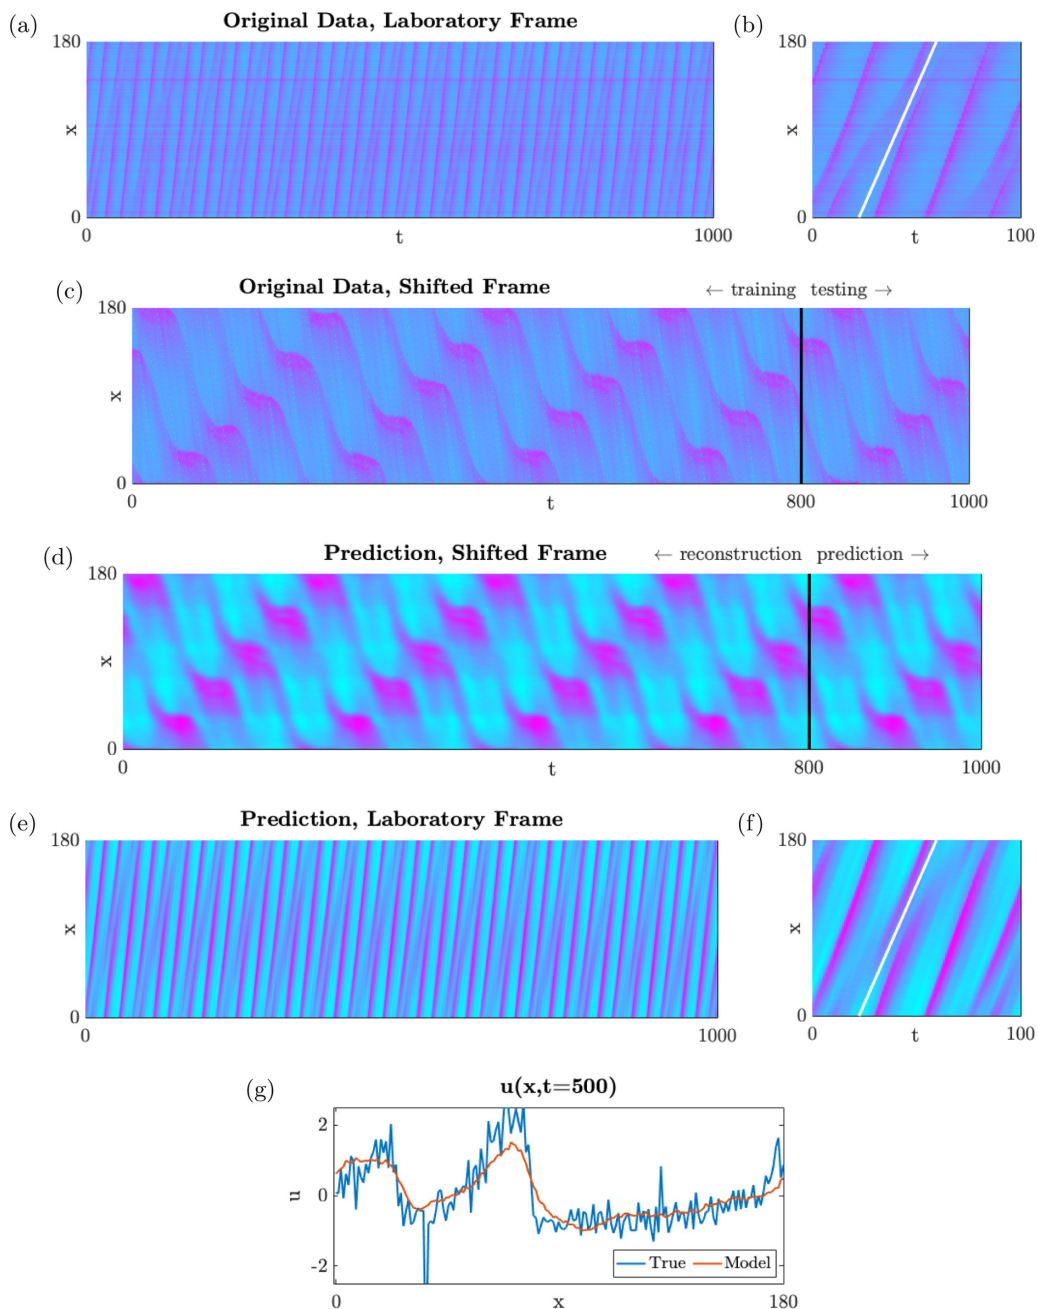


FIG. 11. Comparison of the true measurements to the prediction of the Koopman Forecast algorithm. (a) Full time series of original data in the laboratory frame. (b) Short-time series showcasing the nonlinear interactions between wave fronts, as they do not run parallel to the white guideline. (c) Data shifted into the coordinate frame learned through UnTWIST. (d) Prediction of the same data as in (c) using the Koopman Forecast algorithm. (e) and (f) Koopman Forecast shifted back into the laboratory frame, compared to (a) and (b), respectively. It is clear in (f) that the nonlinear interaction between the waves is well preserved. (g) Comparison of the learned model to the true data at one time slice, $t = 500$, showcasing the robustness to noise and the denoising effect of the model.

approximately linear to strongly nonlinear. Such models provide reductions that enable exploration of the complex and multiscale dynamics of the reactive, compressible fluid dynamics of RDEs.

The architecture presented here emphasizes the critical components necessary 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 [12] 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.

ACKNOWLEDGMENTS

We acknowledge support from the Defense Threat Reduction Agency through Grant No. HDTRA1-18-1-0038. J.N.K. acknowledges support from the Air Force Office of Scientific Research Grant No. FA9550-17-1-0329. S.L.B. acknowledges support from the Army Research Office through Grant No. W911NF-19-1-0045. J.N.K. and S.L.B. acknowledge support from the National Science Foundation (HDR Award No. 1934292). A.M. acknowledges support from Graduate Opportunities and Minority Achievement Program and the Presidential Graduate Fellowship.

-
- [1] V. Anand, A. St. George, R. Driscoll, and E. Gutmark, Investigation of rotating detonation combustor operation with H₂-air mixtures, *Int. J. Hydrogen Energy* **41**, 1281 (2016).
 - [2] M. L. Fotia, F. Schauer, T. Kaemming, and J. Hoke, Experimental study of the performance of a rotating detonation engine with nozzle, *J. Propul. Power* **32**, 674 (2016).
 - [3] C. A. Nordeen, D. Schwer, F. Schauer, J. Hoke, T. Barber, and B. Cetegen, Thermodynamic model of a rotating detonation engine, *Combust. Explos. Shock Waves* **50**, 568 (2014).
 - [4] Y.-T. Shao, M. Liu, and J.-P. Wang, Numerical investigation of rotating detonation engine propulsive performance, *Combust. Sci. Technol.* **182**, 1586 (2010).
 - [5] J. Sousa, G. Paniagua, and E. C. Morata, Thermodynamic analysis of a gas turbine engine with a rotating detonation combustor, *Appl. Energy* **195**, 247 (2017).
 - [6] B. A. Rankin, M. L. Fotia, A. G. Naples, C. A. Stevens, J. L. Hoke, T. A. Kaemming, S. W. Theuerkauf, and F. R. Schauer, Overview of performance, application, and analysis of rotating detonation engine technologies, *J. Propul. Power* **33**, 131 (2017).
 - [7] J. Koch, M. Kurosaka, C. Knowlen, and J. N. Kutz, Mode-locked rotating detonation waves: Experiments and a model equation, *Phys. Rev. E* **101**, 013106 (2020).
 - [8] J. Koch, M. Kurosaka, C. Knowlen, and J. N. Kutz, Multi-scale physics of rotating detonation engines: Autosolitons and modulational instabilities, [arXiv:2003.06655](https://arxiv.org/abs/2003.06655).
 - [9] R. Dyer, A. Naples, T. Kaemming, J. Hoke, and F. Schauer, Parametric testing of a unique rotating detonation engine design, in *Proceedings of the 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, Nashville, 2012* (AIAA, Reston, 2012), paper 121.
 - [10] M. L. Fotia, J. Hoke, and F. Schauer, Experimental performance scaling of rotating detonation engines operated on gaseous fuels, *J. Propul. Power* **33**, 1187 (2017).
 - [11] I. V. Walters, C. L. Journell, A. Lemcherfi, R. Gejji, S. D. Heister, and C. D. Slabaugh, Parametric survey of a natural gas-air rotating detonation engine at elevated pressure, in *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech), San Diego, 2019* (AIAA, Reston, 2019), paper 1510.
 - [12] J. Koch and J. N. Kutz, Modeling thermodynamic trends of rotating detonation engines, *Phys. Fluids* **32**, 126102 (2020).

- [13] H. A. Haus, Mode-locking of lasers, *IEEE J. Sel. Top. Quantum Electron.* **6**, 1173 (2000).
- [14] J. N. Kutz, Mode-locked soliton lasers, *SIAM Rev.* **48**, 629 (2006).
- [15] F. Li, P. Wai, and J. N. Kutz, Geometrical description of the onset of multi-pulsing in mode-locked laser cavities, *J. Opt. Soc. Am. B* **27**, 2068 (2010).
- [16] J. Koch, L. Chang, C. Upadhye, K. Chau, M. Kurosaka, and C. Knowlen, Influence of injector-to-annulus area ratio on rotating detonation engine operability, in *Proceedings of the AIAA Propulsion and Energy 2019 Forum, Indianapolis* (AIAA, Reston, 2019), paper 4038.
- [17] D. Schwer and K. Kailasanath, Numerical investigation of the physics of rotating-detonation-engines, *Proc. Combust. Inst.* **33**, 2195 (2011).
- [18] D. A. Schwer, R. F. Johnson, A. Kercher, D. Kessler, and A. T. Corrigan, Progress in efficient, high-fidelity, rotating detonation engine simulations, in *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech), San Diego, 2019* (Ref. [11]), paper 2018.
- [19] S. Subramanian and J. Meadows, Novel approach for computational modeling of a non-premixed rotating detonation engine, *J. Propul. Power* **36**, 617 (2020).
- [20] 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 Astronaut.* **141**, 64 (2017).
- [21] J. Sun, J. Zhou, S. Liu, Z. Lin, and J. Cai, Effects of injection nozzle exit width on rotating detonation engine, *Acta Astronaut.* **140**, 388 (2017).
- [22] J. Sun, J. Zhou, S. Liu, and Z. Lin, Numerical investigation of a rotating detonation engine under premixed/non-premixed conditions, *Acta Astronaut.* **152**, 630 (2018).
- [23] C. Lietz, N. L. Mundis, S. A. Schumaker, and V. Sankaran, Numerical investigation of rotating detonation rocket engines, in *Proceedings of the 2018 AIAA Aerospace Sciences Meeting, Kissimmee, 2018* (AIAA, Reston, 2018), paper 0882.
- [24] R. Zhou and J.-P. Wang, Numerical investigation of flow particle paths and thermodynamic performance of continuously rotating detonation engines, *Combust. Flame* **159**, 3632 (2012).
- [25] R. T. Fievisohn and K. H. Yu, Steady-state analysis of rotating detonation engine flowfields with the method of characteristics, *J. Propul. Power* **33**, 89 (2017).
- [26] J. Sousa, J. Braun, and G. Paniagua, Development of a fast evaluation tool for rotating detonation combustors, *Appl. Math. Model.* **52**, 42 (2017).
- [27] T. Kaemming, M. L. Fotia, J. Hoke, and F. Schauer, Thermodynamic modeling of a rotating detonation engine through a reduced-order approach, *J. Propul. Power* **33**, 1170 (2017).
- [28] A. R. Mizener and F. K. Lu, Low-order parametric analysis of a rotating detonation engine in rocket mode, *J. Propul. Power* **33**, 1543 (2017).
- [29] M. Bohon, R. Bluemner, A. Orchini, C. Paschereit, and E. Gutmark, Analysis of RDC operation by Dynamic Mode Decomposition (DMD), in *Proceedings of the AIAA Propulsion and Energy 2019 Forum, Indianapolis* (Ref. [16]), paper 4377.
- [30] J. Humble, S. V. Sardeshmukh, and S. D. Heister, Reduced order modeling of rotational detonation engines, in *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech), San Diego, 2019* (Ref. [11]), paper 0201.
- [31] J. N. Kutz, *Data-Driven Modeling & Scientific Computation: Methods for Complex Systems & Big Data* (Oxford University Press, Oxford, 2013).
- [32] S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control* (Cambridge University Press, Cambridge, 2019).
- [33] P. Benner, S. Gugercin, and K. Willcox, A survey of projection-based model reduction methods for parametric dynamical systems, *SIAM Rev.* **57**, 483 (2015).
- [34] K. Taira, S. L. Brunton, S. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley, Modal analysis of fluid flows: An overview, *AIAA J.* **55**, 4013 (2017).
- [35] K. Taira, M. S. Hemati, S. L. Brunton, Y. Sun, K. Duraisamy, S. Bagheri, S. T. M. Dawson, and C.-A. Yeh, Modal analysis of fluid flows: Applications and outlook, *AIAA J.* **58**, 998 (2020).
- [36] A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems* (SIAM, Philadelphia, 2005), Vol. 6.

- [37] J. S. Hesthaven, G. Rozza, and B. Stamm, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations* (Springer, Berlin, 2016).
- [38] A. Quarteroni, A. Manzoni, and F. Negri, *Reduced Basis Methods for Partial Differential Equations: An Introduction*, La Matematica per il 3+2 Vol. 92 (Springer International, Cham, 2015).
- [39] M. Kirby and D. Armbruster, Reconstructing phase space from PDE simulations, *Z. Angew Math. Phys.* **43**, 999 (1992).
- [40] C. W. Rowley and J. E. Marsden, Reconstruction equations and the Karhunen-Loève expansion for systems with symmetry, *Physica D* **142**, 1 (2000).
- [41] D. Rim, S. Moe, and R. J. LeVeque, Transport reversal for model reduction of hyperbolic partial differential equations, *SIAM/ASA J. Uncertainty Quantif.* **6**, 118 (2018).
- [42] J. Reiss, P. Schulze, J. Sesterhenn, and V. Mehrmann, The shifted proper orthogonal decomposition: A mode decomposition for multiple transport phenomena, *SIAM J. Sci. Comput.* **40**, A1322 (2018).
- [43] A. Mendible, S. L. Brunton, W. Lowrie, A. Y. Aravkin, and J. N. Kutz, Dimensionality reduction and reduced-order modeling for traveling wave physics, *Theor. Comput. Fluid Dyn.* **34**, 385 (2020).
- [44] K. Duraisamy, G. Iaccarino, and H. Xiao, Turbulence modeling in the age of data, *Annu. Rev. Fluid Mech.* **51**, 357 (2019).
- [45] M. Brenner, J. Eldredge, and J. Freund, Perspective on machine learning for advancing fluid mechanics, *Phys. Rev. Fluids* **4**, 100501 (2019).
- [46] S. L. Brunton, B. R. Noack, and P. Koumoutsakos, Machine learning for fluid mechanics, *Annu. Rev. Fluid Mech.* **52**, 477 (2020).
- [47] P. Zheng, T. Askham, S. L. Brunton, J. N. Kutz, and A. Y. Aravkin, A unified framework for sparse relaxed regularized regression: SR3, *IEEE Access* **7**, 1404 (2018).
- [48] E. J. Candès, X. Li, Y. Ma, and J. Wright, Robust principal component analysis? *J. ACM* **58**, 11 (2011).
- [49] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5 (2010).
- [50] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. Henningson, Spectral analysis of nonlinear flows, *J. Fluid Mech.* **645**, 115 (2009).
- [51] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz, On dynamic mode decomposition: Theory and applications, *J. Comput. Dyn.* **1**, 391 (2014).
- [52] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems* (SIAM, Philadelphia, 2016).
- [53] T. Askham and J. N. Kutz, Variable projection methods for an optimized dynamic mode decomposition, *SIAM J. Appl. Dyn. Syst.* **17**, 380 (2018).
- [54] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proc. Natl. Acad. Sci. USA* **113**, 3932 (2016).
- [55] J.-C. Loiseau and S. L. Brunton, Constrained sparse Galerkin regression, *J. Fluid Mech.* **838**, 42 (2018).
- [56] J.-C. Loiseau, B. R. Noack, and S. L. Brunton, Sparse reduced-order modeling: Sensor-based dynamics to full-state estimation, *J. Fluid Mech.* **844**, 459 (2018).
- [57] M. Dam, M. Brøns, J. Juul Rasmussen, V. Naulin, and J. S. Hesthaven, Sparse identification of a predator-prey system from simulation data of a convection model, *Phys. Plasmas* **24**, 022310 (2017).
- [58] B. O. Koopman, Hamiltonian systems and transformation in Hilbert space, *Proc. Natl. Acad. Sci. USA* **17**, 315 (1931).
- [59] B. O. Koopman and J. v. Neumann, Dynamical systems of continuous spectra, *Proc. Natl. Acad. Sci. USA* **18**, 255 (1932).
- [60] G. D. Birkhoff, Proof of the ergodic theorem, *Proc. Natl. Acad. Sci. USA* **17**, 656 (1931).
- [61] J. v. Neumann, Proof of the quasi-ergodic hypothesis, *Proc. Natl. Acad. Sci. USA* **18**, 70 (1932).
- [62] G. D. Birkhoff and B. Koopman, Recent contributions to the ergodic theory, *Proc. Natl. Acad. Sci. USA* **18**, 279 (1932).
- [63] J. v. Neumann, Physical applications of the ergodic hypothesis, *Proc. Natl. Acad. Sci. USA* **18**, 263 (1932).
- [64] C. C. Moore, Ergodic theorem, ergodic theory, and statistical mechanics, *Proc. Natl. Acad. Sci. USA* **112**, 1907 (2015).

- [65] I. Mezić, Spectral properties of dynamical systems, model reduction and decompositions, *Nonlinear Dyn.* **41**, 309 (2005).
- [66] M. Budišić, R. Mohr, and I. Mezić, Applied Koopmanism, *Chaos* **22**, 047510 (2012).
- [67] I. Mezić, Analysis of fluid flows via spectral properties of the Koopman operator, *Annu. Rev. Fluid Mech.* **45**, 357 (2013).
- [68] S. E. Otto and C. W. Rowley, Linearly recurrent autoencoder networks for learning dynamics, *SIAM J. Appl. Dyn. Syst.* **18**, 558 (2019).
- [69] B. Lusch, J. N. Kutz, and S. L. Brunton, Deep learning for universal linear embeddings of nonlinear dynamics, *Nat. Commun.* **9**, 4950 (2018).
- [70] E. Yeung, S. Kundu, and N. Hodas, Learning deep neural network representations for Koopman operators of nonlinear dynamical systems, in *Proceedings of the 2019 American Control Conference, Philadelphia* (IEEE, Piscataway, 2019), pp. 4832–4839.
- [71] C. Wehmeyer and F. Noé, Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics, *J. Chem. Phys.* **148**, 241703 (2018).
- [72] N. Takeishi, Y. Kawahara, and T. Yairi, Learning Koopman invariant subspaces for dynamic mode decomposition, in *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran, Red Hook, 2017), pp. 1130–1140.
- [73] S. Pan and K. Duraisamy, Physics-informed probabilistic learning of linear embeddings of non-linear dynamics with guaranteed stability, *SIAM J. Appl. Dyn. Syst.* **19**, 480 (2020).
- [74] K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton, Data-driven discovery of coordinates and governing equations, *Proc. Natl. Acad. Sci. USA* **116**, 22445 (2019).
- [75] H. Lange, S. L. Brunton, and J. N. Kutz, From Fourier to Koopman: Spectral methods for long-term time series prediction, *J. Mach. Learn. Research* **22**, 1 (2021).
- [76] T. Hey, S. Tansley, and K. Tolle, *The Fourth Paradigm: Data-Intensive Scientific Discovery* (Microsoft Research, Redmond, 2009), Vol. 1.