


Patterns

Improving reduced-order models through nonlinear decoding of projection-dependent outputs

Graphical abstract



Highlights

- Undesired projection topology can be formed through dimensionality reduction
- Nonlinear decoding of quantities of interest (QoIs) promotes improved data projections
- Projection topology is improved if some of the QoIs depend on the projection operator
- Reducing nonuniqueness and steep gradients in QoIs enhances data-driven workflows

Authors

Kamila Zdybał, Alessandro Parente, James C. Sutherland

Correspondence

kamilazdybal@gmail.com (K.Z.),
alessandro.parente@ulb.be (A.P.),
james.sutherland@utah.edu (J.C.S.)

In brief

Effective dimensionality reduction is crucial in the data-driven modeling of high-dimensional systems. In this article, the authors show that informing low-dimensional data projections by quantities of interest and incorporating a nonlinear reconstruction error measure during dimensionality reduction leads to improved data representations in latent space. Their dimensionality reduction approach is applicable to multivariate datasets across many disciplines of science and engineering.



Article

Improving reduced-order models through nonlinear decoding of projection-dependent outputs

Kamila Zdybał^{1,2,4,*}, Alessandro Parente^{1,2,*} and James C. Sutherland^{3,*}¹Université Libre de Bruxelles, École Polytechnique de Bruxelles, Aero-Thermo-Mechanics Laboratory, Brussels, Belgium²BRITE: Brussels Institute for Thermal-Fluid Systems and Clean Energy, Brussels, Belgium³Department of Chemical Engineering, University of Utah, Salt Lake City, UT, USA⁴Lead contact*Correspondence: kamilazdybal@gmail.com (K.Z.), alessandro.parente@ulb.be (A.P.), james.sutherland@utah.edu (J.C.S.)<https://doi.org/10.1016/j.patter.2023.100859>

THE BIGGER PICTURE Large datasets are abundant in various scientific and engineering disciplines. Multiple physical variables are frequently gathered into one dataset, leading to high data dimensionality. Visualizing and modeling multivariate datasets can be achieved through dimensionality reduction. However, in many reduction techniques to date, there is no guarantee that the reduced data representation will possess certain desired topological qualities. We show that the quality of reduced data representations can be significantly improved by informing data projections by target quantities of interest (QoIs), some of which are functions of the projection itself. The target QoIs can include closure terms required in modeling, important physical variables, or class labels in the case of categorical data. This work can have particular relevance in data visualization and efficient modeling of dynamical systems with many degrees of freedom, as well as in fundamental research of representation learning.



Development/Pre-production: Data science output has been rolled out/validated across multiple domains/problems

SUMMARY

A fundamental hindrance to building data-driven reduced-order models (ROMs) is the poor topological quality of a low-dimensional data projection. This includes behavior such as overlapping, twisting, or large curvatures or uneven data density that can generate nonuniqueness and steep gradients in quantities of interest (QoIs). Here, we employ an encoder-decoder neural network architecture for dimensionality reduction. We find that nonlinear decoding of projection-dependent QoIs, when embedded in a dimensionality reduction technique, promotes improved low-dimensional representations of complex multiscale and multiphysics datasets. When data projection (encoding) is affected by forcing accurate nonlinear reconstruction of the QoIs (decoding), we minimize nonuniqueness and gradients in representing QoIs on a projection. This in turn leads to enhanced predictive accuracy of a ROM. Our findings are relevant to a variety of disciplines that develop data-driven ROMs of dynamical systems such as reacting flows, plasma physics, atmospheric physics, or computational neuroscience.

INTRODUCTION

High-dimensional datasets are collected in large volumes in various scientific and engineering disciplines. This is due to advances in high-performance computing that allow for the collection of data from numerical simulations (e.g., from multi-component turbulent reacting flow simulations) and advances

in experimental techniques that enable collection of massive numbers of features simultaneously (e.g., the entire genome in single-cell data). Techniques that help us understand multivariate data and build data-driven models have thus become a necessary apparatus in a researcher's toolbox. To this end, efficient dimensionality reduction is crucial, as it allows one to represent high-dimensional data in a lower-dimensional latent space,



where, it is hoped, understanding^{1–5} and modeling^{6–9} of complex systems are easier.

Many dimensionality reduction strategies can provide suboptimal data representations when they are not informed by the intended use case or by the knowledge of the problem at hand. Commonly used techniques do not inform the projection topology by variables other than the independent state variables.^{10–13} This hindrance has an impact on efficient data visualization and hypothesis generation because relevant information can become obscured in latent space. It also has an impact on building data-driven reduced-order models (ROMs) of complex dynamical systems. The predictive capability of a ROM can be hampered by poor-quality low-dimensional data representation.

Independent of the discipline from which the dynamical system originates, ROMs share some salient characteristics.^{7,14–19} First, a low-dimensional basis is computed from the state variables known from either simulations or experiments. Next, we map the governing equations of a dynamical system (e.g., nonlinear, coupled partial differential equations [PDEs]) to the lower-dimensional basis.⁷ When a nonlinear regression map is incorporated into the reduced-order modeling workflow (e.g., for handling the nonlinear terms in the projected PDEs), the projection should be optimized toward representing the reconstructed target variables, different from the state variables. The success of that reconstruction, and hence the success of a ROM, depends on the topological quality of the data projection. We argue that data projections should be informed by the quantities of interest (QoIs) that are the regression targets and include the QoIs whose reconstruction is required at ROM runtime.^{20,21}

We formulate an autoencoder-like neural network architecture with a linear encoder and a nonlinear decoder. The difference from the ordinary autoencoder (AE)^{22,23} is that at the output layer we reconstruct various QoIs, different from the original state variables but crucial for constructing the ROM. When the dimensionality reduction task and the nonlinear reconstruction task communicate with each other through backpropagation, a linear projection created by the encoder is optimized to result in a good nonlinear reconstruction of important QoIs by the decoder. This approach naturally penalizes projections that exhibit nonuniqueness or large gradients, since any difficulties in representing QoIs on a projection immediately increase the mean-squared-error (MSE) loss function during training. Evidence from the existing research^{24–26} suggests that such a joint encoding-decoding^{27–31} approach provides improvements to the topology of a low-dimensional data projection. Here, we demonstrate quantitatively that this is the case.

The recently developed cost function that allows one to assess the quality of a low-dimensional manifold^{32,33} is extensively used in this work to quantify the benefits of the proposed approach. Linear reconstruction error optimality³⁴ often proves an insufficient metric to guarantee the quality of the projection topology.^{33,35} Therefore, we require optimality measures other than those present in, e.g., principal-component analysis (PCA) or its time-series variant, proper orthogonal decomposition (POD). We find that nonlinear reconstruction of QoIs provides significant improvements in (linear) projection topology as measured by the quantitative cost function. Moreover, further improvements are achieved when the very definition of some of the QoIs depends on the projection operator. Our conclusions

hold for different projection dimensionalities and for a variety of reacting flow datasets representing the combustion of different fuels in air.

RESULTS

The outstanding challenges of data-driven reduced-order modeling

It has been recognized that, under some conditions, ROMs work very well, but they can fail dramatically under others. At the same time, many aspects of ROMs are not yet understood. In particular, undesirable topological behaviors can be introduced on a data projection through dimensionality reduction. The most problematic behavior is an overlap on a projection. An overlap can occur on a projection when different observations in state space are projected onto the same location in the latent space. Overlap can lead to nonuniqueness in representing a QoI on a projection. When observations inside the overlapping region correspond to different values in a QoI, they can be mapped onto only one output (QoI) value based on the same input (lower-dimensional coordinates) value. This leads to errors in modeling QoIs over a projection and can hinder the performance of a ROM and data-driven workflows in general.³⁶ Two important questions that ought to be answered to enhance the predictive performance and our understanding of ROMs are the following:

1. How can we define a low-dimensional data projection such that it possesses certain desired topological qualities?
2. How do poor low-dimensional data projections affect reduced-order modeling?

Here, we tackle these outstanding questions and demonstrate that significant improvements in projection topology can be achieved thanks to nonlinear decoding of QoIs.

From a full-order model to a reduced-order model

The PDEs that describe the full-order model (FOM) often have the form of an advection-diffusion equation with source terms shared between various disciplines. In this work, we are focused on a FOM that describes a multicomponent reacting flow. The PDEs can be written in matrix form:

$$\frac{\partial \mathbf{X}^T}{\partial t} = -\nabla \cdot \mathbf{C}^T - \nabla \cdot \mathbf{D}^T + \mathbf{S}^T, \quad (\text{Equation 1})$$

where $\mathbf{X} \in \mathbb{R}^{N \times Q}$ is the thermochemical state, $\mathbf{C} \in \mathbb{R}^{d \times N \times Q}$ is the convective flux, $\mathbf{D} \in \mathbb{R}^{d \times N \times Q}$ is the diffusive flux, and $\mathbf{S} \in \mathbb{R}^{N \times Q}$ are the thermochemical source terms. In a reacting flow, \mathbf{S} are nonlinear functions of the state variables and can be computed explicitly from a known chemical mechanism. Here, Q is the number of state variables (equal to the problem dimensionality), d is the number of spatial dimensions of the problem, and $N \gg Q$ is the number of observations (equal to the number of grid points). The magnitude of Q depends on the number of chemical species in the mechanism. The appropriate formulation of \mathbf{X} , \mathbf{C} , \mathbf{D} , and \mathbf{S} depends on the problem at hand.

Data-driven model order reduction starts with changing the basis to represent the original high-dimensional system as per Equation 1. Let $\mathbf{A} \in \mathbb{R}^{Q \times q}$ be the matrix defining the new




Figure 1. The QoI-aware encoder-decoder dimensionality reduction strategy and its effect on the projection topology

(A) The topology of the low-dimensional projection computed at the bottleneck layer is affected by backpropagating MSEs when nonlinearly reconstructing the relevant QoIs at the output of the decoder. The projection can be explicitly optimized for representing variables required by the ROM; we mix projection-independent variables (selected state variables, $\mathbf{X}^{(S)}$) and projection-dependent variables (the projected source terms, \mathbf{S}_η). This requires an intrusive modification to the neural network training, where at each training step, the current basis matrix, \mathbf{A} , is extracted from the encoder, and \mathbf{S}_η are recomputed. Here, \mathbf{S}_η are the projected source terms, and $\tilde{\mathbf{S}}_\eta$ is their signed square-root transformation (more information is provided in the main text).

(B) We visualize the evolution of one example 2D projection of a reacting flow dataset during neural network training. The overlap initially present in the projection topology is minimized and eventually resolved with a sufficient amount of training. The projection also becomes more spread out in the new coordinate space, reducing gradients in QoIs. Projections are colored by the temperature, T . η denotes the low-dimensional data projection, where η_i are the individual lower-dimensional coordinates.

low-dimensional basis. As long as \mathbf{A} is constant in space and time, the PDEs presented in Equation 1 can be transformed to that new basis:

$$\frac{\partial \eta^\top}{\partial t} = \underbrace{-\nabla \cdot \mathbf{C}_\eta^\top}_{\text{Advection}} - \underbrace{\nabla \cdot \mathbf{D}_\eta^\top}_{\text{Diffusion}} + \underbrace{\mathbf{S}_\eta^\top}_{\text{Source}}, \quad (\text{Equation 2})$$

where $\eta = \mathbf{X} \cdot \mathbf{A}$, $\mathbf{C}_\eta = \mathbf{C} \cdot \mathbf{A}$, $\mathbf{D}_\eta = \mathbf{D} \cdot \mathbf{A}$, and $\mathbf{S}_\eta = \mathbf{S} \cdot \mathbf{A}$. Here, $\eta = [\eta_1, \eta_2, \dots, \eta_q]$ is a matrix defining the low-dimensional data projection and $\mathbf{S}_\eta = [S_{\eta,1}, S_{\eta,2}, \dots, S_{\eta,q}]$ are the q projected source terms. With this transformation, we obtain a much smaller set (q instead of Q) of PDEs compared with Equation 1 that need to be solved numerically. We train a nonlinear regression model, \mathcal{F} , to provide the map $\mathbf{S}_\eta \approx \mathcal{F}(\eta)$. At ROM simulation runtime, \mathbf{S}_η are reconstructed from the current location on a projection using the trained regression model. The success of that reconstruction depends on the topological quality of the data projection, η . The QoI-aware encoder-decoder that we propose in this work addresses this challenge.

The QoI-aware encoder-decoder approach to dimensionality reduction

A schematic illustration of the proposed QoI-aware encoder-decoder approach is presented in Figure 1A. A standard AE uses the same variables at the input and at the output. In this work, we do not restrict ourselves to having the same output as the input. We propose to use the independent state variables at the input layer but use the relevant QoIs at the output layer. The QoIs can be selected as the important state variables,

$\mathbf{X}^{(S)}$, such as temperature or pressure, and any other relevant physical variables, such as density or viscosity. The QoIs can also include any variables required to close the system of the projected governing equations from Equation 2, e.g., the projected source terms, \mathbf{S}_η .

We classify the QoIs as “projection dependent” and “projection independent.” The projection-dependent QoIs are the ones whose very definition depends on the projection operator, \mathbf{A} . These can, for instance, be the projected source terms, \mathbf{S}_η , whose definition changes during training with every change to the basis matrix, \mathbf{A} . This is in contrast to projection-independent QoIs such as $\mathbf{X}^{(S)}$; their definition remains fixed throughout training, regardless of the projection operator.

Due to decoding QoIs that depend on a low-dimensional basis (such as \mathbf{S}_η), we introduce an intrusive modification to the neural network training process. At each training step, the current matrix \mathbf{A} is extracted from the encoder (see Figure 1A) and the projection-dependent QoIs are recomputed as $\mathbf{S}_\eta = \mathbf{S} \cdot \mathbf{A}$. We observe that the MSE loss still converges during training despite this modification. Under “methods”, we provide complete details on the encoder-decoder training. Figures S1 and S2 provide insights into the MSE loss convergence when training with vs. without projection-dependent QoIs.

We use *linear* activation functions in the encoder to obtain a new basis for a straightforward projection of the governing equations (recall Equation 2). But we use *nonlinear* activation functions in the decoder to accurately reconstruct QoIs. Building ROMs becomes less straightforward with nonlinear manifold learning. If the projection operator is not linear or varies in space or time, applying it to transform the original system of PDEs

requires solving nontrivial closure problems. Thus, to the extent possible, a linear projection is desired. We note that the projection dimensionality, q , is specified by the user. The projection becomes optimized for the requested dimensionality q .

Evolution of one example 2D projection of a combustion dataset during neural network training is visualized in Figure 1B. The changing projection is colored by temperature, which is one of the QoIs. The overlap initially present in projection topology generates significant nonuniqueness in the temperature variable and in other important QoIs. The overlap is minimized and eventually resolved during neural network training. The projection also becomes more spread out in the new low-dimensional space defined by η_1 and η_2 , reducing gradients in QoIs. This single illustrative example is representative of changes in a projection topology that we have observed using the QoI-aware encoder-decoder strategy on a variety of datasets.

Quantitative assessment of projection quality

The topological quality of a low-dimensional manifold can be assessed using the recently proposed normalized variance derivative metric.³² The metric scans the data projection for variation in a QoI using Gaussian kernels of different widths. Regions of high variation in a QoI over small widths can be linked to steep gradients or nonuniqueness resulting from an overlap. The recently proposed cost function for characterizing projection quality³³ builds upon the topological metric proposed by Armstrong and Sutherland.³² It distills the topological information to a single number associated with a given projection and a given QoI, giving a compact quantitative measure. The mathematical formulation of the normalized variance derivative metric and the cost function can be found under “methods.”

The cost function is extensively used in this work to assess projections coming from the QoI-aware encoder-decoder approach. Analysis of projection quality can be performed on projections of any dimensionality and from the perspective of any QoI. The cost function becomes particularly powerful when building ROMs. We can assess how well the data projection represents QoIs reconstructed at model runtime, such as the projected source terms. In the remainder of this work, the cost is referred to as \mathcal{L} . High \mathcal{L} comes from high variation in a QoI's values at very small scales on a projection. It is indicative of problems such as overlap, where data points corresponding to very different QoI values can be projected right on top of one another. A smaller \mathcal{L} indicates an improved projection topology. We note that other distance-based metrics for assessing the quality of data projections exist in the literature,^{37–39} but they do not perform the assessment from the perspective of representing various QoIs on a data projection, which is the main focus of this work.

Nonlinear decoding of QoIs improves projection topologies

We first demonstrate the benefits of nonlinear decoding of QoIs on a synthetic 2D dataset generated from a Gaussian bivariate distribution. In Figure 2A, we visualize the synthetic dataset in the original coordinate system defined by x and y . With the dashed line, we mark the 1D subspace, η_1 , that PCA finds on that dataset. In PCA, this subspace is found using information about the original (state) variables, x and y , only. Such subspace

might not be optimal for representing certain QoIs. In Figure 2B, we color the same synthetic dataset by one example QoI, φ , which exhibits a gradient in a direction perpendicular to the main direction of data variation. To represent this particular QoI, a 1D subspace that is perpendicular to the PCA subspace would be ideal, as it would minimize nonuniqueness in representing the QoI. With the dashed line, we mark the 1D subspace, η_1 , that the QoI-aware encoder-decoder finds guided by nonlinear decoding of the QoI.

In Figure 2C, we perform cost function assessments of various 1D subspaces, η_1 , that are at an angle to the original x axis, using the dataset and the QoI from Figure 2B. A smaller cost value, \mathcal{L}_φ , indicates an improved projection topology; a larger cost indicates a more problematic projection that can include overlaps generating nonuniqueness in QoIs and/or steep gradients in QoIs. We find a distinct minimum in \mathcal{L}_φ that happens when η_1 is at an approximately 135° angle to the x axis. This finding agrees with our visualization of the QoI in Figure 2B and with the optimized 1D subspace that the QoI-aware encoder-decoder finds. We mark this range of the “best” 1D subspaces in gray in Figure 2C. For subspaces at an angle outside of the marked range, the cost function detects significant levels of nonuniqueness. The noisy behavior in the graph of \mathcal{L}_φ stems from the data scatter in the synthetic dataset. A small variation in the angle can cause some points to suddenly overlap one another, thereby increasing the \mathcal{L}_φ value.

Finally, we demonstrate the benefits that nonlinear decoding provides as opposed to linear decoding. In Figure 2D, we visualize the same 2D synthetic dataset colored by four new QoIs, φ , each with a different functional form (noted). The first QoI is a function of x only. The three remaining QoIs are functions of both x and y and thus there exists no 1D subspace that can represent these QoIs uniquely. However, there still exists the “best” 1D subspace that represents a given QoI with the smallest amount of nonuniqueness. The goal of our dimensionality reduction strategy is to find that subspace.

In Figure 2E, we generate 1D subspaces at varying angles to the original x axis corresponding to the cases visualized in Figure 2D. For each subspace, we compute the cost, \mathcal{L}_φ , of representing a particular QoI. We measure the final MSE loss of linear vs. nonlinear decoding of QoIs. Costs and MSE losses are visualized in Figure 2E. With the red dashed line, we mark the minimum \mathcal{L}_φ and the minimum MSE loss that happens for each QoI. We note that the minima in \mathcal{L}_φ always coincide with the minima in the MSE loss when using nonlinear decoding and *not* with the minima when using linear decoding. The minima in \mathcal{L}_φ indicate that 1D projections generate the smallest amount of nonuniqueness in each QoI. Unsurprisingly, these 1D projections also lead to the best nonlinear reconstruction performance in terms of the MSE. This is an expected outcome, since nonuniqueness in a QoI cannot be overcome during training a regression model, no matter how accurate the regression technique is otherwise.

Application to reacting flow dynamical system: nonlinear decoding of projection-dependent outputs

We now apply our QoI-aware dimensionality reduction strategy to reacting flow datasets representing the combustion of various fuels in air. The training datasets are generated from numerical




Figure 2. Nonlinear decoding of QoLs promotes improved low-dimensional data representations

(A) A synthetic 2D dataset with x and y representing the original (state) variables. With the dashed line, we mark the 1D subspace, η_1 , that PCA finds on that dataset.

(B) The same synthetic 2D dataset colored by a QoL, ϕ . With the dashed line, we mark the 1D subspace, η_1 , that a QoL-aware encoder-decoder finds. This subspace minimizes nonuniqueness in representing ϕ after projection and is thus an optimal subspace for that particular QoL.

(C) Cost function, \mathcal{L}_ϕ , assessments of 1D subspaces at an angle to the original x axis for the dataset and the QoL from (B). We achieve the minimum \mathcal{L}_ϕ when the 1D subspace is at approximately 135° to the x axis, which is consistent with the visualization in (B).

(D) Various QoLs, ϕ , represented on the same synthetic 2D dataset. We note the exact functional form for each ϕ .

(E) Costs, \mathcal{L}_ϕ , and MSEs for various 1D subspaces at an angle to the x axis corresponding to QoLs visualized in (D). We show the final MSE loss from nonlinear and linear decoding of each of the four QoLs. With the red dashed line, we mark the global minimum in \mathcal{L}_ϕ or MSE. The minimum \mathcal{L}_ϕ always coincides with the minimum MSE using nonlinear decoding and *not* using linear decoding.

simulations of simplified combustion models with varying stoichiometric and flow conditions to cover a wide range of possible thermochemical states of a burning flame (see “[methods](#)” for more information on how the datasets were generated). In a ROM, these states are accessed at specific locations on a low-dimensional projection. At ROM simulation runtime, projection-dependent QoLs are reconstructed from the current location on a projection, e.g., using a trained regression model.

Here, the relevant QoLs are the important state variables, such as temperature and major chemical species, $\mathbf{X}^{(S)}$, and the projected source terms, \mathbf{S}_η , required by the reduced model (recall [Figure 1](#)). For reacting flow datasets, \mathbf{S}_η are highly nonlinear functions of temperature, pressure, and chemical composition. We also include a nonlinear signed square-root transformation (see “[methods](#)”) of the projected source terms, $\tilde{\mathbf{S}}_\eta$, which highlights the importance of states $\mathbf{S}_\eta \approx 0$, important from the physical point of view but easy to miss in a regression model due to the highly nonlinear nature of \mathbf{S}_η .⁴⁰

We find that including projection-dependent QoLs, \mathbf{S}_η and $\tilde{\mathbf{S}}_\eta$, at the output of the decoder yields further improvements in projection topologies. This result can explain the improved ROM performance reported in the literature.^{24–26} In [Figure 3A](#), we demonstrate this using three reacting flow datasets for the com-

bustion of hydrogen, methane, and ethylene in air. We generate 2D and 3D projections of high-dimensional thermochemical state spaces. We highlight, however, that the projection dimensionality, q , can be arbitrarily selected by the user. For each fuel and each dimensionality, we compute 100 projections resulting from a different random initialization of the weights in the neural network (drawn from the Glorot uniform distribution⁴¹). We again quantify the projection quality using the cost function. We visualize probability density functions (PDFs) of averaged costs, \mathcal{L} , across target QoLs. Each PDF is thus computed from 100 values of \mathcal{L} corresponding to 100 different neural network initializations. The PDF for each fuel is plotted with a different line style in [Figure 3A](#). We test two scenarios for the QoLs included at the output of the decoder: (1) selected state variables, $\mathbf{X}^{(S)}$, and the projected source terms, \mathbf{S}_η and $\tilde{\mathbf{S}}_\eta$, and (2) only the selected state variables, $\mathbf{X}^{(S)}$. We note that the mean of the PDFs is generally shifted toward smaller \mathcal{L} when \mathbf{S}_η and $\tilde{\mathbf{S}}_\eta$ are included at the decoder output. For 2D projections of the 32-dimensional state space corresponding to the most complex fuel tested, the hydrocarbon ethylene, 91% of random initializations lead to better projection topologies when the projected source terms are included at the decoder output compared with when they are not included. We have, however, encountered one combustion




Figure 3. Including projection-dependent QoIs at the decoder output promotes improved projection topologies across three combustion datasets for various fuels

(A) Projection qualities are assessed using the cost function, \mathcal{L} . PDFs of averaged \mathcal{L} for each fuel are plotted with a different line style.

(B) Example 2D projections from different ranges in the cost function as tagged in (A). All projections are colored by the corresponding first projected source term, $S_{\eta,1}$.

(C) Correlation between the cost, \mathcal{L} , and the coefficient of determination, R^2 , for kernel regression reconstruction of the two projected source terms for 2D and 3D projections from (A).

dataset for a syngas fuel for which nonlinear decoding of \mathbf{S}_η and $\hat{\mathbf{S}}_\eta$ did not bring further benefits.

Since the thermochemical source terms, \mathbf{S} , are nonlinear functions of state variables, the projected source terms, \mathbf{S}_η , are also nonlinear functions of the state variables. Much different features show up in \mathbf{S}_η represented on a projection compared with the selected state variables, $\mathbf{X}^{(S)}$, represented on a projection. We hypothesize that including \mathbf{S}_η as QoIs at the decoder output highlights those new features during encoder-decoder training. The projection now has to be well behaved with respect to these

additional regions of steep variation in \mathbf{S}_η . This promotes resolving any remaining overlaps that state variables alone could not resolve because they do not exhibit much variation across certain regions on a projection. We highlight that, also in this context, the projection-dependent QoIs are distinctly different from other, projection-independent, variables.

In Figure 3B, we visualize representative 2D projections coming from different ranges in the cost function's PDF for each fuel. The best projections formed when including \mathbf{S}_η and $\hat{\mathbf{S}}_\eta$ at the decoder output are characterized by minimized overlaps and a

more spread-out topology. The offset in the PDFs between cases of including vs. not including \mathbf{S}_η and $\tilde{\mathbf{S}}_\eta$ is enough to resolve significant overlaps on a projection. Finally, we note that there is no guarantee that the basis obtained with the QoI-aware encoder-decoder approach is orthogonal (unlike, e.g., in PCA). We have observed, however, that many of the best projections are associated with nearly orthogonal bases. This indicates that the encoder-decoder dimensionality reduction strategy can be successfully coupled with an orthogonality constraint^{42,43} in future research.

A good-quality low-dimensional system representation facilitates building models on top of a projection. In Figure 3C, we quantify the benefits that improved low-dimensional projections bring in ROMs. We build nonlinear kernel regression models for the projected source terms, \mathbf{S}_η , on top of 2D and 3D projections computed in Figure 3A (more information on the kernel regression model is provided under “methods”). We visualize the correlation between the coefficient of determination, R^2 , and the cost function assessment for each individual $S_{\eta,j}$. For 2D projections, there are two projection-dependent variables, $S_{\eta,1}$ and $S_{\eta,2}$, and for 3D projections, there are three: $S_{\eta,1}$, $S_{\eta,2}$, and $S_{\eta,3}$. We compare two scenarios where we include or do not include projection-dependent QoIs at the decoder output. We note that, generally, regression models perform better at reconstructing the highly nonlinear projected source terms when improved projections are used as regressors. Cases corresponding to including \mathbf{S}_η and $\tilde{\mathbf{S}}_\eta$ at the decoder output tend to be clustered in the region of lowest \mathcal{L} and highest R^2 . This result is of particular relevance in systems without advection and diffusion, where the problem dynamics is evolved solely based on the reconstructed projected source terms (see Equation 2). We are now going to look further into one such system.

The effect of projection topology in reduced-order modeling

We solve a simplified system of an adiabatic, incompressible, 0D reactor. All spatial derivatives vanish for a 0D reactor, and the FOM from Equation 1 simplifies to:

$$\frac{d\mathbf{X}^\top}{dt} = \mathbf{S}^\top. \quad (\text{Equation 3})$$

This set of ordinary differential equations (ODEs) is closed by appropriate initial conditions (ICs), $\mathbf{X}(t = 0)$. We can project Equation 3 onto a lower-dimensional basis, \mathbf{A} :

$$\frac{d\boldsymbol{\eta}^\top}{dt} = \mathbf{S}_\eta^\top, \quad (\text{Equation 4})$$

similar to what was done in Equation 2. We formulate a deep neural network (DNN) regression model, \mathcal{F} , to predict \mathbf{S}_η from the low-dimensional projection, $\boldsymbol{\eta}$, at the current time step. More details on the DNN model are presented under “methods.” The ODEs solved by the numerical solver are effectively:

$$\frac{d\boldsymbol{\eta}^\top}{dt} = \mathcal{F}(\boldsymbol{\eta})^\top. \quad (\text{Equation 5})$$

Evolving q -dimensional parameters is thus described by a set of q ODEs. The transport equations for the low-dimensional parameters are a set of coupled ODEs; the model used for predict-

ing each $S_{\eta,j}$ is a function of all low-dimensional parameters, $\boldsymbol{\eta}$. We use `scipy.odeint` for numeric integration of the system of ODEs.⁴⁴

We generate training data by solving Equation 3 for a variety of ICs that specify different initial stoichiometries in the reactor. Temporal trajectories that form the dataset are visualized in Figure 4A. The mixture initially present in the reactor autoignites within a few milliseconds and shortly then reaches chemical equilibrium (which we denote as the steady state, SS). By varying the IC, we cover a wide range of thermochemical states, which helps broaden the training manifold.

Dimensionality reduction techniques that are indifferent to projection topology can severely affect the ROM. For example, overlapping regions on a projection can cause unphysical behavior at simulation runtime, e.g., prematurely move the solution toward the steady state or push the solution outside of the training manifold. This is conceptually demonstrated in Figures 4B and 4C, where we numerically transport a single trajectory in time according to Equation 5 over an overlapping region on a 2D projection of a combustion dataset. This particular data projection visualized in Figure 4B is generated with PCA and thus without incorporating information about the relevant projection-dependent QoIs. With the black dashed line in Figure 4C, we show the true trajectory that the dynamical system with a given IC should follow. With the blue solid line, we show the path that the ROM takes. In the gray background in Figure 4C, we visualize the 2D projection used for training the DNN regression model. We note that the projection exhibits significant overlap between the nearly horizontal (lighter shade of gray) and the nearly vertical (darker shade of gray) training trajectories. When the ROM solution enters the overlapping region, it is being incorrectly deflected away from the true solution, influenced by the nonunique values of the projected source terms, \mathbf{S}_η . In the zoomed-in box, we show how the ROM solution snaps to the overlapping vertical trajectories on the projection instead of following the trajectories lying underneath. From that point on, the solution continues in a diverging direction. This behavior can lead to large errors in reconstructing physical variables from the transported trajectory. It can also move the ROM solution past the termination point at steady state, causing unphysical results when the simulation ventures into an undefined region of a low-dimensional space.

For comparison, we generate an improved-quality 2D projection of the dataset from Figure 4A using the QoI-aware encoder-decoder. The improved, unique projection is visualized in Figure 4D. In Figure 4E, we show the behavior of the ROM when the improved projection topology is used. We observe that the selected temporal trajectory maintains the correct path along the projection. In Figure 4F, we visualize the predictive performance of a poor ROM vs. an improved ROM. We train a separate DNN model, \mathcal{F}_χ , to predict the thermochemical state variables, \mathbf{X} , from the low-dimensional projection, $\boldsymbol{\eta}$. More details on the DNN model are presented under “methods.” In Figure 4F, we predict the formation of one of the main combustion products, H_2O , during autoignition from the numerically evolved ROMs corresponding to using the projection with overlap from Figure 4B vs. using the improved, unique projection from Figure 4D. We note that a significant error is incurred in predicting the H_2O profile when the ROM is forced to traverse the




Figure 4. The effect of projection topology on a reduced-order model

(A) Dataset generated from the full-order model (FOM) representing autoignition in a 0D reactor. Different temporal trajectories correspond to different initial stoichiometries in the reactor. “IC” denotes initial condition, and “SS” denotes steady state.

(B) 2D projection of the 0D reactor dataset from (A) computed by PCA. This projection exhibits significant overlap affecting temporal trajectories in the vicinity of IC and SS. The projection is colored by the first projected source term, $S_{\eta,1}$.

(C) Behavior of a single time trajectory solved by a ROM when the projection exhibits overlap. The training manifold is visualized as the gray background. In the zoomed-in box, we show how the time trajectory of a ROM snaps to the nearly vertical training trajectories once it reaches the region of overlap. From that point on, the ROM continues in a direction diverging from the true transport direction.

(D) An improved 2D projection of the 0D reactor dataset computed using the QoI-aware encoder-decoder. The projection is colored by the first projected source term, $S_{\eta,1}$.

(E) Behavior of a single time trajectory solved by a ROM using the improved projection. The selected trajectory, evolved using a ROM, follows the true path.

(F) Prediction of H_2O formation during autoignition from the numerically evolved ROMs corresponding to using the projection with overlap from (B) vs. using the improved, unique projection from (D).

nonunique projection. Conversely, we obtain an excellent prediction of the H_2O profile using the ROM solved for the improved projection. While in this work we build separate DNN models, \mathcal{F} and \mathcal{F}_X , we note that one could use the already trained decoder to reconstruct \mathbf{S}_η at ROM runtime and to predict the most important state variables, $\mathbf{X}^{(S)}$.

Finally, we note that a steep gradient in a QoI can also be difficult to model with nonlinear regression. Not every available regression technique is able to adapt to rapid changes in a regressed variable. Steep gradients can occur in a model due to the large compression of certain regions of a projection in a sub-optimal dimensionality reduction. It can also occur for highly nonlinear QoIs whose large absolute values are present only on small portions of a projection. In the context of reacting flows, which exhibit a highly multiscale structure, these QoIs can, for example, be minor species, present only in a small region of a projection associated with their formation and short lifetime.⁴⁵ In the autoignition problem tackled here, mispredictions of steep gradients in the projected source terms near the IC can cause mispredictions of the ignition time (recall Figure 4A). The error from a regression model incurred at the start of the simulation propagates throughout the entire transported trajectory. The simulation originating from the IC mispredicts the ignition time in a reactor due to a poor prediction of \mathbf{S}_η at the start of the simulation. What follows is a shift along the time domain in the *a posteriori* reconstruction of the thermochemical state variables.

DISCUSSION

We argue that if nonlinear regression models are required in the downstream use of reduced-order modeling, we may want to pay attention to aspects such as reducing nonuniqueness and smoothing out gradients in QoIs on a low-dimensional data projection. We demonstrate that a dimensionality reduction strategy that uses nonlinear decoding of QoIs to inform the low-dimensional projection can enhance ROMs across various disciplines. Improved projection topologies naturally emerge during training of the proposed QoI-aware encoder-decoder. We show the benefits of including projection-dependent QoIs, in addition to projection-independent QoIs, at the output of a decoder. Our results are relevant to many scientific and engineering disciplines that construct ROMs.^{46–50} Our findings can also be applicable in the broader field of representation learning.^{4,51–55}

We now discuss the bigger picture of the current trends in reduced-order modeling and latent space data representations for understanding and modeling high-dimensional datasets. Latent space data representations are abundantly used, e.g., in fluid dynamics, atmospheric physics, climate science, computational neuroscience, or single-cell omics. We foresee a number of potential multidisciplinary applications for the QoI-aware encoder-decoder approach proposed in this work. Below, we provide avenues for future work where the QoI-aware dimensionality reduction strategy can become useful.

Dimensionality reduction for categorical data

Future research can focus on incorporating class label information as the relevant QoI for categorical datasets to improve data projections and reduce class overlap. This idea falls within the family of supervised dimensionality reduction^{56–58} and can be used for efficient data visualization^{3,58–60} given user-defined labels. For datasets where class labels exhibit a hierarchical, multiscale structure, such as single-cell data,^{61,62} the exact label values become meaningful when training the QoI-aware encoder-decoder. For example, during training that uses the MSE loss, there is a larger penalty for placing classes “1” and “9” close to each other in the latent space compared with placing classes “1” and “2” close to each other (because $1/n\sum_{i=1}^n(1-9)^2 > 1/n\sum_{i=1}^n(1-2)^2$ for n overlapping samples). Therefore, providing a meaningful class hierarchy at the decoder output can promote projections where that hierarchy is preserved in the latent space. It is also expected that the QoI-aware strategy will minimize class overlap in the latent space by penalizing nonuniqueness in representing a QoI.

From dimensionality reduction to manifold learning using nonlinear encoders

In this work, we focus entirely on *linear* data projections. While linear dimensionality reduction may be sufficient for some classes of problems,^{58,63} future research can also explore a fully nonlinear QoI-aware encoder-decoder. This can be achieved by using nonlinear activation, a deep architecture in the encoder, and possibly other constraints.⁴² Nonlinear encoding can have applications in efficient data visualization and hypothesis generation based on latent space representations.^{64,65} Nonlinear data transformations to latent space are still underdeveloped for building ROMs due to challenges in formulating closure for transformed PDEs, although models for simple dynamical systems have recently been formulated.⁶⁶

Toward reduced-order model adaptivity

Model adaptivity has been noted as a desired characteristic that ROMs should possess.⁶⁷ This might mean, in practice, being able to increase prediction accuracy in a region or for a variable of interest, even with the trade-off of losing accuracy in a region (or for a variable) where high accuracy is not required. A commonly invoked characteristic of data-driven ROMs that steers the model accuracy is the projection dimensionality (or, equivalently, the number of transported low-dimensional parameters, q). In PCA, this is a straightforward parameter to use, since principal components are ordered; increasing q will always improve the linear reconstruction errors of the original state variables. However, when both state variables and other dependent variables (such as the projected source terms) are to be nonlinearly regressed based on the q low-dimensional parameters, the quality of the parameterization becomes another important factor.

Strategies for improving projection quality developed in this work can be utilized in the context of model adaptivity. For example, we can target a good reconstruction of specific state variables by calibrating the projection topology directly from the perspective of that variable. Moreover, the recently proposed metrics for assessing the quality of data projections^{32,33} allow us to perform assessments of isolated regions of the projection.

This can, for instance, include regions corresponding to certain important physical phenomena that we would like to capture particularly well in a ROM. In the context of reacting flows, these regions can include ignition, extinction, or reaction zones and can be isolated with data clustering techniques.⁶⁸ We can then focus on improving those local regions on a projection, for example, by reducing overlap, data compression, or curvature. In the context of model adaptivity, we may also be willing to sacrifice the orthogonality of a low-dimensional basis in favor of an improved projection topology.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Kamila Zdybal (kamila.zdybal@gmail.com).

Materials availability

This study did not generate new unique reagents or materials.

Data and code availability

Original code has been deposited at <https://github.com/kamilazdybal/nonlinear-decoding> and is publicly available as of the date of publication. Datasets from numerical simulations of reacting flow for various fuels are available in the same public repository. Code and datasets are archived on Zenodo: <https://doi.org/10.5281/zenodo.8319055>.⁶⁹ To enable researchers to apply the QoI-aware encoder-decoder approach to their own datasets, we have implemented all of the relevant functionalities in our open-source Python library, *PCAfold*.^{70,71} The user can easily implement the QoI-aware encoder-decoder by instantiating an object of the `QoIAwareProjection` class from the `utilities` module. The *PCAfold* library is required to reproduce our results. More information can be found in the documentation: <https://pcafold.readthedocs.io>. Figures presented in this paper use sequential color maps designed by Cramer et al.⁷² The neural network schematic in Figure 1A was drawn using the NN-SVG tool.⁷³

Methods

Reacting flow datasets

Reacting flow datasets used in this work were generated using the *Spitfire* Python library.⁷⁴ We numerically solved the steady laminar flamelet model,⁷⁵ which provides a simplified description of a reacting flow system. The steady laminar flamelet model allows for varying the mixture stoichiometry, f , and the strain rate, χ , imposed on the flow conditions. The mixture stoichiometry specifies the ratio at which the fuel and oxidizer are mixed. The strain rate mimics the effect that turbulence has in a chemically reacting flow. For example, a large strain acts to extinguish the flame. By varying the two model parameters, f and χ , we are able to cover a wide range of possible thermochemical states of a burning flame, from equilibrium toward steady extinction and from a pure oxidizer stream toward a pure fuel stream. Those states are collected into a state-space matrix, $\mathbf{X} \in \mathbb{R}^{(N \times Q)}$, defined as:

$$\mathbf{X} = \begin{bmatrix} \vdots & \vdots & \vdots & \vdots & \vdots \\ T & Y_1 & Y_2 & \dots & Y_{n_s-1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix}, \quad (\text{Equation 6})$$

where T is temperature and Y_i is a mass fraction of the i th chemical species. Only $n_s - 1$ chemical species are included in the matrix, since mass fractions satisfy the relationship $\sum_{i=1}^{n_s} Y_i = 1$. Thus, of n_s chemical species mass fractions, only $n_s - 1$ are independent. The corresponding matrix of thermochemical source terms, $\mathbf{S} \in \mathbb{R}^{(N \times Q)}$, is defined as:

$$\mathbf{S} = \begin{bmatrix} \vdots & \vdots & \vdots & \vdots & \vdots \\ -\frac{1}{\rho c_p} \sum_{i=1}^{n_s} h_i \omega_i & \frac{\omega_1}{\rho} & \frac{\omega_2}{\rho} & \dots & \frac{\omega_{n_s-1}}{\rho} \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix}, \quad (\text{Equation 7})$$

where ρ is mixture density, c_p is mixture specific heat capacity, h_i is the enthalpy of species i , and ω_j is the net mass production rate of species i . Matrices \mathbf{X} and \mathbf{S} are high dimensional. Below, we note the sizes of those matrices for each dataset:

- hydrogen/air: $N = 58,101$ and $Q = 9$;
- methane/air: $N = 54,000$ and $Q = 52$;
- ethylene/air: $N = 60,000$ and $Q = 32$.

The number N denotes the number of observations in a dataset and can be thought of as the number of grid points. In a steady laminar flamelet model, the grid points are not related to physical space. Instead, each observation can be linked to a specific stoichiometry and a specific strain rate, the two model variables that act as proxies for a spatial location in a turbulent reacting flow. The number Q denotes system dimensionality.

In Figure 4, we use a 0D reactor dataset also generated using the *Spitfire* Python library.⁷⁴ The structure of the state space matrix is the same as in Equation 6 and the structure of the thermochemical source terms matrix is the same as in Equation 7. We use the combustion of syngas in air. The dataset has $N = 68,140$ and $Q = 11$.

Training the QoI-aware encoder-decoder

Here, we provide details on how the QoI-aware encoder-decoder neural network was built and trained on the datasets presented in the main text. We used the *Keras* Python library⁷⁶ to set up the DNN architecture. We used the Adam optimizer with an initial learning rate of 0.001 and an MSE loss function. We used a linear activation function in the encoder and hyperbolic tangent activation functions in the decoder. The only exception was for generating results from Figure 2E, where linear decoding was accomplished by using a linear activation function in the entire encoder-decoder. The weights were initialized in the neural network using the Glorot uniform distribution⁴¹ and the biases were initialized to zeros. The inputs to the neural network were scaled to a $(0, 1)$ range.⁷⁷ The outputs of the neural network were always scaled to a $(-1, 1)$ range to match the hyperbolic tangent output range. For results reproducibility, we used fixed random seeds for network initialization. The exact values for random seeds can be retrieved from the code that we provide in the public GitHub repository.

We did not train using mini-batches, following the premise that at any given time during training we should be evaluating the MSE loss from the entire data projection and not from small samples (batches) of it. Evaluating predictions from a sampled projection can miss the severity of nonuniqueness or can underresolve steep gradients in a QoI. We have observed that training on small batches of data instead of a full dataset at once can lead to suboptimal projections. We thus used the entire 90% of the data observations (N) at each training step. The remaining 10% was used as validation data. For reacting flow datasets, we trained the network for 20,000 epochs. For a much simpler synthetic dataset in Figure 2 we used 1,000 epochs. The optimized basis matrix, \mathbf{A} , was retrieved once the MSE loss reached convergence, and it always corresponded to the epoch at which the MSE loss was smallest. This follows from the argument that if the projection topology is well behaved, nonlinear regression performed by the decoder should lead to small reconstruction errors. The basis matrix was the weights from the encoding layer, normalized to yield unit length. The basis matrix had q columns that denoted projection dimensionality (or, equivalently, the number of neurons in the bottleneck layer).

For 2D projections, we used the following network architectures for each dataset:

- hydrogen/air: 9-2-6-9-10-10;
- methane/air: 52-2-6-9-10-10;
- ethylene/air: 32-2-6-9-10-10.

For 3D projections, we used the following network architectures for each dataset:

- hydrogen/air: 9-3-7-10-11-12;
- methane/air: 52-3-7-10-11-12;
- ethylene/air: 32-3-7-10-11-12.

These were the architectures when the projection-dependent QoIs were included at the decoder output. When projection-dependent QoIs were not

included at the decoder output, only the number of neurons at the last decoding layer changed, and we used the same interior architecture.

The signed square-root transformation of each individual projected source term, S_{η_j} , was computed as per Armstrong⁴⁰ in the following way:

$$\tilde{S}_{\eta_j} = \text{sign}(S_{\eta_j} + 10^{-4}) \cdot \sqrt{|S_{\eta_j} + 10^{-4}|}. \quad (\text{Equation 8})$$

This is a way of stretching values $S_{\eta_j} \approx 0$, maintaining sign changes, and reducing the overall variation in the projected source terms. We call $\tilde{\mathbf{S}}_{\eta} = [\tilde{S}_{\eta,1}, \tilde{S}_{\eta,2}, \dots, \tilde{S}_{\eta,q}]$.

Cost function evaluations

The recently proposed normalized variance metric can be used to assess the quality of a low-dimensional manifold.³² The data projection was first scaled to a unit box such that each $\eta_i \in (0, 1)$. Next, the projection was scanned at length scales given by $\sigma \in (\sigma_{min}, \sigma_{max})$, for any variation in a QoI's values. The normalized variance, $\mathcal{N}(\sigma)$, was computed for a single QoI, φ , as:

$$\mathcal{N}(\sigma) = \frac{\sum_{k=1}^N (\varphi_k - \mathcal{K}_{\varphi}(\eta_k, \sigma))^2}{\sum_{k=1}^N (\varphi_k - \bar{\varphi})^2}, \quad (\text{Equation 9})$$

where N is the number of observations in a dataset, $\bar{\varphi}$ is an arithmetic average of φ , and \mathcal{K} is a weighted average of observations of φ computed as:

$$\mathcal{K}_{\varphi}(\eta, \sigma) = \frac{\sum_{j=1}^N W_j \varphi_j}{\sum_{j=1}^N W_j}, \quad (\text{Equation 10})$$

where the weights, w_j , are determined using a Gaussian kernel. We then computed a derivative of $\mathcal{N}(\sigma)$ as per Armstrong and Sutherland³²:

$$D(\sigma) = \frac{d\mathcal{N}(\sigma)}{d\log_{10}(\sigma)} + \lim_{\sigma \rightarrow 0} \mathcal{N}(\sigma), \quad (\text{Equation 11})$$

and normalized it by its maximum value:

$$\hat{D}(\sigma) = \frac{D(\sigma)}{\max(D(\sigma))}. \quad (\text{Equation 12})$$

The normalized variance derivative, $\hat{D}(\sigma)$, provides a graph for each assessed QoI that can be analyzed visually. We took it one step further and turned this graph into a single number that provides a "cost" of representing a QoI on a projection. The cost function for low-dimensional manifold topology assessment³³ is therefore defined as:

$$\mathcal{L}_{\varphi} = \int_{\bar{\sigma}_{min}}^{\bar{\sigma}_{max}} P(\sigma, \sigma_{peak}) \cdot \hat{D}(\sigma) d\bar{\sigma}, \quad (\text{Equation 13})$$

where $P(\sigma, \sigma_{peak})$ is the penalty function defined as:

$$P(\sigma, \sigma_{peak}) = |\bar{\sigma} - \bar{\sigma}_{peak}|^r + b \cdot \frac{\bar{\sigma}_{max} - \bar{\sigma}_{min}}{\bar{\sigma}_{peak} - \bar{\sigma}_{min}}, \quad (\text{Equation 14})$$

and σ_{peak} is the length scale that determines the largest-scale variation of a QoI on a projection. Tilde over σ denotes a log-transformed quantity. We performed log transformation of length scales to give equal importance to scales that were orders of magnitude smaller than the characteristic length scale of the projection. Throughout this work, we performed assessments of projection quality taking $\sigma \in (10^{-7}, 10^3)$ in order to span a wide range of length scales at which a QoI could vary on a (normalized) projection.

The cost function has two hyperparameters, as seen in Equation 14: power, r , and vertical shift, b . Increasing the power allows for a higher penalty for projections that exhibit nonuniqueness, while increasing the vertical shift allows for a higher penalty for projections that exhibit steep gradients. Steep gradients in a QoI can be modeled with a well-trained regression model, but nonuniqueness cannot be overcome even with the best regression models available in the research community. We thus used power $r = 4$ and vertical shift $b = 1$ to better discriminate between unique and nonunique projections while maintaining a moderate penalty for small feature sizes in QoIs.

The averaged cost value as reported in Figure 3 was computed as the root-mean-squared norm over the individual costs for the relevant QoIs. The

relevant QoIs taken for each fuel included temperature; mass fractions of major chemical species (fuel, oxidizer, and major combustion products) and selected minor chemical species; the projected thermochemical source terms, \mathbf{S}_p ; and the nonlinearly transformed projected source terms. For each fuel, these relevant QoIs were the same as variables included at the output of a decoder. For 2D projections, there were 10 QoIs, and for 3D projections, there were 12 QoIs.

Kernel regression model

In Figure 3C, we used Nadaraya-Watson kernel regression to predict the projected source terms, \mathbf{S}_p , from the 2D and 3D low-dimensional projections. Predictions were made as a weighted average of the neighboring observations. Weights were established using a Gaussian kernel with an isotropic bandwidth. The bandwidth was adapted locally based on the 100 nearest neighbors.

Deep neural network model for \mathbf{S}_p

In Figures 4B–4E, a DNN was trained as a regression model, \mathcal{F} , to close the system of projected PDEs in Equation 5. Here, we provide details on how this model was set up. We used a DNN with a 2-5-10-15-10-5-2 architecture and hyperbolic tangent activation functions. We used the Adam optimizer with an MSE loss and a learning rate of 0.001. We trained the model for 1,500 epochs. The weights were initialized in the neural network using the Glorot uniform distribution,⁴¹ and the biases were initialized to zeros. The input to the neural network was the data projection, η , scaled to a $(-1, 1)$ range. The outputs of the neural network were the projected source terms, \mathbf{S}_p , scaled to a $(-1, 1)$ range to match the hyperbolic tangent output range. For results reproducibility, we used a fixed random seed of 100 for network initialization.

Deep neural network model for \mathbf{X}

In Figure 4F, a separate DNN was trained as a regression model, \mathcal{F}_X , to predict the thermochemical state variables from the evolved ROM solution. Here, we provide details on how this model was set up. We used a DNN with a 2-5-10-10-10-11 architecture. We used hyperbolic tangent activation functions in the first three layers and sigmoid activation functions in the last two layers. We used the Adam optimizer with an MSE loss and a learning rate of 0.001. We trained the model for 1,500 epochs. The weights were initialized in the neural network using the Glorot uniform distribution,⁴¹ and the biases were initialized to zeros. The input to the neural network was the data projection, η , scaled to a $(-1, 1)$ range. The outputs of the neural network were the thermochemical state variables, \mathbf{X} , scaled to a $(0, 1)$ range to match the sigmoid output range. For results reproducibility, we used a fixed random seed of 100 for network initialization.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.patter.2023.100859>.

ACKNOWLEDGMENTS

The authors would like to thank R. Malpica Galassi and J. Phillips for helpful discussions. The research of K.Z. is supported by an F.R.S.-FNRS Aspirant Research Fellow grant. A.P. has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement 714605. J.C.S. acknowledges the support from the National Science Foundation under grant 1953350.

AUTHOR CONTRIBUTIONS

K.Z. and J.C.S. developed the idea. K.Z. developed the code and performed simulations. K.Z., A.P., and J.C.S. designed the research, analyzed the results, and wrote the paper. A.P. and J.C.S. supervised the work.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: June 9, 2023

Revised: July 12, 2023

Accepted: September 14, 2023

Published: October 10, 2023

REFERENCES

- Allen, E.A., Erhardt, E.B., and Calhoun, V.D. (2012). Data visualization in the neurosciences: Overcoming the curse of dimensionality. *Neuron* 74, 603–608. <https://doi.org/10.1016/j.neuron.2012.05.001>.
- Goldstone, R.L., Pestilli, F., and Böerner, K. (2015). Self-portraits of the brain: Cognitive science, data visualization, and communicating brain structure and function. *Trends Cognit. Sci.* 19, 462–474. <https://doi.org/10.1016/j.tics.2015.05.012>.
- Moon, K.R., van Dijk, D., Wang, Z., Gigante, S., Burkhardt, D.B., Chen, W.S., Yim, K., Elzen, A.v. d., Hirn, M.J., Coifman, R.R., et al. (2019). Visualizing structure and transitions in high-dimensional biological data. *Nat. Biotechnol.* 37, 1482–1492. <https://doi.org/10.1038/s41587-019-0336-3>.
- Kopf, A., and Claassen, M. (2021). Latent representation learning in biology and translational medicine. *Patterns* 2, 100198. <https://doi.org/10.1016/j.patter.2021.100198>.
- Eivazi, H., Le Clairche, S., Hoyas, S., and Vinuesa, R. (2022). Towards extraction of orthogonal and parsimonious non-linear modes from turbulent flows. *Expert Syst. Appl.* 202, 117038. <https://doi.org/10.1016/j.eswa.2022.117038>.
- Berkooz, G., Holmes, P., and Lumley, J.L. (1993). The proper orthogonal decomposition in the analysis of turbulent flows. *Annu. Rev. Fluid Mech.* 25, 539–575. <https://doi.org/10.1146/annurev.fl.25.010193.002543>.
- Sutherland, J.C., and Parente, A. (2009). Combustion modeling using principal component analysis. *Proc. Combust. Inst.* 32, 1563–1570. <https://doi.org/10.1016/j.proci.2008.06.147>.
- Boulakia, M., Schenone, E., and Gerbeau, J.-F. (2012). Reduced-order modeling for cardiac electrophysiology. Application to parameter identification. *Int. J. Numer. Method. Biomed. Eng.* 28, 727–744. <https://doi.org/10.1002/cnm.2465>.
- Lehtimäki, M., Paunonen, L., and Linne, M.-L. (2022). Accelerating neural ODEs using model order reduction. *IEEE Transact. Neural Networks Learn. Syst.* 1–13. <https://doi.org/10.1109/TNNLS.2022.3175757>.
- Parente, A., and Sutherland, J.C. (2013). Principal component analysis of turbulent combustion data: Data pre-processing and manifold sensitivity. *Combust. Flame* 160, 340–350. <https://doi.org/10.1016/j.combustflame.2012.09.016>.
- Biglari, A., and Sutherland, J.C. (2015). An a-posteriori evaluation of principal component analysis-based models for turbulent combustion simulations. *Combust. Flame* 162, 4025–4035. <https://doi.org/10.1016/j.combustflame.2015.07.042>.
- Malik, M.R., Isaac, B.J., Coussement, A., Smith, P.J., and Parente, A. (2018). Principal component analysis coupled with nonlinear regression for chemistry reduction. *Combust. Flame* 187, 30–41. <https://doi.org/10.1016/j.combustflame.2017.08.012>.
- Malik, M.R., Coussement, A., Echehki, T., and Parente, A. (2022). Principal component analysis based combustion model in the context of a lifted methane/air flame: Sensitivity to the manifold parameters and subgrid closure. *Combust. Flame* 244, 112134. <https://doi.org/10.1016/j.combustflame.2022.112134>.
- Benner, P., Gugercin, S., and Willcox, K. (2015). A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Rev.* 57, 483–531. <https://doi.org/10.1137/130932715>.
- Mendez, M.A., Balabane, M., and Buchlin, J.-M. (2019). Multi-scale proper orthogonal decomposition of complex fluid flows. *J. Fluid Mech.* 870, 988–1036. <https://doi.org/10.1017/jfm.2019.212>.
- Kadeethum, T., Ballarin, F., O'Malley, D., Choi, Y., Bouklas, N., and Yoon, H. (2022). Reduced order modeling for flow and transport problems with Barlow twins self-supervised learning. *Sci. Rep.* 12, 20654. <https://doi.org/10.1038/s41598-022-24545-3>.
- Recanatani, S., Bradde, S., Balasubramanian, V., Steinmetz, N.A., and Shea-Brown, E. (2022). A scale-dependent measure of system dimensionality. *Patterns* 3, 100555. <https://doi.org/10.1016/j.patter.2022.100555>.

18. Dutta, S., Rivera-Casillas, P., Styles, B., and Farthing, M.W. (2022). Reduced order modeling using advection-aware autoencoders. *Math. Comput. Appl.* 27, 34. <https://doi.org/10.3390/mca27030034>.
19. Mendez, M.A., Ianiro, A., Noack, B.R., and Brunton, S.L. (2023). *Data-Driven Fluid Mechanics: Combining First Principles and Machine Learning* (Cambridge University Press). <https://doi.org/10.1017/9781108896214>.
20. Hesthaven, J.S., Rozza, G., Stamm, B., et al. (2016). *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, 590 (Springer). <https://doi.org/10.1007/978-3-319-22470-1>.
21. Hesthaven, J.S., and Ubbiali, S. (2018). Non-intrusive reduced order modeling of nonlinear problems using neural networks. *J. Comput. Phys.* 363, 55–78. <https://doi.org/10.1016/j.jcp.2018.02.037>.
22. Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986). Learning representations by back-propagating errors. *Nature* 323, 533–536. <https://doi.org/10.1038/323533a0>.
23. Mehlig, B. (2019). *Machine Learning with Neural Networks* (Cambridge University Press). Chapter 10. <https://doi.org/10.1017/9781108860604>.
24. Perry, B.A., Henry de Frahan, M.T., and Yellapantula, S. (2022). Co-optimized machine-learned manifold models for large eddy simulation of turbulent combustion. *Combust. Flame* 244, 112286. <https://doi.org/10.1016/j.combustflame.2022.112286>.
25. Scherding, C., Rigas, G., Sipp, D., Schmid, P.J., and Sayadi, T. (2023). Data-driven framework for input/output lookup tables reduction: Application to hypersonic flows in chemical nonequilibrium. *Phys. Rev. Fluids* 8, 023201. <https://doi.org/10.1103/PhysRevFluids.8.023201>.
26. Salunkhe, A., Deighan, D., DesJardin, P.E., and Chandola, V. (2023). Physics informed machine learning for chemistry tabulation. *J. Comput. Sci.* 69, 102001. <https://doi.org/10.1016/j.jocs.2023.102001>.
27. Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., and Battaglia, P.W. (2020). Learning mesh-based simulation with graph networks. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2010.03409>.
28. Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., and Battaglia, P. (2020). Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning (PMLR)*, pp. 8459–8468. <https://doi.org/10.5555/3524938.3525722>.
29. Cheng, S., Jin, Y., Harrison, S.P., Quilodr n-Casas, C., Prentice, I.C., Guo, Y.-K., and Arcucci, R. (2022). Parameter flexible wildfire prediction using machine learning techniques: Forward and inverse modelling. *Rem. Sens.* 14, 3228. <https://doi.org/10.3390/rs14133228>.
30. Jacobsen, C., and Duraisamy, K. (2022). Disentangling generative factors of physical fields using variational autoencoders. *Front. Physiol.* 10, 890910. <https://doi.org/10.3389/fphys.2022.890910>.
31. Cheng, S., Chen, J., Anastasiou, C., Angeli, P., Matar, O.K., Guo, Y.-K., Pain, C.C., and Arcucci, R. (2023). Generalised latent assimilation in heterogeneous reduced spaces with machine learning surrogate models. *J. Sci. Comput.* 94, 11. <https://doi.org/10.1007/s10915-022-02059-4>.
32. Armstrong, E., and Sutherland, J.C. (2021). A technique for characterising feature size and quality of manifolds. *Combust. Theor. Model.* 25, 646–668. <https://doi.org/10.1080/13647830.2021.1931715>.
33. Zdybal, K., Armstrong, E., Sutherland, J.C., and Parente, A. (2022). Cost function for low-dimensional manifold topology assessment. *Sci. Rep.* 12, 14496. <https://doi.org/10.1038/s41598-022-18655-1>.
34. Baldi, P., and Hornik, K. (1989). Neural networks and principal component analysis: Learning from examples without local minima. *Neural Network* 2, 53–58. [https://doi.org/10.1016/0893-6080\(89\)90014-2](https://doi.org/10.1016/0893-6080(89)90014-2).
35. Zdybal, K., Sutherland, J.C., and Parente, A. (2023). Manifold-informed state vector subset for reduced-order modeling. *Proc. Combust. Inst.* 39, 5145–5154. <https://doi.org/10.1016/j.proci.2022.06.019>.
36. Dong, X., and Bacher, R. (2022). Data-driven assessment of dimension reduction quality for single-cell omics data. *Patterns* 3, 100465. <https://doi.org/10.1016/j.patter.2022.100465>.
37. Zhang, Y., Shang, Q., and Zhang, G. (2021). pyDRMetrics — A Python toolkit for dimensionality reduction quality assessment. *Heliyon* 7, e06199. <https://doi.org/10.1016/j.heliyon.2021.e06199>.
38. Glielmo, A., Macocco, I., Doimo, D., Carli, M., Zeni, C., Wild, R., d’Errico, M., Rodriguez, A., and Laio, A. (2022). DADAPy: Distance-based analysis of data-manifolds in Python. *Patterns* 3, 100589. <https://doi.org/10.1016/j.patter.2022.100589>.
39. Johnson, E.M., Kath, W., and Mani, M. (2022). EMBEDR: distinguishing signal from noise in single-cell omics data. *Patterns* 3, 100443. <https://doi.org/10.1016/j.patter.2022.100443>.
40. Armstrong, E. (2023). *Development of Improved Parameterizations and Nonlinear Regression for Reduced-Order Modeling in Combustion* (Department of Chemical Engineering, The University of Utah). Ph.D. thesis.
41. Glorot, X., and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, JMLR Workshop and Conference Proceedings*, pp. 249–256. <https://proceedings.mlr.press/v9/glorot10a.html>.
42. Ladjal, S., Newson, A., and Pham, C.-H. (2019). A PCA-like autoencoder. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1904.01277>.
43. Chen, Y., Wen, Y., Xie, C., Chen, X., He, S., Bo, X., and Zhang, Z. (2023). MOCSS: Multi-omics data clustering and cancer subtyping via shared and specific representation learning. *iScience*. <https://doi.org/10.1016/j.isci.2023.107378>.
44. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat. Methods* 17, 261–272. <https://doi.org/10.1038/s41592-019-0686-2>.
45. Pr fert, U., Hartl, S., Hunger, F., Messig, D., Eiermann, M., and Hasse, C. (2015). A constrained control approach for the automated choice of an optimal progress variable for chemistry tabulation. *Flow, Turbul. Combust.* 94, 593–617. <https://doi.org/10.1007/s10494-015-9595-3>.
46. Umakantha, A., Morina, R., Cowley, B.R., Snyder, A.C., Smith, M.A., and Yu, B.M. (2021). Bridging neuronal correlations and dimensionality reduction. *Neuron* 109, 2740–2754.e12. <https://doi.org/10.1016/j.neuron.2021.06.028>.
47. Iyer, K.K., Hwang, K., Hearne, L.J., Muller, E., D’Esposito, M., Shine, J.M., and Cocchi, L. (2022). Focal neural perturbations reshape low-dimensional trajectories of brain activity supporting cognitive performance. *Nat. Commun.* 13, 4. <https://doi.org/10.1038/s41467-021-26978-2>.
48. Malik, M.R., Khamedov, R., P rez, F.E.H., Coussemant, A., Parente, A., and Im, H.G. (2022). Dimensionality reduction and unsupervised classification for high-fidelity reacting flow simulations. *Proc. Combust. Inst.* <https://doi.org/10.1016/j.proci.2022.06.017>.
49. Jia, S., Li, X., Huang, T., Liu, J.K., and Yu, Z. (2022). Representing the dynamics of high-dimensional data with non-redundant wavelets. *Patterns* 3, 100424. <https://doi.org/10.1016/j.patter.2021.100424>.
50. Rustam, Gunawan, A.Y., and Kresnowati, M.T.A.P. (2022). Data dimensionality reduction technique for clustering problem of metabolomics data. *Heliyon* 8, e09715. <https://doi.org/10.1016/j.heliyon.2022.e09715>.
51. Recanatesi, S., Farrell, M., Lajoie, G., Deneve, S., Rigotti, M., and Shea-Brown, E. (2021). Predictive learning as a network mechanism for extracting low-dimensional latent space representations. *Nat. Commun.* 12, 1417. <https://doi.org/10.1038/s41467-021-21696-1>.
52. Zeng, X., Xiang, H., Yu, L., Wang, J., Li, K., Nussinov, R., and Cheng, F. (2022). Accurate prediction of molecular properties and drug targets using a self-supervised image representation learning framework. *Nat. Mach. Intell.* 4, 1004–1016. <https://doi.org/10.1038/s42256-022-00557-6>.
53. Ma, R., Sun, E.D., and Zou, J. (2023). A spectral method for assessing and combining multiple data visualizations. *Nat. Commun.* 14, 780. <https://doi.org/10.1038/s41467-023-36492-2>.
54. Wang, X., Cheng, Y., Yang, Y., Yu, Y., Li, F., and Peng, S. (2023). Multitask joint strategies of self-supervised representation learning on biomedical networks for drug discovery. *Nat. Mach. Intell.* 5, 445–456. <https://doi.org/10.1038/s42256-023-00640-6>.

55. Johnston, W.J., and Fusi, S. (2023). Abstract representations emerge naturally in neural networks trained to perform multiple tasks. *Nat. Commun.* *14*, 1040. <https://doi.org/10.1038/s41467-023-36583-0>.
56. Paiva, J.G.S., Schwartz, W.R., Pedrini, H., and Minghim, R. (2012). Semi-supervised dimensionality reduction based on partial least squares for visual analysis of high dimensional data. Wiley Online Library. *Comput. Graph. Forum* *31*, 1345–1354. <https://doi.org/10.1111/j.1467-8659.2012.03126.x>.
57. Vogelstein, J.T., Bridgeford, E.W., Tang, M., Zheng, D., Douville, C., Burns, R., and Maggioni, M. (2021). Supervised dimensionality reduction for big data. *Nat. Commun.* *12*, 2872. <https://doi.org/10.1038/s41467-021-23102-2>.
58. Amouzgar, M., Glass, D.R., Baskar, R., Averbukh, I., Kimmey, S.C., Tsai, A.G., Hartmann, F.J., and Bendall, S.C. (2022). Supervised dimensionality reduction for exploration of single-cell data by HSS-LDA. *Patterns* *3*, 100536. <https://doi.org/10.1016/j.patter.2022.100536>.
59. Becht, E., McInnes, L., Healy, J., Dutertre, C.-A., Kwok, I.W.H., Ng, L.G., Ginhoux, F., and Newell, E.W. (2018). Dimensionality reduction for visualizing single-cell data using UMAP. *Nat. Biotechnol.* *37*, 38–44. <https://doi.org/10.1038/nbt.4314>.
60. Etheredge, R.I., Schartl, M., and Jordan, A. (2021). Decontextualized learning for interpretable hierarchical representations of visual patterns. *Patterns* *2*, 100193. <https://doi.org/10.1016/j.patter.2020.100193>.
61. Kobak, D., and Berens, P. (2019). The art of using t-SNE for single-cell transcriptomics. *Nat. Commun.* *10*, 5416. <https://doi.org/10.1038/s41467-019-13056-x>.
62. Han, L., Wei, X., Liu, C., Volpe, G., Zhuang, Z., Zou, X., Wang, Z., Pan, T., Yuan, Y., Zhang, X., et al. (2022). Cell transcriptomic atlas of the non-human primate *Macaca fascicularis*. *Nature* *604*, 723–731. <https://doi.org/10.1038/s41586-022-04587-3>.
63. Christodoulou, E., Ma, J., Collins, G.S., Steyerberg, E.W., Verbakel, J.Y., and Van Calster, B. (2019). A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *J. Clin. Epidemiol.* *110*, 12–22. <https://doi.org/10.1016/j.jclinepi.2019.02.004>.
64. Ji, Y., Lotfollahi, M., Wolf, F.A., and Theis, F.J. (2021). Machine learning for perturbational single-cell omics. *Cell Syst.* *12*, 522–537. <https://doi.org/10.1016/j.cels.2021.05.016>.
65. Perl, Y.S., Pallavicini, C., Piccinini, J., Demertzi, A., Bonhomme, V., Martial, C., Panda, R., Alnaggar, N., Annen, J., Gosseries, O., et al. (2023). Low-dimensional organization of global brain states of reduced consciousness. *Cell Rep.* *42*, 112491. <https://doi.org/10.1016/j.celrep.2023.112491>.
66. Zhang, P., and Sankaran, R. (2022). Autoencoder neural network for chemically reacting systems. *J. Mach. Learn. Model. Comput.* *3*, 1–28. <https://doi.org/10.1615/JMachLearnModelComput.2022045133>.
67. Pope, S.B. (2013). Small scales, many species and the manifold challenges of turbulent combustion. *Proc. Combust. Inst.* *34*, 1–31. <https://doi.org/10.1016/j.proci.2012.09.009>.
68. Zdybal, K., D'Alessio, G., Attili, A., Coussement, A., Sutherland, J.C., and Parente, A. (2023). Local manifold learning and its link to domain-based physics knowledge, *Appl. Energy and Combust. Sci.* *14*, 100131. <https://doi.org/10.1016/j.jaecs.2023.100131>.
69. Zdybal, K. (2023). Improving reduced-order models through nonlinear decoding of projection-dependent outputs. *Zenodo*. <https://doi.org/10.5281/zenodo.8319055>.
70. Zdybal, K., Armstrong, E., Parente, A., and Sutherland, J.C. (2020). PCAfold: Python software to generate, analyze and improve PCA-derived low-dimensional manifolds. *SoftwareX* *12*, 100630. <https://doi.org/10.1016/j.softx.2020.100630>.
71. Zdybal, K., Armstrong, E., Parente, A., and Sutherland, J.C. (2023). PCAfold 2.0—Novel tools and algorithms for low-dimensional manifold assessment and optimization. *SoftwareX* *23*, 101447. <https://doi.org/10.1016/j.softx.2023.101447>.
72. Cramer, F., Shephard, G.E., and Heron, P.J. (2020). The misuse of colour in science communication. *Nat. Commun.* *11*, 5444. <https://doi.org/10.1038/s41467-020-19160-7>.
73. LeNail, A. (2019). NN-SVG: Publication-ready neural network architecture schematics. *J. Open Source Softw.* *4*, 747. <https://doi.org/10.21105/joss.00747>.
74. Hansen, M.A., Armstrong, E., Sutherland, J.C., McConnell, J., Hewson, J.C., and Knaus, R. (2022). Spitfire. <https://github.com/sandialabs/Spitfire>.
75. Peters, N. (1988). Laminar flamelet concepts in turbulent combustion. *International Symposium on Combustion* *21*, 1231–1250. *Twenty-First International Symposium on Combustion*. [https://doi.org/10.1016/S0082-0784\(88\)80355-2](https://doi.org/10.1016/S0082-0784(88)80355-2).
76. Chollet, F., et al. (2015). Keras. <https://keras.io>.
77. Ash, J., and Adams, R.P. (2020). On warm-starting neural network training. *Adv. Neural Inf. Process. Syst.* *33*, 3884–3894. <https://doi.org/10.5555/3495724.3496051>.