

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE 10 February 2020		2. REPORT TYPE Conference Paper		3. DATES COVERED (From - To) 14 January 2020 - 28 February 2020	
4. TITLE AND SUBTITLE Attractor Reconstruction and Empirical Parameter Inference for Hydrogen-Oxygen Chemistry				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Mykhaylo Malakhov, Brianna Fitzpatrick, Rebecca Lopez, Abhishek Shivkumar				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER Q18B	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Laboratory (AFMC) AFRL/RQRS 1 Ara Drive Edwards AFB, CA 93524-7013				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory (AFMC) AFRL/RQR 5 Pollux Drive Edwards AFB, CA 93524-7048				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RQ-ED-TR-2020-012	
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution Statement A: Approved for Public Release; Distribution is Unlimited. PA Clearance Number: 20032 Clearance Date: 27 January 2020.					
13. SUPPLEMENTARY NOTES For presentation at RIPS (Research in Industrial Projects for Students) Prepared in collaboration with Andrews University, Carleton College, Marist College, University of California - Berkeley. The U.S. Government is joint author of the work and has the right to use, modify, reproduce, release, perform, display, or disclose the work.					
14. ABSTRACT Robust tools for characterizing nonlinear dynamical systems are indispensable in the development of in-space thrusters and other technologies of interest to the Air Force Research Laboratory. Although combustion can be easily simulated, the difficulty of experimentally observing a large number of chemical species complicates traditional methods for identifying system dynamics and ascertaining reaction rate coefficients. We utilize the attractor reconstruction procedure from convergent cross mapping to reconstruct the complete behavior of a continuously stirred hydrogen-oxygen tank reactor model from time-lagged observations (shadow manifolds) of individual species. Having demonstrated that a shadow manifold can effectively capture the information present in the entire attractor, we describe a novel optimization metric for data-driven parameter inference that only requires knowledge of a single observable. The proposed method infers parameters by minimizing the Wasserstein distance between binned shadow manifolds of a given reference data set and trial solutions. We demonstrate the superiority of our metric over standard approaches and present proof- of-concept results for reaction coefficient inference.					
15. SUBJECT TERMS N/A					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Daniel Eckhardt
Unclassified	Unclassified	Unclassified	SAR	69	19b. TELEPHONE NUMBER (Include area code) N/A

Research in Industrial Projects for Students



Sponsor

Air Force Research Laboratory

Final Report

Attractor Reconstruction and Empirical Parameter Inference for Hydrogen-Oxygen Chemistry

Student Members

Mykhaylo M. Malakhov (Project Manager), *Andrews University*
mykhaylo@andrews.edu

Brianna Fitzpatrick, *Carleton College*
brianna.r.fitzpatrick@gmail.com

Rebecca Alexandra Lopez, *Marist College*
rebecca.lopez1@marist.edu

Abhishek Shivkumar, *University of California, Berkeley*
ashivkum@berkeley.edu

Academic Mentor

An Do, an.do@cgu.edu

Sponsoring Mentors

Daniel Q. Eckhardt, daniel.eckhardt.3@us.af.mil

Robert S. Martin, robert.martin.101@us.af.mil

August 21, 2019

This project was jointly supported by the Air Force Research Laboratory and National Science Foundation grant DMS 1440415.

Abstract

Robust tools for characterizing nonlinear dynamical systems are indispensable in the development of in-space thrusters and other technologies of interest to the Air Force Research Laboratory. Although combustion can be easily simulated, the difficulty of experimentally observing a large number of chemical species complicates traditional methods for identifying system dynamics and ascertaining reaction rate coefficients. We utilize the attractor reconstruction procedure from convergent cross mapping to reconstruct the complete behavior of a continuously stirred hydrogen-oxygen tank reactor model from time-lagged observations (shadow manifolds) of individual species. Having demonstrated that a shadow manifold can effectively capture the information present in the entire attractor, we describe a novel optimization metric for data-driven parameter inference that only requires knowledge of a single observable. The proposed method infers parameters by minimizing the Wasserstein distance between binned shadow manifolds of a given reference data set and trial solutions. We demonstrate the superiority of our metric over standard approaches and present proof-of-concept results for reaction coefficient inference.

Acknowledgements

This project would have never existed without the support of the Air Force Research Laboratory (AFRL), which provided our research problem and sponsored the Research in Industrial Projects for Students (RIPS) program. We are deeply grateful to our sponsoring mentors, Robert Martin and Daniel Eckhardt, whose esteemed expertise and helpful counsel accelerated our progress every time we managed to get stuck on difficult theorems or fruitless simulations. We correspondingly thank An Do, our academic mentor, for her valuable guidance and assistance with our presentations and report. None of us would have survived the summer if not for RIPS Program Director Susana Serna's and Program Coordinator Dimi Mavalski's substantial support and organization. We also thank the 2018 AFRL RIPS team for their Final Report and code, which aided the initial stages of our work. Lastly, we acknowledge AFRL and the National Science Foundation for funding this research.

Contents

1	Introduction	13
1.1	The Proposed Problem	13
1.2	Our Solution and Approach	14
1.3	Report Overview	15
2	Mathematical Background	17
2.1	Dynamical Systems Theory	17
2.2	Mutual Information	21
2.3	Wasserstein Distance	22
3	Chemical Background	25
3.1	Hydrogen-Oxygen Combustion	25
3.2	<i>Cantera</i> Reactor Setup	26
4	Algorithms and Implementation	29
4.1	Introduction to Convergent Cross Mapping	30
4.2	Attractor Reconstruction	30
4.3	Parameter Inference	35
5	Results	41
5.1	Hydrogen-Oxygen Dynamics	41
5.2	Time Series Reconstructions	44
5.3	Parameter Inference	49
6	Conclusions and Discussion	53
Appendices		
A	Complete Reconstructions	55
B	Abbreviations	65
Selected Bibliography Including Cited Works		68

List of Figures

2.1	Lorenz attractor	19
2.2	Box-counting dimension	20
2.3	Koch snowflake	20
3.1	Reactor model setup	27
4.1	Flowchart of attractor reconstruction algorithm	31
4.2	Importance of time lag τ	33
4.3	Minima of mutual information	33
4.4	Flowchart of parameter inference process	36
4.5	Overview of SM-PDF construction	37
4.6	Dynamic framing issues	38
5.1	Time series: natural oscillations and sinusoidal perturbation	42
5.2	Time series: trivial and nontrivial equilibria	43
5.3	Time series: pressure perturbation	43
5.4	Time series: chaotic mass inflow perturbation	44
5.5	Effect of transients on reconstruction quality	45
5.6	Discarding transients before training	46
5.7	Reconstructions: H_2 from O_2 and O_2 from H_2	46
5.8	Reconstructions: H_2 from HO_2 and H_2O from H_2O_2	47
5.9	Reconstructions: temperature from OH and O from pressure	48
5.10	Reconstructions: H from H_2 and O_2 from temperature	48
5.11	Comparison of parameter inference with W_1 and L_2 metrics	50
5.12	Simultaneous parameter inference of A and n	51
5.13	Convergence of gradient descent algorithm	51

List of Tables

5.1	Correlations for all reconstructions in a chaos-perturbed reaction	47
-----	--	----

Chapter 1

Introduction

The Rocket Propulsion Division of the Aerospace Systems Directorate at the Air Force Research Laboratory (AFRL) is tasked with advancing novel technologies for rocket and satellite propulsion, fuels and propellants, and system analysis. As such, the Rocket Propulsion Division dedicates much of its basic research efforts toward high pressure combustion, energetic material design and synthesis, and modeling and simulation of propulsion experiments. Since accurate knowledge of chemical reactions is key for modeling and simulation of in-space combustion, the Division is highly interested in methods for characterization and calibration of the nonlinear dynamics ubiquitous in chemical reaction networks.

1.1 The Proposed Problem

The governing equations of chemical kinetics are well understood, and chemical reactions can be modeled with relatively simple ordinary differential equations. The primary challenge of modeling combustion and other chemical reactions lies in identifying a physically correct set of parameters for these equations. Traditional model calibration techniques require complete and precise measurements of the entire reaction in order to determine correct parameters. However, many of the chemical species that participate in combustion reactions are short-lived and only exist in minute quantities, making them difficult to access experimentally. Moreover, complex reaction mechanisms involve a very large number of species, further impeding experimental measurement. Due to these observational difficulties, model coefficients for most combustion reactions remain poorly known.

Recently researchers at AFRL have become interested in applying convergent cross mapping (CCM) to study noisy, chaotic systems that are difficult to measure experimentally. CCM was initially developed for detecting causality in complex ecosystem population dynamics by checking for the existence of convergent maps between time-lagged observations of different system observables [18]. However, the existence of convergent bijective maps between time-lagged observations of one observable and all other observables in a given system indicates that the first time-lagged signal has captured the information necessary to describe the entire

system. Thus, CCM can be adapted to create a powerful method for reconstructing noisy and chaotic data from limited observations. Indeed, AFRL researchers have already successfully employed it in Hall-effect thruster (HET) research for fusing time-varying measurements from multiple spatially-distributed probes into a unified reconstruction [5] and for model calibration for high-dimensional HET systems [12].

Although preliminary results indicate that CCM-derived techniques work surprisingly well for high-dimensional dynamical systems such as those exhibited by HET dynamics, CCM has only been comprehensively tested for “toy” problems with a small number of state variables. Thus, the goal of our project was to provide an intermediate step between these extremes by exploring the applicability of CCM-derived techniques to intermediate-scale nonlinear dynamics. In particular, we considered data from a 10-state hydrogen-oxygen reactor model.

1.2 Our Solution and Approach

We proposed two numerical algorithms: the first reconstructs the complete behavior of a dynamical system from a small set of observations, and the second performs model fitting using limited data. We also developed Python implementations of both procedures.

Our attractor reconstruction algorithm utilizes theorems from dynamical systems theory and information theory to reconstruct one time-dependent signal from another, granted that the two signals represent dynamically coupled observables. Sequentially, the algorithm accepts two time series as inputs, splits them into training and testing data sets, estimates an optimal time lag and embedding dimension, creates time-lagged embeddings (shadow manifolds) of the training data sets, constructs maps between the training shadow manifolds, and uses those maps to reconstruct one testing time series from the other. Since we only considered simulated time series, both testing time series were always known. Therefore, our algorithm concludes by evaluating the quality of the reconstruction against the actual time series.

Our parameter inference algorithm extends the attractor reconstruction algorithm to find a set of parameters that results in the closest possible “fit” between a mathematical model and a given reference data set. The algorithm accepts a reference time series and a mathematical model as inputs. First, the algorithm creates a shadow manifold from the reference time series. Next, the algorithm iterates over some set of potential parameter values. Each iteration of the process simulates the model, creates a shadow manifold from the trial time series, bins the shadow manifold as a probability distribution function (PDF), and evaluates the discrepancy between the reference and trial PDFs. After an initial arbitrary parameter guess, every subsequent iteration uses gradient descent to repeat the process with a new, better parameter set until a discrepancy of zero is achieved.

We demonstrated the success of both methods on simulated hydrogen-oxygen reaction data. Using the open source *Cantera* package [4], we generated time series data for hydrogen-oxygen combustion in a zero-dimensional continuously stirred tank reactor (CSTR) model. We considered a range of input parameters (e.g. mass inflow rate, temperature, pressure, etc.) and perturbed the mass inflow and piston boundary conditions. Subsequently, we

confirmed the applicability of our attractor reconstruction algorithm across all scenarios and analyzed how distance between species within the reaction network affects reconstruction quality. Finally, we attained proof-of-concept parameter inference results: our empirical parameter inference method recovered a number of the preset parameter values in *Cantera*.

Our techniques enable robust fusion of experimental and modeling studies of combustion, thereby providing AFRL with a tool for accurate calibration of in-space propulsion systems. Although we demonstrated the success of these algorithms for simulated hydrogen-oxygen reactions, the methods and tools presented in this report can be applied to a wide range of coupled dynamical systems. We believe that our work provides an important addition to the existing body of CCM research and would serve to guide future investigations into data-driven algorithms at AFRL.

1.3 Report Overview

Our report is organized in the following fashion:

Chapter 2 presents the mathematical background necessary for the methods we describe. We begin by introducing relevant concepts from dynamical systems theory and explaining Takens' Theorem, which serves as the basis for CCM. Next, we present the rudiments of information theory, which we later employ for estimating an optimal time lag. We also briefly present key notions from optimal transport theory and define the Wasserstein metric, which is essential in our parameter inference method.

Chapter 3 explains hydrogen-oxygen chemistry and the setup of our zero-dimensional CSTR model. For the purpose of this project, we sought a steady state reaction and then perturbed the input boundary conditions to create a variety of oscillatory behavior. Hence, we created a *Cantera* reactor model to represent such a system and allow for various perturbations and initial conditions.

Chapter 4 details the justification, requirements, and implementation of our attractor reconstruction and parameter inference algorithms. We first summarize the CCM technique, and then explain how our attractor reconstruction algorithm adapts CCM to reconstruct one time series from another. Complete attractor reconstruction from a single observable, in turn, facilitates our method for empirical parameter inference. This chapter also explains the challenges and limitations of the described approaches.

Chapter 5 presents the results of our research. First, we share the variety of system dynamics produced through manipulation of our reactor model. Second, we show examples of high-quality and lower-quality time series reconstructions, explaining the factors that determine reconstruction quality. Lastly, we portray the response surface for reaction coefficients and demonstrate several examples of successful parameter recovery by our process.

Chapter 6 concludes the report with a summary of our key observations and contributions. Additionally, we discuss the difficulties we encountered, possible routes for future research, and potential improvements to the proposed algorithms.

Chapter 2

Mathematical Background

This chapter presents the mathematical background necessary to understand the methods applied to the problems introduced in Chapter 1. Chemical reaction networks can be treated as dynamical systems, and we develop some tools generically applicable to certain dynamical systems that we will use to analyze chemistry. We begin with the foundations of topology required to discuss dynamical systems theory, and proceed to an exposition of the result that forms the backbone of our algorithm: Takens' Theorem.

At several stages in our project, we have encountered arbitrary parameters in our algorithms that have no canonical means of selection. There exists a selection criteria for one such parameter that involves information theory, which we introduce at a basic level in the section following the discussion of topology and dynamical systems.

Lastly, we discuss optimal transport theory, which becomes relevant in model calibration as we compare two probability distributions generated by our parameter reconstruction algorithm. In particular, we utilize the Wasserstein distance as a metric to quantify the discrepancy between a reference solution and our trial solution (where we perturb some parameter to generate each given trial solution). The Wasserstein distance can be thought of as the minimum cost it takes to convert one probability density function into another, the minimum cost of moving the “mass” of one distribution to the other.

2.1 Dynamical Systems Theory

Definition 2.1. A **dynamical system** is a tuple (T, M, Φ) with T an open subset of \mathbb{R} , M a smooth manifold, and $\Phi : T \rightarrow M$ a continuous function.

For our purposes, a smooth manifold can be thought of as a space that is locally isomorphic to \mathbb{R}^n at each time for some fixed n . For example, an ordinary sphere S^2 would be a 2-manifold as a sphere appears locally flat when you “zoom in” far enough. In this case, when describing the state space of some deterministic system as a manifold, it is best to think of

M as \mathbb{R}^n itself, an ambient space that encodes all possible values of n state variables that completely define the state of the system.

Definition 2.2. The **state** of a dynamical system is the smallest set of variables (called *state variables* or *observables*) such that the knowledge of these variables at $t = t_0$ together with the input for $t \geq t_0$ completely determines the behavior of the system for any time $t \geq t_0$.

As should be familiar from ordinary physical situations, many dynamical systems evolve toward some end state. For example, a damped harmonic oscillator (as in a pendulum where drag is present) will eventually lie motionless, and will never leave this state without some external perturbation. For an arbitrary dynamical system, we consider a more general limiting case:

Definition 2.3. An **attractor** of a dynamical system (T, M, Φ) is a subset A of M such that

- If $\Phi(t) \in A$, then $\Phi(t + \tau) \in A$ for all $\tau \geq 0$
- There exists a neighborhood of A called the **basin of attraction** for A denoted $B(A)$ such that if $\Phi(t) \in B(A)$ for any time t , then $\Phi(t)$ will eventually be in A . That is, there exists $T_0 \in T$ such that for $t > T_0$, $\Phi(t) \in A$.
- A is maximal, that is, there is no nonempty subset of A having the first two properties.

An attractor can be a finite set of points, a curve, a manifold, or even a set with a fractal structure, in which case it is referred to as a **strange attractor**.

Definition 2.4. A dynamical system exhibits **chaos** if it satisfies all of the following properties:

- Contains dense periodic orbits, that is, every point is arbitrarily close to a point of a periodic orbit
- Is hypersensitive to the initial condition of the system, in that points arbitrarily close to one another will eventually be far apart
- Is topologically transitive, that is, points in any neighborhood will eventually wind up in any other neighborhood

As an example of a dynamical system that exhibits chaos, we consider the Lorenz system [13]. The Lorenz system was originally proposed for modeling atmospheric convection, but is now typically used for providing an example of a smooth but chaotic broadband.

Definition 2.5. The three-equation **Lorenz system** is defined by

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y - x), \\ \frac{dy}{dt} &= x(\rho - z) - y, \\ \frac{dz}{dt} &= xy - \beta z.\end{aligned}$$




Figure 2.1: The Lorenz system simulated with initial condition $(x_0, y_0, z_0) = (1, 1, 1)$ and parameter values $(\sigma, \rho, \beta) = (10, 28, 8/3)$ for $0 \leq t \leq 50$.

In all following examples where we employ the Lorenz system, we use the initial condition $(x_0, y_0, z_0) = (1, 1, 1)$ and base parameter values $(\sigma, \rho, \beta) = (10, 28, 8/3)$, which produce a chaotic attractor. Figure 2.1 displays the Lorenz attractor generated from this parameter set.

It is often useful to consider the dimension of an attractor, which is easy to define when an attractor has the structure of a manifold, and slightly more complicated with strange attractors; Takens' Theorem utilizes the commonly used **box-counting** or **Minkowski-Bouligand dimension**, though other measures for fractal dimension are often used.

Definition 2.6. Suppose $N(\epsilon)$ is the number of boxes (generalized cubes) required to cover a set S lying in \mathbb{R}^n , where ϵ is the spacing of a grid. Then, the **box-counting dimension** of S is given by

$$\dim(S) = \lim_{\epsilon \rightarrow 0^+} \frac{\log N(\epsilon)}{\log \frac{1}{\epsilon}}$$

A motivating example for this definition is the unit square. If we pick $\epsilon_n = \frac{1}{n}$ to be the sequence along which we take the above limit, it is clear that n^2 small squares can fit in the unit square. Hence,

$$\dim_{\text{square}} = \lim_{n \rightarrow \infty} \frac{\log n^2}{\log n} = 2$$

as we would expect.




Figure 2.2: n^2 squares of side length $\frac{1}{n}$ can fit in the unit square, and $\dim_{\text{square}} = 2$ accurately captures the growth rate of the number of squares as $\frac{1}{n}$ decreases

Hence, the box-counting dimension for a set estimates the rate of growth of the number of boxes needed to cover the set as the grid of boxes grows finer. This dimension agrees with intuitive ideas of dimension, and extends to fractals and other sets that we would not be able to classically assign a dimension. An example of this is the Koch Snowflake.




Figure 2.3: First four iterations of the Koch Snowflake. The limit of these figures is the true fractal. Image from <https://en.wikipedia.org/wiki/File:KochFlake.svg> used under CC BY-SA 3.0.

The Koch Snowflake is generated iteratively; the beginning state is an equilateral triangle, and then the following rule is applied infinitely many times: take each line segment and replace the middle third with two legs of an equilateral triangle as seen in Figure 2.3. The limiting figure of this process is the fractal Koch Snowflake, and we can find its box-counting dimension (note that the fractal here is only the boundary, the interior set will have dimension 2 as usual). The “boxes” we need here are line segments, and we can see that each step, the number of line segments is multiplied by a factor of 4, with each length multiplied by $\frac{1}{3}$, so it is easy to see that

$$\dim_{\text{Koch}} = \frac{\log 4}{\log 3} \approx 1.2619$$

The dimension of a set not being a positive integer is classically meaningless, but is well-defined via box-counting; the thrust of this is that we are always able to assign a dimension to an attractor. We are now ready to state Takens’ Theorem:

Theorem 2.1. (*Takens’ Theorem*) *Let A be a compact attractor of dimension D belonging to a dynamical system (T, M, Φ) , and let $x : A \rightarrow \mathbb{R}$ be smooth. Then the **shadow manifold** consisting of all vectors of the form*

$$\vec{x}(t) = \begin{pmatrix} x(t) \\ x(t - \tau) \\ \vdots \\ x(t - (E - 1)\tau) \end{pmatrix}$$

living in \mathbb{R}^E is diffeomorphic to A , where $E > 2D$ is the **embedding dimension** and τ is the **delay time**.

In Theorem 2.1, $x : M \rightarrow \mathbb{R}$ can be thought of as an observable for the dynamical system lying on the attractor. A compact attractor here can be thought of as one that can be contained in a ball of finite radius. A diffeomorphism is an isomorphism of smooth manifolds; equivalently, a diffeomorphism is a smooth function with a smooth inverse. The definition of smooth varies with context; since it does not matter too much here, we will take smooth to mean C^∞ .

As we will discuss shortly, Takens' Theorem provides the theoretical backbone for CCM. In particular, let A_x denote the shadow manifold for an observable x on the attractor A . Since $A_x \cong A$ for any observable x , we have the following diagram:

$$\begin{array}{ccc} & A & \\ \cong \nearrow & & \nwarrow \cong \\ A_x & \overset{\cong}{\dashrightarrow} & A_y \end{array}$$

The transitivity and invertibility of isomorphisms ensures that the shadow manifolds of any two observables in a system satisfying the conditions of Takens' Theorem are diffeomorphic to each other, allowing us to reconstruct the time series for one variable from the time series of another.

2.2 Mutual Information

In theory, the shadow manifold is diffeomorphic to the original attractor for any value of τ and sufficiently large embedding dimension E , but in practice, with finite data, the choice of τ determines the goodness of fit of a reconstruction. Algorithms for finding good values of τ depend on information theory, which we will briefly explore here.

Definition 2.7. The **surprisal** of an event x is given by

$$I(x) = -\log P(x)$$

where $P(x)$ is the probability of the event, and the surprisal is always positive and larger for rare events.

Note that the surprisal of an event grows dramatically larger as its rarity increases, hence, surprisal can be regarded as a metric of the surprise of an observer upon measuring some event.

The **entropy** $H(X)$ of a signal X is the expected value of surprisal, e.g

$$H(X) = E[I] = -\sum_i P_i \log P_i$$

where the sum is replaced by an integral as necessary. The sum is (in theory) taken over all possible outputs or (in practice) taken over a certain binning of all possible outputs of the

signal. The units of entropy are *nats* if the logarithm is taken in base e , *bits* in base 2, or *dits* in base 10.

Entropy is the expected value of surprisal, hence, informally, it is the expected amount of surprise in a signal.

Definition 2.8. Given two signals X and Y , the **mutual information** between them $M(X, Y)$ is given by

$$M(X, Y) = \sum_y \sum_x P(x, y) \log \frac{P(x, y)}{P_X(x)P_Y(y)}$$

where $P(x, y)$ is a joint probability distribution over X and Y , and $P_X(x)$ is a single variable distribution, e.g.,

$$P_X(x) = \sum_y P(x, y)$$

where the sums are replaced by integrals as necessary.

The mutual information of two signals measures their dependence on each other, and is related to entropy in that

$$M(X, Y) = H(X) + H(Y) - H(X, Y)$$

where $H(X, Y) = -\sum_x \sum_y P(x, y) \log P(x, y)$ is the joint entropy of X and Y . The definition of $M(X, Y)$ in terms of entropy is easier to motivate: mutual information tells us how much information we lose when we look at two variables together as opposed to looking at them individually. For example, the mutual information of two signals is 0 if and only if they are identical; moreover, one can show using Jensen's inequality and the concavity of $\log x$ that $M(X, Y) \geq 0$, and as such, $M(X, Y)$ is a good metric for the similarity of two signals.

2.3 Wasserstein Distance

As will be explained below, one of the goals of our project is to infer a set of coefficients which govern the rates of various chemical reactions. Doing so will require the use of optimal transport theory, which we will develop now. Here we define the Wasserstein metric and its advantages to our project.

Definition 2.9. Given a set E of events and a set Σ of subsets of E which satisfies:

- $E \in \Sigma$.
- Σ is closed under complement.
- Σ is closed under countable union.

We call the pair (E, Σ) a σ -algebra. A **probability measure** on (E, Σ) is a function $\gamma : \Sigma \rightarrow [0, 1]$ such that

- For $A \in \Sigma$, $\gamma(\bar{A}) = 1 - \gamma(A)$.
- $\gamma(E) = 1$.

A probability measure is merely a general and rigorous formulation for assigning probabilities to events and combinations of events. In our formulation, the probability for a pair of points will correspond to the amount of mass transported between the points up to a total normalization factor.

Definition 2.10. Let μ and ν be **probability density functions** defined on X and Y respectively. A **transport plan** from μ to ν is a **probability measure** $\gamma \in \Gamma(\mu, \nu)$ where

$$\Gamma(\mu, \nu) = \{\gamma | \gamma(X, A) = \nu(A) \text{ and } \gamma(Y, B) = \mu(B)\}$$

Let $c : X \times Y \rightarrow \mathbb{R}$ define the **transportation cost** from a point $x \in X$ to a point $y \in Y$, then the p^{th} **Wasserstein Distance** between μ and ν is

$$W(\mu, \nu) = \left(\inf_{\gamma \in \Gamma(\mu, \nu)} \int_{X \times Y} c(x, y)^p d\gamma(x, y) \right)^{\frac{1}{p}}$$

The problem of optimal transport is to evaluate $W(\mu, \nu)$. In practice, $c(x, y)$ can be given by any metric (i.e. must be positive definite, symmetric, obey the triangle inequality, etc.)

Informally, the Wasserstein distance between two probability density functions (PDFs) is the minimum cost of transporting the “mass” of the first distribution to the second. The cost function in our use case will be the Euclidean distance between points, and in general serves as a generic metric that must obey the usual criteria (positive definiteness, symmetry, subadditivity).

Interpreting $W(\mu, \nu)$ is perhaps easier when one considers the discrete limiting case: consider $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_m\}$ where $x_i, y_j \in \mathbb{R}^n$ and d_{ij} is the Euclidean distance between x_i and y_j . Each point x_i and y_j has a “mass” associated to it, and we want to move the appropriate amount of mass from the x_i to each y_j , while minimizing the total mass-distance. Mathematically, this is

$$\min \sum_i \sum_j d_{ij} \gamma_{ij}$$

where γ_{ij} is the amount of mass moved from x_i to y_j .

The idea of our inquiry is that we can produce PDFs from reconstructed time series data, then, for any given reaction, produce a PDF from a set of rate coefficients. Then using the Wasserstein metric and optimal transport tools to optimize those coefficients to match the “experimentally produced” PDF. This is advantageous over curve-fitting techniques for several reasons:

1. To find rate coefficients using curve-fitting algorithms, we would need time series data for the rate of a certain reaction, where all we have (and all we could reasonably expect to obtain from an experiment) is the time series for the mole concentration of each species.

2. Using CCM, we can infer rate coefficients of species that are far removed in the reaction chain from our test species (e.g transients that do not appear at significant concentrations).

Chapter 3

Chemical Background

Our project is established in providing an intermediate step between high dimensional and low dimensional systems. Hydrogen-oxygen chemistry can be represented by a dynamic system and allows us to observe the application of CCM techniques in that transition. Hydrogen-oxygen combustion is an exothermic reaction and one of the few “clean burning” reactions. This is of increasing interest to AFRL and so we discuss the reaction mechanism and the chemistry behind it. Additionally, we discuss our setup in the Python interface of the *Cantera* software [4] for a CSTR model, which includes the manipulation of input reservoirs of individual reactants as a perturbation to our system. Here we detail the assumptions our model makes and the processes it represents.

3.1 Hydrogen-Oxygen Combustion

Combustion can be defined as a high temperature exothermic reaction between a fuel and an oxidant. For the purposes of developing an intermediate-scale model, we will be working with a common combustion example of hydrogen, H_2 , as the fuel and oxygen, O_2 , as our oxide, which reacts to form water and energy. Generally, combustion is a series of elementary reactions forming a complex reaction network. It requires a temperature high enough to ignite our elements. In comparison to other fuels, hydrogen has a lower ignition energy and also burns faster than natural gases. Hydrogen also has a wider flammability range in comparison to other hydrocarbon fuels.

Most often, the use of this reaction alludes to the fuel for rocket engines. As one of the few nonpolluting fuels that are plentiful in the environment, hydrogen has become extensively favored as a reactant. An acceptable and common concern is the backfire of the combustion resulting in the complete destruction of entire fueling systems [3]. It is difficult to fully observe what is occurring in an experiment due to lab technology limitations. Researchers are only aware of the destruction once it has already occurred. A common way to reduce backfire is to reduce the temperature of the residual gases.

3.2 *Cantera* Reactor Setup

Cantera is a chemical software used through Python to model chemical kinetics reactions [4]. In our case we use it to simulate hydrogen-oxygen combustion, since it is a simple, well-understood combustion reaction that typifies more complex forms of combustion. *Cantera* is able to calculate states that would be hard to assess experimentally, such as equilibrium, reaction rates, the pre-exponential factor, molarity of each species, etc. It is a software that can replicate a reaction and is able to attain calculations. *Cantera* begins by setting the solution of gases that will be “fed” into the reactor. There are various parameters that can influence the reaction that can be added to the program, for example temperature, pressure, volume, etc. *Cantera* is also able to add more complex parameters, such as having a wall or a piston in the reaction. We use *Cantera* to obtain data in our reaction by having a file that contains each species and time produced by *Cantera*. We are then able to use this data to reconstruct other species. We first simulate the chemical reaction to generate data, and then use that data to test our algorithms.

As the dynamics of chemical reactions are governed by a system of ordinary differential equations, a reaction network can be treated as a dynamical system. In our case, the rate of a given reaction is given by the **Generalized Arrhenius equation**

$$k = AT^n e^{\frac{-E_a}{k_B T}}$$

where k_B is the Boltzmann constant, T the temperature, and E_a the activation energy. The reaction network we chose was hydrogen-oxygen chemistry, which has several crucial advantages for our project. There are 8 chemical species, leading to a dynamical system of dimension at least 8 (in fact, the parameters of temperature and pressure make our system at least 10-dimensional). The various reactions in this network are relatively well-understood. There are 21 elementary reactions contributing to the overall reaction network for hydrogen-oxygen combustion. The species are H_2 , O_2 , H_2O , HO_2 , H_2O_2 , OH^- , O^{2-} , and H^- . Hydrogen-oxygen combustion is used to fuel rocket engines as it produces a relatively high amount of energy for a low ignition energy.

Our *Cantera* reactor setup, as seen in Figure 3.1, consists of several reservoirs of individual reactants with separate mass flow controllers to allow for separate perturbation of the influx of each reactant. For example, consider a sinusoidal input perturbation of the following form:

$$\dot{m}(t) = \begin{cases} c & t < t_p \\ c + A \sin k(t - t_p) & t \geq t_p \end{cases}$$

where t_p is the time when the perturbation turns on and $\dot{m}(t)$ is the mass flow rate. Input perturbations are reflected in the behavior of the system dynamics.

In a CSTR all mass inflow is immediately homogenized into the reactor mixture. The reactants enter through a mass flow controller (MFC) to the reactor. A valve is present to facilitate outflow of reactants and products, governed by a valve coefficient, which is the ratio of mass influx to the reactor against pressure drop across the valve. Also, the value

feeds into an infinite reservoir of gas held at a fixed composition, temperature, and pressure. We are then able to produce data with arbitrary initial conditions and perturbations for analysis through our techniques. The reactor models we produce using this *Cantera* setup are presented in Chapter 5. In the following chapter, we discuss our chosen algorithm methods and the process we take for implementation.

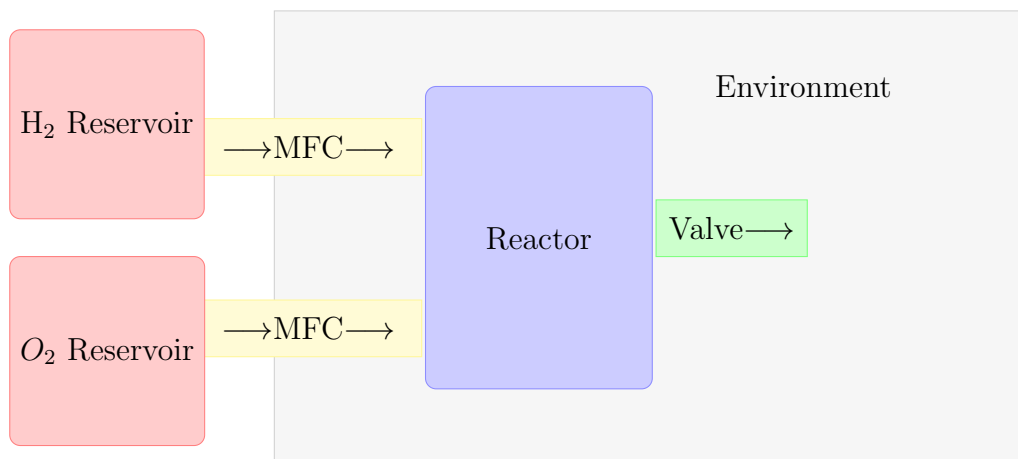


Figure 3.1: Reservoirs of H_2 and O_2 feed into the continuously stirred reactor via MFCs. The reactor has a valve that facilitates the outflow of reactants and products into the environment.

Chapter 4

Algorithms and Implementation

In this Chapter we describe data-driven approaches for attractor reconstruction and parameter inference in coupled dynamical systems, and explain how we applied these methods to simulated hydrogen-oxygen combustion data.

Attractor reconstruction facilitates characterization of the complete behavior of a dynamical system from a limited set of observables. The method we describe adapts the underlying procedure of convergent cross mapping in order to predict any unknown time series in a given dynamical system from any known time series in the same system and historical values of the unknown time series. As implemented, our algorithm accepts two time series as inputs, splits them into training and testing datasets, reconstructs the testing portion of each time series from the shadow manifold of the other time series, and evaluates the quality of the two reconstructions. We provide both MATLAB[®] (MathWorks[™], R2019a) and Python 3.7 implementations of the phase space reconstruction algorithm.

Attractor reconstruction justifies our method for parameter inference. We assume that successful reconstruction of all time series constituting the attractor of a given dynamical system from some shadow manifold M_X indicates that M_X contains the information necessary to describe all observations of the system. If attractor reconstruction succeeds, model calibration algorithms can solely consider a shadow manifold of a single time-dependent observable instead of attempting to work with the presumably more complicated system manifold. Consequently, we extend the attractor reconstruction algorithm summarized above to propose a novel optimization metric for data-driven parameter inference.

Generally speaking, any empirical method for parameter inference seeks to minimize the “distance” between a known reference solution and generated trial solutions from a mathematical model. The parameter set yielding the smallest possible distance is the correct one. Our parameter inference method adheres to this standard process, but employs an innovative measure of distance between dynamical systems. The optimization metric we describe offers several advantages over existing approaches. As already noted, our metric only requires knowledge of a single observable, thus allowing for parameter inference from limited observational data. Moreover, we demonstrate that the approach we describe remains robust to noise and chaos while providing a meaningful sense of “closer” even for data that are very

“far” in state space. We provide a Python 3.7 implementation of our parameter inference algorithm.

4.1 Introduction to Convergent Cross Mapping

CCM was developed as a technique for distinguishing between causality and correlation in complex systems such as those exhibited in ecological settings [18]. The method succeeds in identifying causal networks for nonseparable weakly connected dynamical systems, where earlier approaches such as the Granger causality test [11] fail.

Intuitively, CCM tests whether two time series belong to the same dynamical system by evaluating the skill with which they are able to reconstruct each other. For example, consider two (not necessarily coupled) observables X, Y and their corresponding shadow manifolds M_X, M_Y . We have two cases:

Presence of causality: Without loss of generality, suppose that X causes Y . Then Y contains information about X and hence the two state variables belong to the same dynamical system. By Theorem 2.1 bijective maps between M_X and M_Y exist. CCM first employs a numerical algorithm to find maps between M_X and M_Y , and then those maps are used to reconstruct M_X from M_Y . Finally, an estimated time series of X is obtained from M_X . The reconstructed time series of X should converge to its original time series as the size of the “library” of data used to construct the cross maps (the time series length) increases. Therefore, when X causes Y the shadow manifold of Y can be used to reliably approximate a time series of X .

Absence of causality: Suppose that construction of convergent cross maps between the shadow manifolds of X and Y fails. This implies that no diffeomorphism exists between the two manifolds M_X, M_Y , and therefore by the contrapositive of Theorem 2.1 we have that X and Y do not belong to the same dynamical system. It can be concluded that X and Y are not causally coupled.

Instead of using CCM to detect causality, the algorithms we describe assume mutual causality between state variables. CCM begins with state variables whose coupling or lack thereof is unknown, but our approach only applies to systems where all state variables are already known to be coupled. Under this assumption, Theorem 2.1 implies that diffeomorphisms exist between the shadow manifolds of any two state variables. The present methods hinge on the existence of such bijections between the shadow manifolds of state variables in a given dynamical system.

4.2 Attractor Reconstruction

Figure 4.1 provides a brief overview of the attractor reconstruction process. In what follows, we explain each step of the algorithm precisely.

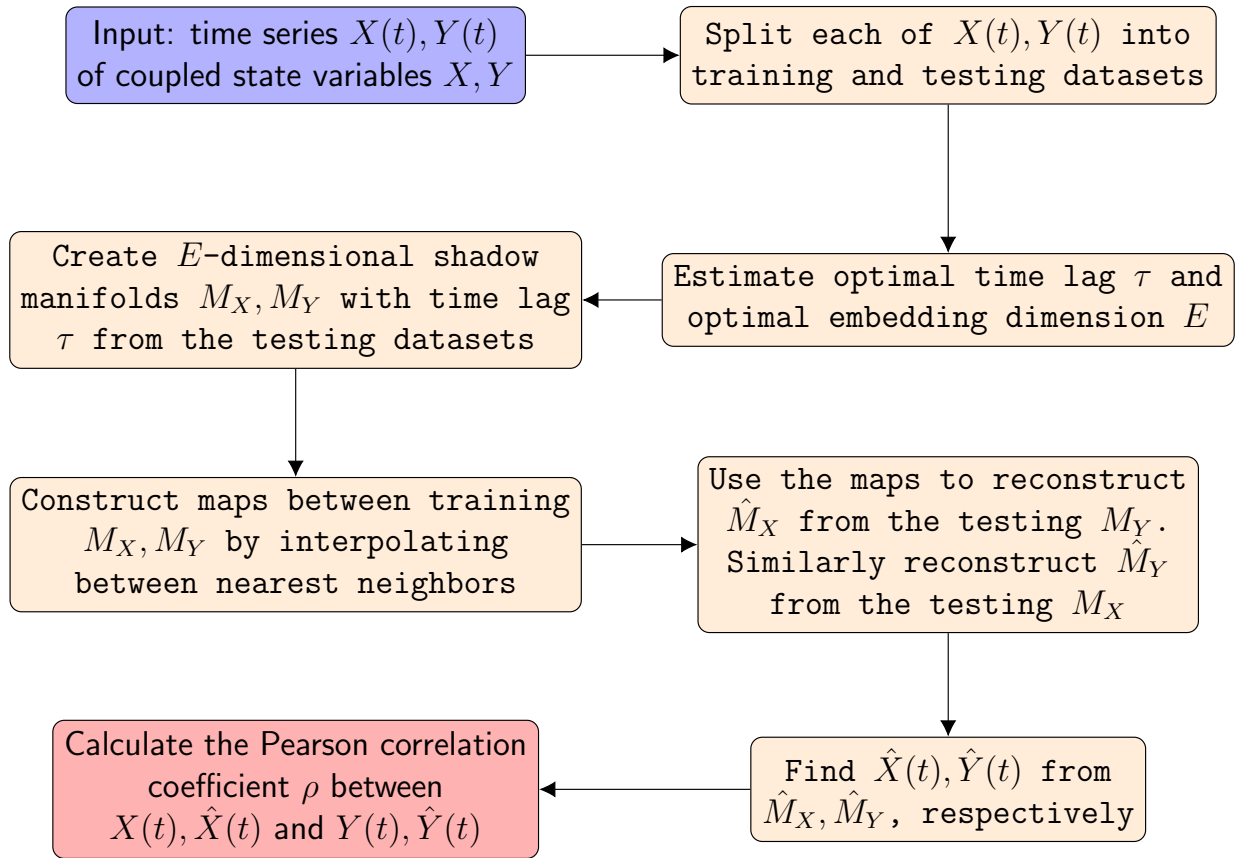


Figure 4.1: Flowchart of our attractor reconstruction algorithm.

Consider a smooth, real dynamical system (T, M, Φ) with attractor dimension D that satisfies the criteria for Theorem 2.1. Let X, Y be any two state variables of (T, M, Φ) , and let $X(t), Y(t)$ be finite time series of equal length for X and Y , respectively. Our phase space reconstruction algorithm takes $X(t), Y(t)$ as inputs and is comprised of the following consecutive stages:

4.2.1 Split data into training and testing

If the length of $X(t)$ is not even, the algorithm discards the first data points of both $X(t)$ and $Y(t)$ before proceeding further¹. If the input time series length is already even, no data is discarded. Henceforth, we use the notation $X(t), Y(t)$ to denote the (potentially shortened) even-length versions of the input time series.

Subsequently the algorithm splits each time series $X(t), Y(t)$ in half. Our implementation of the algorithm allows using the first half of the time series for training and the second half for testing, or the second half of the time series for training and the first half for testing. By default, our algorithm uses the first halves of $X(t)$ and $Y(t)$ for training, and the second halves for testing.


4.2.2 Estimate optimal lag and embedding dimension

Although Theorem 2.1 holds for any value of time lag τ as long as the embedding dimension E is high enough, on a practical level the choices of τ and E significantly impact the performance and quality of phase space reconstruction.


For an illustration, consider Figure 4.2. Clearly, setting $\tau = 0.12$ seconds results in a shadow manifold that appears much more similar to the Lorenz attractor than the shadow manifold created when $\tau = 5.6$ seconds. Since our algorithm constructs maps between shadow manifolds via numerical interpolation, it will succeed more readily if the shadow manifolds of $X(t)$ and $Y(t)$ closely resemble each other. Therefore, we want to find a value of τ that is sufficiently small to capture local behavior while remaining large enough to capture global behavior. Such values of τ will result in a “nice” shadow manifold, as in Figure 4.2b. The method proposed in [7] attempts to use mutual information to find this balance for τ , so we adapt it in our algorithm. Our method sweeps through a range of possible time lags with a step size equal to the time series step size, and returns the smallest τ at which the average mutual information between $X(t)$ and $X(t - \tau)$ stops decreasing. See Figure 4.3 for an example of the algorithm output.

Although the time lag τ produced by the above method usually results in an excellent reconstruction, the algorithm we employ does not always find the best possible value of τ . More consequentially, the algorithm sometimes fails entirely. When a time series $X(t)$ is excessively noisy or stochastic, increasing the time lag τ would always decrease predictive

¹Since $X(t)$ and $Y(t)$ are of equal length, one can equivalently check if the length of $Y(t)$ is even.



(a) $\tau = 5.6$ seconds.



(b) $\tau = 0.12$ seconds.

Figure 4.2: Shadow manifolds of the Lorenz attractor constructed from the X time series with embedding dimension $E = 3$. The two figures differ only in choice of τ .

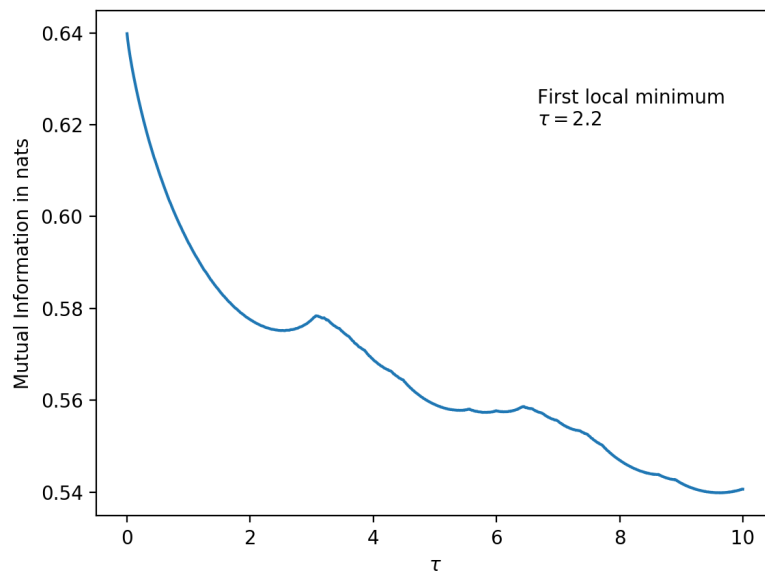


Figure 4.3: Mutual information between an observable $X(t)$ and its time lag $X(t - \tau)$ with respect to the time lag value τ .

power. In this case, the mutual information between $X(t)$ and $X(t - \tau)$ would exhibit no local minima.

According to Theorem 2.1, an embedding dimension $E = 2D + 1$ will be sufficiently high to ensure diffeomorphism between shadow manifolds of X, Y . However, in practice a much lower embedding dimension will often suffice. We want to minimize the computational time of our algorithm to ensure scalability to high-dimensional dynamical systems, so we estimate the minimum embedding dimension. To do this, we employ the algorithm described in [2].

Our algorithm estimates an optimal time lag τ and an optimal embedding dimension E twice. The same procedure is used both times, but on different datasets. We find τ_X from the training portion of $X(t)$ and find τ_Y from the training portion of $Y(t)$. Similarly, we find E_X from the training portion of $X(t)$ and E_Y from the training portion of $Y(t)$. Later on, τ_X and E_X are used when reconstructing the testing portion of $Y(t)$ from $X(t)$, while τ_Y and E_Y are used when reconstructing the testing portion of $X(t)$ from $Y(t)$.

4.2.3 Construct training shadow manifolds

By definition, the shadow manifold M_X of $X(t)$ is the set of all vectors of the form

$$\vec{x}(t) = \begin{pmatrix} x(t) \\ x(t - \tau) \\ \vdots \\ x(t - (E - 1)\tau) \end{pmatrix}$$

living in \mathbb{R}^E . The shadow manifold M_Y of $Y(t)$ is constructed analogously.

4.2.4 Construct maps

Once the training shadow manifolds M_X and M_Y are created, maps are numerically constructed between them by building a library of points on M_X with their corresponding points on M_Y . That is, our algorithm associates each point on the training portion of M_X with the point on the training portion of M_Y that has the same timestamp. Since M_X, M_Y have the same number of data points, this results in a bijection between the two shadow manifolds.

4.2.5 Reconstruct testing shadow manifolds

For an unknown input point x on M_X , the k -nearest neighbors algorithm finds the closest neighbors of x on M_X . Then the weighted mean of corresponding points on M_Y is taken to find the approximate point \hat{y} corresponding to x . Essentially, we build a library of correspondence between M_X and M_Y and then “interpolate” between known input points of training M_Y to find output points, which together constitute \hat{M}_Y . \hat{M}_X is created similarly, except we begin with an input point y on M_Y and then “interpolate” between known input points of training M_X .

4.2.6 Reconstruct testing time series

After \hat{M}_X is created the procedure restricts it to the first dimension, producing a 1-dimensional time series. In other words, we remove the lags of each vector to flatten the manifold back into a single dimension. This yields the reconstructed time series $\hat{X}(t)$. Similarly, our algorithm flattens \hat{M}_Y to create the reconstructed time series $\hat{Y}(t)$.

4.2.7 Output correlation coefficients

In practice, our algorithm for attractor reconstruction would be employed for studying dynamical systems whose complete attractor manifold is unknown. However, complete time series of all observables are always known in the proof-of-concept simulations we examine. In this project our goal is to evaluate the success of the described algorithm by comparing the reconstructed testing time series $\hat{X}(t)$ to the original testing time series $X(t)$.

We evaluate the success and quality of our reconstructions by calculating the Pearson correlation coefficient ρ between $\hat{X}(t)$ and $X(t)$ after subtracting the mean of each one. Since the Pearson correlation coefficient is a measure of linear independence, the correlation coefficient between a dataset and its mean is often high. Therefore, a reconstruction algorithm that performs little better than the approximation achieved by taking the mean of the original data would still yield “good” correlation coefficients. In order to provide more meaningful and sensitive correlation values, we instead compute

$$\rho^* \left(\hat{X}(t), X(t) \right) \doteq \rho \left(X(t) - \langle X(t) \rangle, \hat{X}(t) - \langle \hat{X}(t) \rangle \right)$$

where $\langle X(t) \rangle$ and $\langle \hat{X}(t) \rangle$ denote the arithmetic means of $X(t)$ and $\hat{X}(t)$, respectively.

4.3 Parameter Inference

Figure 4.4 provides a brief overview of the general parameter inference process we followed for this project. Since system identification is not the focus of our work, we assume that the dynamics of the system in question are already understood; that is, the governing equations for which we seek to find parameters are known. Each iteration of the parameter inference process begins by running the known mathematical model with a set of values for the unknown input parameters $\vec{\beta}$ to produce a trial solution $\vec{x}_{trial}(t)$. The generated trial solution is then compared to the reference solution $\vec{x}_{ref}(t)$. In particular, the algorithm uses some optimization metric ε to calculate a measure of discrepancy between $\vec{x}_{trial}(t)$ and $\vec{x}_{ref}(t)$. If the optimization metric, also known as a cost function, yields a discrepancy value below a predefined convergence criterion ε_0 , we conclude that the current parameter values are sufficiently optimal. Otherwise, the model simulations are redone with a new set of parameters and the discrepancy between the new trial solution and the reference solution is computed. This iterative process is repeated until a $\vec{\beta}$ is found such that $\varepsilon(\vec{x}_{ref}(t), \vec{x}_{trial}(t)) < \varepsilon_0$.

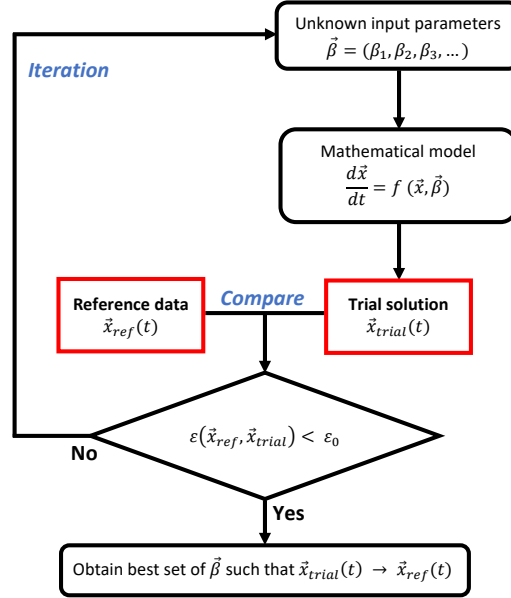


Figure 4.4: Flowchart of our parameter inference process.

The optimization metric we describe takes a time series from the trial solution and a time series from the reference data as inputs, and returns a measure of discrepancy between them. The method requires that both time series represent the same state variable. For example, when inferring parameters for hydrogen-oxygen chemistry we might use an O_2 time series from the trial solution and an O_2 time series from the reference data. The function is comprised of two steps. First, we create a shadow manifold representation of each input time series and bin it as a normalized probability distribution function to create what we refer to as a shadow manifold probability distribution function (SM-PDF). Next, we compute the Wasserstein distance between the two SM-PDFs. This distance constitutes the output of our cost function. In what follows, we explain each step of the process in detail.

4.3.1 Construction of Shadow Manifold Probability Distribution Function (SM-PDF)

Figure 4.5 provides an overview of the process for constructing SM-PDFs.

First, our method calculates an optimal time lag τ and embedding dimension E for the given time series as described in 4.2.2. Then, the algorithm uses these calculated values of τ, E to construct a shadow manifold from the time series as described in 4.2.3. Note that although Figure 4.5 displays a 2-dimensional shadow manifold for illustration, higher dimensions may be used as deemed necessary by the algorithm for estimating E .

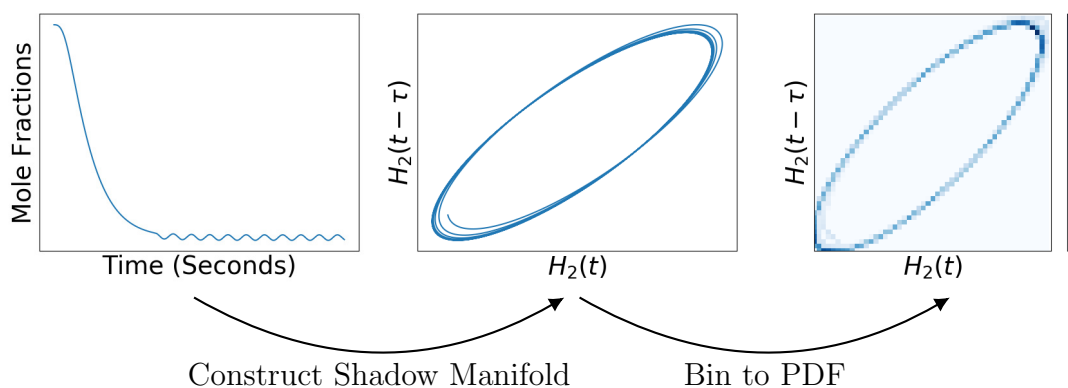


Figure 4.5: This figure provides a high-level overview of the process for constructing a shadow manifold probability distribution function (SM-PDF).

Next, we use the `numpy.histogramdd` function from the NumPy Python library to bin the shadow manifold into a discrete PDF. The `numpy.histogramdd` function computes the multidimensional histogram of an array using given values for bin edges. To ensure that SM-PDFs can be accurately compared when computing the Wasserstein distance between them in the following step, we use a constant number of bins for each dimension across all iterations of the parameter inference process. However, more must be said on the choice of bin edges, since some binning schemes can result in falsely similar SM-PDFs.

We consider two methods for selecting the bin edges: dynamic framing and static framing. Dynamic framing uses the minimum and maximum values along each dimension of every shadow manifold to define the bin edges for that particular manifold. This approach recalculates the bin edges for each iteration of the parameter inference process and is generally a poor choice because it is susceptible to outliers. Since the multidimensional binning grid must include all data from the shadow manifold, a single point that is removed from the bulk of the manifold in state space might affect the level of detail captured by the binning. Additionally, dynamic framing can create very similar SM-PDFs for shadow manifolds that are far apart in state space. As an illustration of this problem, Figure 4.6a displays a pair of 2-dimensional shadow manifolds of the temperature time series from two different *Cantera* simulations. When shown on the same coordinate system, the two shadow manifolds are clearly distinct. However, when binned using the dynamic framing method and plotted on top of each other (Figure 4.6b) the two SM-PDFs now appear indistinguishable. In this case, any distance metric between the two SM-PDFs would falsely yield a value close to zero.

To avoid the issues stemming from dynamic framing, we instead employ a single, static frame for all iterations of the parameter inference process. That is, our algorithm sets constant bin edges and uses them for every iteration of the parameter inference process. Static framing avoids the pitfalls of dynamic framing, but can only be used when all shadow manifolds fall within the set bin edges. In general, finding a suitable set of bin edges to capture every possible shadow manifold for the given observable is an open problem. Trial solution shadow manifolds can drastically change position in state space when sweeping over




Figure 4.6: (a) Two shadow manifolds that, although similar in shape, are located far apart in state space when plotting on the same coordinate system. (b) Associated probability distribution functions of shadow manifolds plotted on a single set of axes.

a range of parameters, potentially moving outside of the pre-set bin edges. If this happens, the trial solution SM-PDF will not capture the entire shadow manifold and any metric of discrepancy between the reference SM-PDF and the trial SM-PDF will fail to provide any useful information. This problem can be easily resolved, however, if the trial solutions are known to be bounded. For hydrogen-oxygen chemistry, all species state variables represent mole fractions and hence fall in the range $[0, 1]$. Consequently, all n -dimensional shadow manifolds created from time series of species densities will be bounded by an n -dimensional unit hypercube. In this project we only used species time series for the parameter inference process, and set the bin edges along each dimension to $[0, 1]$.

4.3.2 Computation of Wasserstein Distance W_1

Having represented the behavior of the entire system by binned SMs, the final step is to identify an appropriate measure for the discrepancy between the SM-PDF of the reference data and the SM-PDF of the current trial solution. We use optimal transport tools to find the Wasserstein distance. In particular, we use `ot.emd2` from the Python Optimal Transport library to return the Wasserstein distance between two PDFs (passed as `numpy.ndarray` objects flattened to vectors) given a distance matrix (which itself can be calculated using an arbitrary metric). The minimum of the Wasserstein metric across some sample space is used to obtain the optimal set of parameters to reproduce the reference solution.

4.3.3 Parameter Inference

With a robust framework for constructing and evaluating reconstructions in place, we employed various techniques to infer parameters of our systems. The basic method involves generating time-series data for a trial set of parameters and comparing the resulting SM-PDF with the SM-PDF for a reference solution, corresponding to the ground truth. The difficult step in this procedure is to infer a new set of parameters after a given run.

As typically implemented, parameter inference algorithms employ an optimization scheme to choose a new set of parameters after every iteration that failed to achieve a sufficiently low value from the cost function. The most popular optimization schemes involve some variant of machine learning or genetic algorithms, which must be chosen depending on the topological properties of the given dynamical system and cost function. In this project we employed two basic methods: brute force and gradient descent.

For the former, the method is simple: we iterate over some region of parameter space at some spacing, say $\pm 50\%$ of the known parameters, and look for the minimum of W_1 . This strategy was applied to some success (as we will discuss in Chapter 5), but has the obvious drawback of being extraordinarily slow. In particular, the problem scales extremely poorly with dimension, e.g, the time complexity is exponential in dimension.

The latter method is also fairly well-understood: we evaluate W_1 at a point in parameter space, as well as at some of this base point's neighbors, and pick our new point in the direction of steepest descent. We employed this algorithm as well to great success, although it has certain drawbacks in determinacy with respect to the brute force method. Our results within this area serve as a proof of concept, and the gradient descent methods we used are far from optimal.

Chapter 5

Results

In hydrogen-oxygen combustion all species densities are coupled to varying degrees through temperature and belong to the same dynamical system. Therefore, our motivating example satisfies the coupling criterion explained in Section 4.2 and is hence suitable for the proposed algorithms.

We applied our attractor reconstruction and parameter inference technique to various time series data with a variety of input perturbations. First, we note the system's sensitivity to various setups and conditions.

An intermediate goal was to construct bidirectional convergent cross maps between various species densities and temperatures. Since combustion is a series of elementary reactions that are often favored one way for producing trace amounts of a species, we analyzed the impact this has on our reconstructed time series data. Additionally, we found that reconstruction quality is impacted by the distance in the reaction network as well as the heat released during the reaction.

Finally, we demonstrate the success of our method for parameter inference by predicting reaction rate constants for the reaction mechanism of hydrogen-oxygen chemistry.

5.1 Hydrogen-Oxygen Dynamics

As mentioned in Section 3.2, we are able to manipulate the following parameters within our model for hydrogen-oxygen combustion in *Cantera*:

- Pressure
- Temperature
- Initial ratio of hydrogen to oxygen
- Mass inflow rate
- Angular frequency of sinusoidal perturbations

- Amplitude of perturbation
- Time at which the perturbation begins

Here we generalize the behavior observed in the reaction dynamics from changes in the initial conditions.

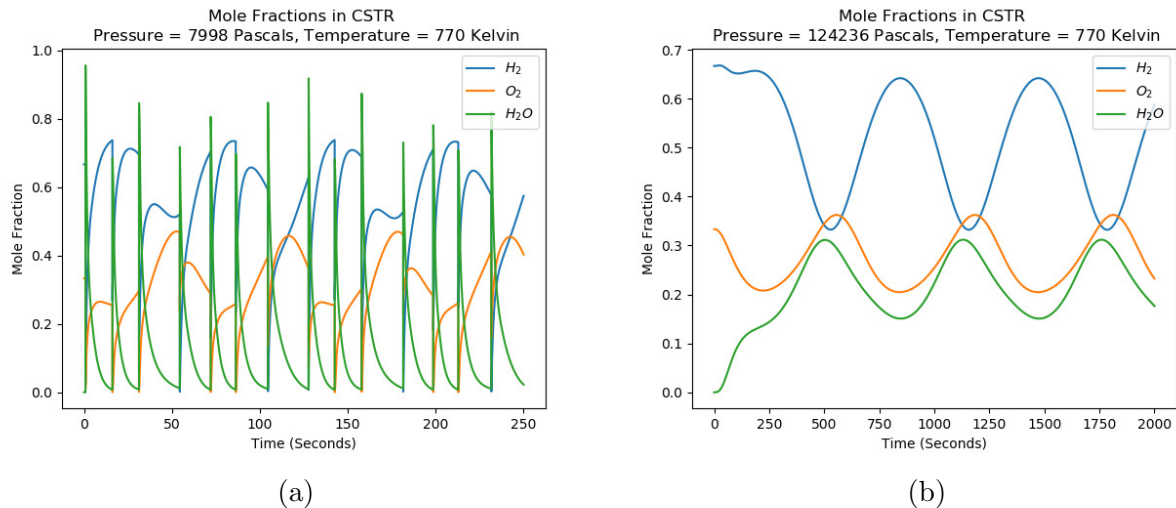
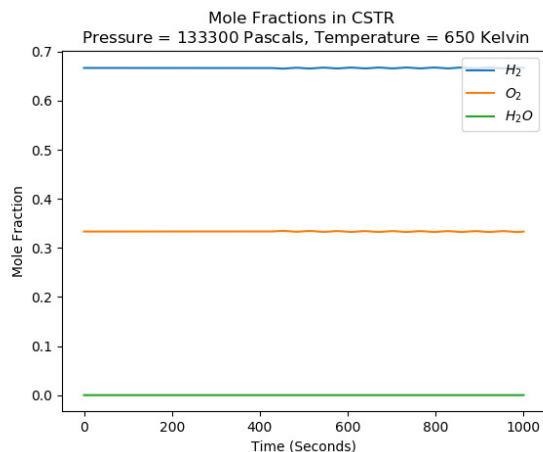


Figure 5.1: (a) A low pressure system of 7998 Pascals, a temperature of 770 Kelvin, and no input perturbation. There is natural oscillatory behavior between the reactions. (b) A high pressure system of 124236 Pascals and 770 Kelvins. In this system, both H_2 and O_2 are being perturbed with a sinusoidal function immediately.

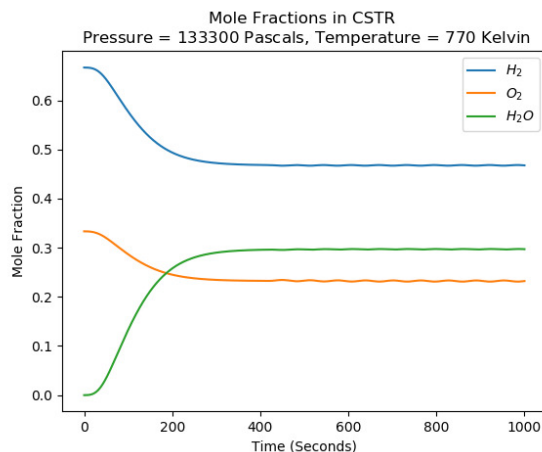
We observed that low pressure systems usually require a high temperature in order to ignite the reactants and produce a combustion reaction. Consider Figure 5.1a as an example where our low pressure of 7998 Pascals and high temperature of 770 Kelvin lead to consistent combustion. We are able to observe natural oscillatory behavior without an artificial perturbation.

On the contrary, high pressure systems drive the reaction towards a steady state rapidly. For example, as seen in Figure 5.2b, an initial reaction occurs and the input perturbation becomes apparent beyond the steady state. A sufficiently high enough temperature is still required in a high pressure system in order to ignite the reactants and initiate the reaction. Consider Figure 5.2a where the temperature is too low to produce a reaction. And so, in turn we observe a trivial equilibrium where no reactions occur.

We also considered a variety of input perturbations. To begin, we started with a low-frequency sinusoidal perturbation of the mass flow rate of our limiting reactant: oxygen. For example, consider Figure 5.2b. The input perturbation is apparent in the time series data since the reaction follows a sinusoidal movement. We found that in systems with both input reservoirs of H_2 and O_2 being perturbed by a sinusoidal function, the system immediately entered oscillatory behavior past a rapidly reached steady state. In Figure 5.1b, observe that the system dynamics follow the same period and behavior.




(a)

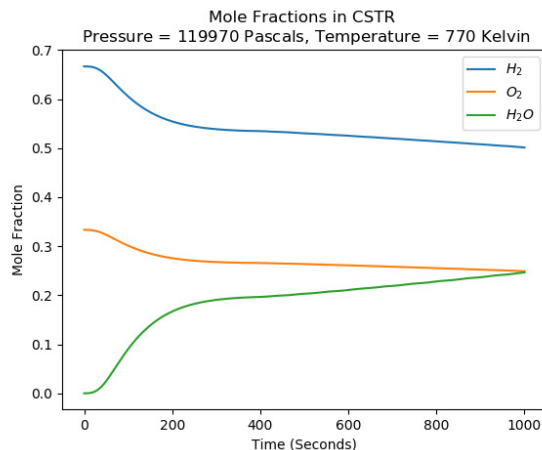


(b)

Figure 5.2: (a) A high pressure system of 13300 Pascals and a low temperature of 650 Kelvin. The input perturbation of O_2 begins past 450 seconds. (b) A high pressure system of 13300 Pascals and a temperature of 770 Kelvin with an input sinusoidal perturbation of the O_2 mass flow rate past 420 seconds.



(a)



(b)

Figure 5.3: (a) Time series data for a pressure perturbation. The piston motion begins past 350 seconds and thus we observe the oscillatory behavior of pressure past the initial increase of pressure. (b) A system with an initial pressure of 119970 Pascals and a temperature of 770 Kelvin. The piston motion begins past 350 seconds and thus we observe a shift in the behavior of the system.

If we increase the frequency for the perturbation of the mass flow rate then the reaction should become faster and result in decreased wavelengths. We expected this to then pull apart the reaction time scales, pushing the slower chemistry off equilibrium. However, the faster we perturbed the mass flow rate, the lower the resulting amplitude became. We tried

the same process with perturbing O_2 and H_2 but the same results followed. Therefore, mass inflow perturbation did not achieve its intended goal.

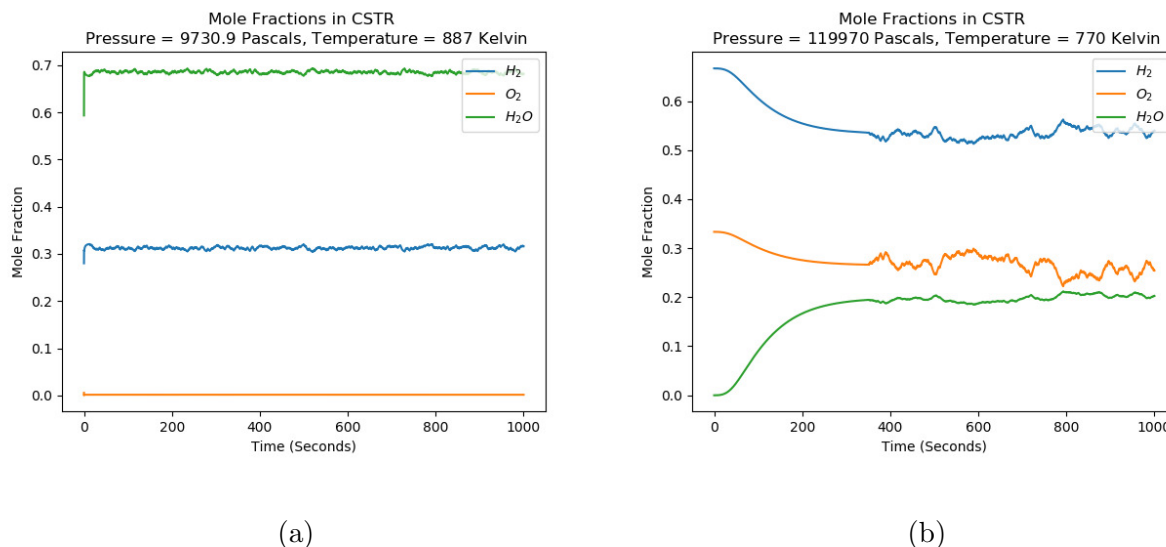


Figure 5.4: (a) Time series data with the Lorenz system output immediately being put in as the input perturbation in a system of 9730.9 Pascals and a temperature of 887 Kelvin. (b) Time series data with the Lorenz system output immediately being put in as the input perturbation in a system of 119970 Pascals and a temperature of 770 Kelvin.

In our search for interesting behavior, we also implemented a piston-pulsed wall in the reactor model. This causes a perturbation that drives the pressure in a sinusoidal manner as seen in Figure 5.3a. However, this method did not produce interesting behavior (Figure 5.3b).

We also designed a reactor model in which the inflow of the reactants uses the Lorenz system time series to set the mass inflow rate. This is intended to create chaotic behavior. Figure 5.4 demonstrates two reactions in which the O_2 mass inflow rate is set by the X time series output of the Lorenz system.

5.2 Time Series Reconstructions

The CCM technique is used to evaluate the degree to which different species population time series data from a hydrogen-oxygen reaction are coupled and thus can be reconstructed from other observed times series data from the system. We use the modified Pearson correlation coefficient ρ^* between the original time series $X(t)$ and the reconstructed time series $\hat{X}(t)$ in order to determine the effects of the application 4.2.7. A correlation of 1 implies there is a perfect relationship between our two variables. And so, we would like to achieve a correlation coefficient of 1 or a value close to that. Allow us to emphasize that we are not evaluating the correlation between both observed time series data. Rather, we are evaluating the goodness of fit of our reconstruction against our testing data of the original time series data.

Consider the time series data from a reactor model with a sinusoidal input perturbation of O_2 in a high pressure system. For this reconstruction, we consider species $X(t)$ to be H_2 and species $Y(t)$ to be O_2 . The direct application of CCM to the data set provided reconstructions of H_2 and O_2 , as seen in Figure 5.5, initially produced a fairly low Pearson coefficients of $\rho^* = -0.42963$ and $\rho^* = 0.093599$ for H_2 and O_2 respectively. As these values are far from 1, it implies a lack of relationship between our X and Y or more specially, that our reconstruction is a bad fit of our original time series data. However, we found that discarding the data before a steady state is reached in the reaction will result in less noise and a high correlation reconstruction. For the purposes of this data, we discard 50000 data points before splitting our data set. The reconstruction of the time series vastly improves as seen in Figure 5.6. We hypothesize that the initial system reaction before the steady state impacts the k -nearest neighbors aspect of the algorithm and thus accounts for variation that is not present in the remaining time series.




Figure 5.5: (a) Reconstruction and original time series of H_2 using the first half of the full data set as training data. This is a high pressure system of 126635 Pascals and temperature of 770 Kelvin. (b) Reconstruction and original time series of O_2 using the first half of the full data set as training data. This is a high pressure system of 126635 Pascals and temperature of 770 Kelvin.

We continued to apply the technique to various other data sets and observed continued success. For example, on a nonlinear time series data set from a model with low pressure, we see that despite increased noise, we are able to reconstruct O_2 . However, our reconstruction of H_2 seen in Figure 5.7b does not have as high of correlation. Since the species are closely related within the reaction network, we assume that the poor reconstruction quality stems from a sub-optimal choice of τ or E .

Consider a system in which the Lorenz output values are fed as a perturbation into our hydrogen-oxygen model. This system produced interesting results as seen in Figure 5.8. The H_2 reconstruction consistently provided variation to accurately represent the original time series. In this case, the reaction does not fully demonstrate the relationship between H_2 and O_2 as coupled. Again, this may imply a non-optimal τ or E for this reconstruction.




Figure 5.6: (a) Reconstruction and original time series of H_2 using the first half for training after discarding the first 500 seconds. This is a high pressure system of 126635 Pascals and temperature of 770 Kelvin. $\rho^* = 0.99166$. (b) Reconstruction and original time series of O_2 using the first half for training after discarding the first 500 seconds. This is a high pressure system of 126635 Pascals and temperature of 770 Kelvin. $\rho^* = 0.99817$.




Figure 5.7: (a) This plot is the reconstruction of H_2 and the original times series data of H_2 for a low pressure system of 7998 Pascals and temperature of 770 Kelvin. $\rho^* = 0.85668$. (b) This plot is the reconstruction of O_2 and the original times series data of O_2 for a low pressure system of 7998 Pascals and temperature of 770 Kelvin. $\rho^* = 0.88191$.

We reconstructed all ninety possible pairings of observables for a low pressure Lorenz system where oxygen was perturbed with Lorenz output. During this process, we found that pressure, O_2 , HO_2 , H_2 , O_2 , OH , O , and H reconstruct other species well, while temperature, H_2 and H_2O do not as shown by the Pearson correlation coefficients in Table 5.1.

We will begin discussing the species that constructed other species well. In Figure 5.8, the two reconstructions consisting of H_2 from HO_2 and H_2O from H_2, O_2 . Looking at Figure 5.8a we can see that $\text{HO}_2 + \text{H} \longleftrightarrow \text{H}_2 + \text{O}_2$, which has an energy release of about 5, has a high correlation coefficient of 0.99. This agrees with what we know of the close causal relationship between H_2 and HO_2 . Next, consider Figure 5.8b where the species are related

	Temperature	Pressure	H2	H	O	O2	OH	H2O	HO2	H2O2
Temperature	1	0.99096214	0.67713799	0.95102556	0.99248812	0.93373808	0.99678087	0.67789841	0.93747772	0.96162824
Pressure	0.94153736	1	0.56781647	0.90522579	0.9419758	0.98899	0.95535412	0.56881688	0.97801161	0.9315285
H2	0.4121609	0.9693973	1	0.84458815	0.97385279	0.99359548	0.96386741	0.99123315	0.9944405	0.96432019
H	0.84223121	0.9856382	0.66155668	1	0.97132263	0.93952494	0.97178042	0.66338885	0.94450914	0.98638581
O	0.9396447	0.9873007	0.7878483	0.8924923	1	0.9495714	0.995776	0.787724	0.9494535	0.9490924
O2	0.45576688	0.97792911	0.89434755	0.86979136	0.96283131	1	0.95717667	0.8942072	0.99305316	0.96670665
OH	0.96755346	0.98796414	0.75721249	0.90793334	0.99484636	0.94667778	1	0.75711415	0.94983763	0.95516958
H2O	0.41087832	0.96937296	0.99122481	0.84485469	0.97379792	0.99360393	0.96379534	1	0.99445187	0.96439178
HO2	0.35604134	0.97445267	0.93974018	0.86823654	0.96622923	0.99511197	0.95769995	0.93970326	1	0.97235453
H2O2	0.64726225	0.98174168	0.72530366	0.97682331	0.9578301	0.96335408	0.95631881	0.72638712	0.96999542	1

Table 5.1: This is a table of all the Pearson correlations from the all the possible reconstructions for a Lorenz-perturbed low pressure system. The x -axis value is the original time series species and the y -axis value is the reconstruction. Thus, the reconstruction of temperature from pressure has Pearson coefficient $\rho^* = 0.94153736$.

by the reactions $\text{H}_2\text{O}_2 + \text{H} \longleftrightarrow \text{H}_2\text{O} + \text{OH}$ and $\text{H}_2 + \text{O}_2 + \text{OH} \longleftrightarrow \text{H}_2\text{O} + \text{HO}_2$. This proves the notion that H_2O_2 and H_2O are directly related causing a high Pearson correlation of 0.96.




Figure 5.8: (a) Reconstruction of H_2 from the original time series data of HO_2 for a Lorenz low pressure system of 9730.9 Pascals and temperature of 887 Kevin. $\rho^* = 0.99$. (b) Reconstruction of H_2O from the original time series data of H_2O_2 for a Lorenz low pressure system of 9730.9 Pascals and temperature of 887 Kevin. $\rho^* = 0.96$.

Continuing our well-constructed species, we will analyze reconstructions of temperature from OH and O from pressure (Figure 5.9). Looking at Figure 5.9a we can see that $\text{H} + \text{OH} \longleftrightarrow \text{H}_2\text{O}$, which has a energy release of about 436 thus showing OH results in change of temperature, results in a high Pearson correlation of 0.99. Next in 5.9b we see that pressure and O must be related because of the high Pearson correlation of 0.98. We found that pressure reconstructs all species well, showing that a change in pressure affects the reactions. This is consistent with theory because at high or low pressures reactions might have little energy, leading to higher-quality reconstructions from pressure than from temperature.

Now we will be switching focus on the species that did not reconstruct others well (Figure 5.10). Here the two reconstructions consist of H from H_2 and O_2 from temperature. Looking at Figure 5.10a we can see that $\text{H}_2 \longleftrightarrow 2 \text{H}$, which has a energy gain of about 436 showing

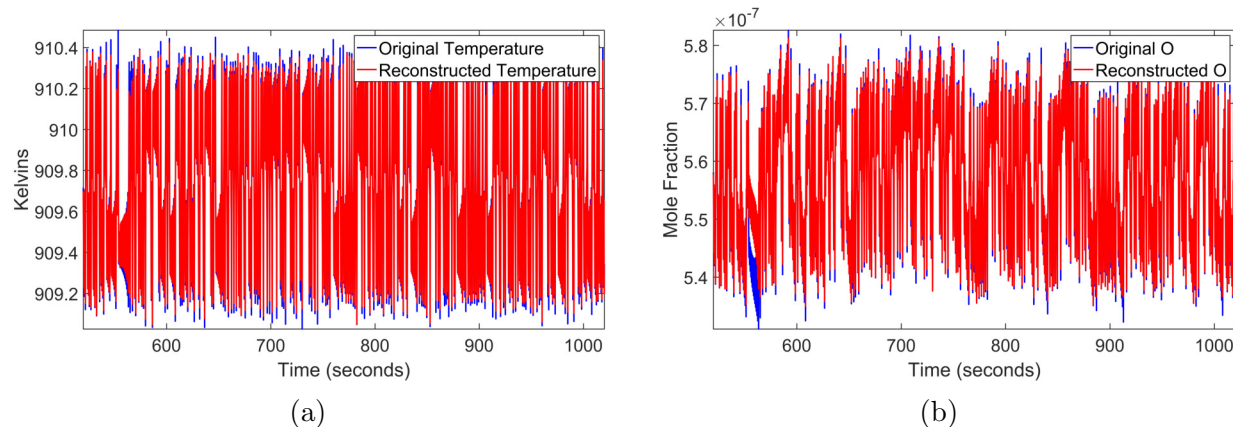


Figure 5.9: (a) Reconstruction of temperature from the original time series data of OH for a Lorenz low pressure system of 9730.9 Pascals and temperature of 887 Kelvin. $\rho^* = 0.99$. (b) Reconstruction of O from the original time series data of pressure for a Lorenz low pressure system of 9730.9 Pascals and temperature of 887 Kelvin. $\rho^* = 0.98$.

that the reaction is not favorable, results in a low Pearson correlation of 0.66. Next in Figure 5.10b we see $\text{HO}_2 + \text{H} \longleftrightarrow \text{H}_2 + \text{O}_2$, which as shown above has a release of 5, hence implying that the reverse will have a gain of 5. This reaction will effect temperature very slightly, so the reconstruction yields a low Pearson correlation of 0.45.




Figure 5.10: (a) Reconstruction of H from the original time series data of H_2 for a Lorenz low pressure system of 9730.9 Pascals and temperature of 887 Kelvin. There is a correlation of 0.66. (b) Reconstruction of O_2 from the original time series data of temperature for a Lorenz low pressure system of 9730.9 Pascals and temperature of 887 Kelvin. There is a correlation of 0.45.

Upon further investigation, we calculated the respective energy release for each reaction that did not produce a high quality reconstruction by temperature. We did this in order to see whether the species that were not reconstructed well by temperature participated in reactions whose energy/heat release is favorable. We will be focusing in on O_2 , HO_2 , H_2O_2 , H_2 and H_2O , since these were not reconstructed well by temperature.

The results discovered: O breaks H₂ bond turning H₂ into H⁺, which is not favorable for H₂ thus not releasing energy to affect temperature. H₂O remains H₂O throughout all the reactions that H₂O is reacting in, thus resulting in no energy release from H₂O. H₂O₂ reactions have the same bond energy of two hydrogen-oxygen bond and one oxygen-oxygen bond resulting in no energy release. O₂ reacts with H, where H breaks the oxygen double bond releasing the energy from the reaction, which not very favorable to begin with. HO₂ reacts with O or H, which enables the reaction thus resulting in the energy release for O or H. Thus, these species do not release energy into the reaction so temperature is not being affected in their reactions. This causes them to be reconstructed less effectively than other species from temperature.

Next, we analyzed why H₂ and H₂O cannot reconstruct other species well. H₂O does not affect other reactions resulting in the decreased ability to reconstruct other species. H₂ only reacts in two reactions and is created as a product in four thus it will not be able to reconstruct species that are not affected by H₂. Therefore, H₂ and H₂O cannot reconstruct as successfully as the other species.

We have been able to provide consistently favorable correlation coefficients for the reconstructions for most species. With the construction of bidirectional cross maps proving successful between closely related species in our reaction network, we move into a metric to infer parameters.

5.3 Parameter Inference

Using the methods detailed in Chapter 4, we were able to infer parameters of chemical reaction simulations in various setups. First, as a test case, we attempted to recover one parameter (the initial ratio of H₂ to O₂) by uniformly sampling the underlying space and evaluating test distributions against the reference solution (corresponding to H₂:O₂ = 2:1). We also evaluate the L_2 distance between PDFs X and Y given by

$$\sqrt{\sum_{\text{bins}} (X(\text{bin}) - Y(\text{bin}))^2}$$

Overall, the plot for W_1 is much better behaved than the plot for L_2 . In particular, W_1 has one true minimum (corresponding to the reference solution) in this plot, whereas L_2 has many local minima away from the true solution. The global minimum of both plots corresponds to the reference solution, but the presence of many false local minima in the L_2 metric would make this distance much worse than W_1 for gradient descent methods, which will be discussed in greater detail below. In this test and in the following test, we employ a “brute force” search, where we run simulations over some range of phase space and compute the distance for each test from the reference solution.

In the next test, we altered two of the Arrhenius parameters (A and n) for a certain reaction in our reaction network and evaluated the Wasserstein distance between each test solution and the reference solution to produce the following plot.




Figure 5.11: Parameter inference with the initial ratio of $\text{H}_2\text{:O}_2$ being altered to test for ground truth

As in the previous test case, the true global minimum corresponding to the reference solution was correctly identified, and the plot appears relatively smooth which bodes well for gradient descent methods; however, in this case, there was more than one local minimum in our data. As is evident from Figure 5.12, the parameter n has a much more dramatic effect on the Wasserstein distance, and there are points along fixed n where various values of A are local minima of W_1 . We suspect this is indicative of the fact that A affects the dynamics far less than n , which will lead to issues in inferring A and other parameters that have vanishing effects on dynamics.

Finally, we implemented a gradient descent search along the Wasserstein metric in order to infer a full set of parameters at once. The basic procedure is as standard: for a given point in phase space, we calculate the Wasserstein distance of the PDFs (with respect to a reference solution) corresponding to all of the neighbors of our point, and choose our next point in phase space to head in the direction of steepest descent. This method (as seen in Figure 5.13) converges to a full triple of Arrhenius parameters (A, n, E_a) much more quickly, with 10 – 15 iterations, than the brute force method. The drawback of this method is that it will not find the “true” global minimum, as identifying this exact point in phase space is impossible without the step size and initial point conspiring to allow this to happen.

Moreover, we have found, as above, that it is much harder to predict the parameter A than n or E_a . For example, for a given setup with reference parameters $(A, n, E_a) = (2.2 \times 10^{16}, -2, 0)$, gradient descent search found $(A, n, E_a) = (1.89 \times 10^{16}, -1.98025, 0)$, where the result for n is much closer than the result for A . These results serve as proof of concept for gradient descent along the Wasserstein response curve, and much work should be done in this area to fine-tune this method and generalize these results.




Figure 5.12: Wasserstein distance of test PDFs against the two parameters being altered, A and n

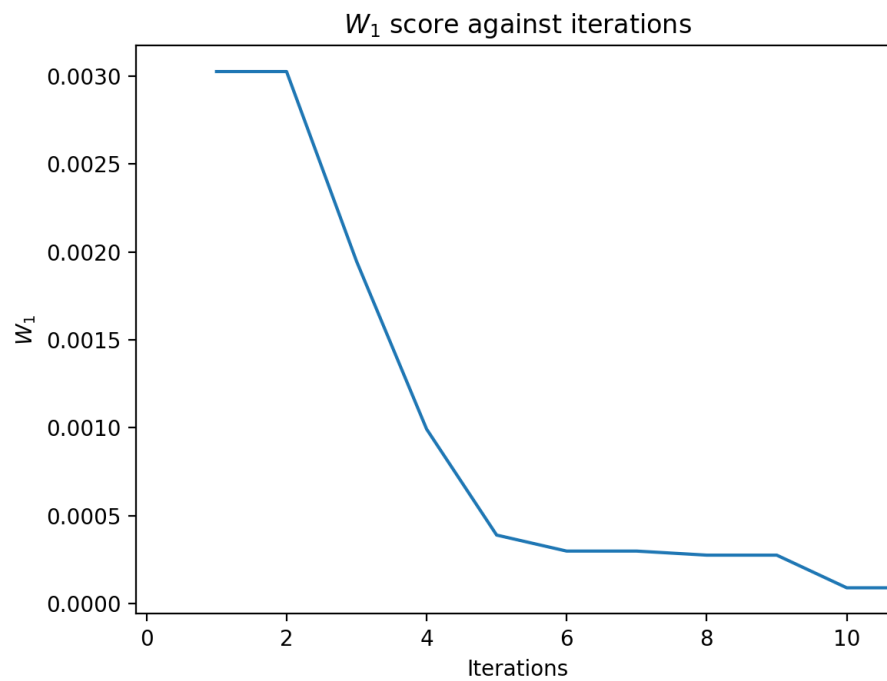


Figure 5.13: Convergence of gradient descent algorithm

Chapter 6

Conclusions and Discussion

The methods presented in this report provide AFRL with new and improved tools for characterizing and calibrating chemical reactions.

We described relevant setups for simulating hydrogen-oxygen combustion in a zero-dimensional CSTR model using the *Cantera* software package, focusing on the initial conditions necessary to “push” the system off equilibrium. While our attempts to perturb the mass inflow and pressure boundary conditions with a variety of sinusoidal and Lorenz-output perturbations did not succeed in pulling apart the finite rate chemistry timescales, those results illuminate the limitations of *Cantera* simulations.

More importantly, we described a method for attractor reconstruction inspired by the CCM algorithm, and demonstrated the success of this method for hydrogen-oxygen chemistry. We were able to provide successful reconstructions from pressure, O_2 , HO_2 , H_2O_2 , OH , O , and H , while temperature, H_2 and H_2O did not provide as successful of reconstructions. This was explained by having a direct favorable reaction from the original time series species to the reconstructed species. Thus, before reconstructing species it is important to make sure the species are related before obtaining a reconstruction. However, pressure being able to reconstruct every species successfully is a very interesting find that we would like to look into further to see if this is always the case. If so, this could be a great way to predict species from if it is always consistent, otherwise looking for the species that is the most dominant in the reactions would be the best species to use for reconstruction.

The existence of convergent cross maps between shadow manifolds of different species, as indicated by our successful cross-species reconstructions, confirms that shadow manifolds of individual species capture the information necessary to describe all observations of the reactor system. Consequently, we can propose a method for parameter inference that relies only on shadow manifolds of single species instead of requiring knowledge of the entire phase space dynamics. The optimization metric for parameter inference that we describe is further innovative in its use of the Wasserstein distance. We first create shadow manifolds from the trial and reference solutions, bin them to create a PDF, and then find the Wasserstein distance between the resulting SM-PDFs.

We demonstrated that the Wasserstein distance yields smoother response surfaces relative to other standard distance metrics. In addition, the Wasserstein distance provides meaningful notions of “closer” and “farther” even for SM-PDFs that have few points in common. Our method for parameter inference is robust to noise, thus allowing for accurate inference of reaction rate coefficients from noisy, limited time series data.





We concluded the project by confirming the success of our method for parameter inference by recovering several parameters from our simulated reference solution: the inflow ratio of hydrogen, H_2 , to oxygen, O_2 ; the fitted rate temperature constant n of the Arrhenius equation; and the Arrhenius equation pre-exponential factor A . Moreover, we were able to infer two reaction rate coefficients at once, and conjecture that the empirical parameter inference method we describe can successfully infer even larger sets of parameters when given sufficient computation resources.











Although we employed the best available algorithms for estimating a minimum embedding dimension E , finding an optimal time lag τ , choosing bin sizes when creating probability distribution functions from shadow manifolds, and computing the Wasserstein distance, no perfect algorithms exist for any of these steps. As described in the body of the report, each choice of algorithm or method involves some tradeoff between factors such as computational speed, robustness to noise, scalability, and universality of application. These choices might need to be revisited in future work once improved numerical techniques become available for the information theoretic and topological methods that our algorithms rely on.











Appendix A











Complete Reconstructions











These are all the possible reconstructions of a low pressure Lorenz perturbation. Each graph is paired with its reverse reconstruction. Thus, the first two consist of Pressure from Temperature and Temperature from Pressure.



















































Appendix B

Abbreviations

AFRL: The Air Force Research Laboratory. A scientific research organization operated by the U.S. Air Force.

CCM: Convergent Cross Mapping. A technique for distinguishing causality from correlation in time series data.

CSTR: Continuously stirred tank reactor. A single zero-dimensional Cantera reactor with an inlet, an outlet, and constant volume.

HET: Hall-effect thruster.

IPAM: The Institute for Pure and Applied Mathematics. A National Science Foundation Mathematical Sciences Institute, located at UCLA.

MFC: Mass flow controller. A valve that regulates the amount of mass flowing into a reactor from a reservoir.

PDF: Probability distribution function.

RIPS: Research in Industrial Projects for Students. A regular summer program at IPAM, in which teams of undergraduate students work on research projects proposed by sponsors from industry or the public sector.

SM: Shadow manifold.

SM-PDF: Shadow manifold probability distribution function.

UCLA: The University of California, Los Angeles. A public research university located in Los Angeles, California and operated by the University of California system.

Selected Bibliography Including Cited Works

- [1] M. ABURDENE AND J. DORBAND, *Parallel computation of discrete legendre transforms*, in 1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings, IEEE, May 1996.
- [2] L. CAO, *Practical method for determining the minimum embedding dimension of a scalar time series*, Physica D: Nonlinear Phenomena, 110 (1997), p. 43–50.
- [3] L. DAS, *Hydrogen-oxygen reaction mechanism and its implication to hydrogen engine combustion*, International Journal of Hydrogen Energy, 21 (1996), p. 703–715.
- [4] C. DEVELOPERS, *Cantera*. <https://cantera.org>, Aug 2018. Version 2.4.0.
- [5] D. ECKHARDT, J. KOO, R. MARTIN, M. HOLMES, AND K. HARA, *Spatiotemporal data fusion and manifold reconstruction in hall thrusters*, Plasma Sources Science and Technology, 28 (2019), p. 45005.
- [6] R. FLAMARY AND N. COURTY, *Pot python optimal transport library*, 2017.
- [7] A. M. FRASER AND H. L. SWINNEY, *Independent coordinates for strange attractors from mutual information*, Physical Review A, 33 (1986), p. 1134–1140.
- [8] A. GEORGE, R. CHAKMAK, C. CHAN, G. DIMAND, C. FALCON, D. ECKHARDT, AND R. MARTIN, *Decomposition of nonlinear system dynamics into multiple time scales*, August 2018. RIPS 2018 AFRL Final Report.
- [9] J. F. GIBSON, J. DOYNE FARMER, M. CASDAGLI, AND S. EUBANK, *An analytic approach to practical state space reconstruction*, Physica D: Nonlinear Phenomena, 57 (1992), p. 1–30.
- [10] I. GLASSMAN, R. A. YETTER, AND N. G. GLUMAC, *Combustion*, Elsevier, 5 ed., 2014.
- [11] C. W. J. GRANGER, *Investigating causal relations by econometric models and cross-spectral methods*, Econometrica, 37 (1969), p. 424.
- [12] C. M. GREVE, K. HARA, R. S. MARTIN, D. Q. ECKHARDT, AND J. W. KOO, *A data-driven approach to model calibration for nonlinear dynamical systems*, Journal of Applied Physics, 125 (2019), p. 244901.

- [13] E. N. LORENZ, *Deterministic nonperiodic flow*, Journal of the Atmospheric Sciences, 20 (1963), p. 130–141.
- [14] R. MARTIN, J. KOO, AND D. ECKHARDT, *Impact of embedding view on cross mapping convergence*, 2019.
- [15] J. M. MCCracken AND R. S. WEIGEL, *Convergent cross-mapping and pairwise asymmetric inference*, Physical Review E, 90 (2014).
- [16] L. M. PECORA, L. MONIZ, J. NICHOLS, AND T. L. CARROLL, *A unified approach to attractor reconstruction*, AIP Publishing, (1970).
- [17] Y. ROBIN, P. YIOU, AND P. NAVEAU, *Detecting changes in forced climate attractors with wasserstein distance*, Nonlinear Processes in Geophysics, 24 (2017), p. 393–405.
- [18] G. SUGIHARA, R. MAY, H. YE, C.-H. HSIEH, E. DEYLE, M. FOGARTY, AND S. MUNCH, *Detecting causality in complex ecosystems*, Science, 338 (2012), p. 496–500.
- [19] F. TAKENS, *Detecting strange attractors in turbulence*, Springer Berlin Heidelberg, 1981, p. 366–381.
- [20] X. XU, X. LIU, AND X. CHEN, *The cao method for determining the minimum embedding dimension of sea clutter*, in 2006 CIE International Conference on Radar, IEEE, Oct 2006.
- [21] H. YE, E. R. DEYLE, L. J. GILARRANZ, AND G. SUGIHARA, *Distinguishing time-delayed causal interactions using convergent cross mapping*, Scientific Reports, 5 (2015).
- [22] H. YE AND G. SUGIHARA, *Information leverage in interconnected ecosystems: Overcoming the curse of dimensionality*, Science, 353 (2016), p. 922–925.