Artificial Intelligence Exploration (AIE)
DARPA-PA-18-02-02
Artificial Intelligence Research Associate (AIRA)

Milestone Report #10
Reporting Period: 21 May 2020 – 20 June 2020

Agreement No. HR00111990028

Date Prepared
20 June 2020

Prepared by:
PI: Krishna Garikipati, Professor
PI: Xun Huan, Assistant Professor
Karthik Duraisamy, Associate Professor
Vikram Gavini, Associate Professor
Alex Gorodetsky, Assistant Professor
Emmanuelle Marquis, Associate Professor
University of Michigan
2350 Hayward
Ann Arbor, MI 48109-2125
734-936-0414
krishna@umich.edu

Prepared for:
DARPA Program Manager (PM):
Dr. Jiangying Zhou
Phone: 703-526-2797
jiangying.zhou@darpa.mil

Agreements Officer (AO):
Dr. Peter Donaghue
Phone: 703-526-2705
desmond.donaghue@darpa.mil

Agreements Officers Representative (AOR):
Dr. Daniel M. Sparkman, AFRL/RXCA
Phone: 937-255-1340
daniel.sparkman.1@us.af.mil

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From – To)*
20-06-2020	Phase 2 Final Report	20-09-2019 – 20-06-2020

4. TITLE AND SUBTITLE
ARTIFICIAL INTELLIGENCE EXPLORATION (AIE)
DARPA-PA-18-02-02
Artificial Intelligence Research Associate (AIRA)
Milestone 10 Report
Artificial Intelligence Guided Multi-scale Multi-physics Framework for Discovering Complex Emergent Material Phenomena (AIMM Discover)

5a. CONTRACT NUMBER
HR00111990028

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Krishna Garikipati, Xun Huan, Karthik Duraisamy, Vikram Gavini, Alex Gorodetsky, Emmanuelle Marquis

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Michigan
3003 S. State Street
Ann Arbor, MI 48109-1274

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Defense Advanced Research Projects Agency
Defense Sciences Office
675 North Randolph Street
Arlington, VA 22203-2114

10. SPONSOR/MONITOR'S ACRONYM(S)
DARPA/DSO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Distribution Statement A: Approved for public release.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
We propose a Bayesian framework to develop new machine learning and operator inference methods to aid the discovery of physical phenomena and the prediction of material properties and responses. We specifically target the challenges in material physics associated with systematic attempts to (a) abstract complexity from a hierarchy of scales into predictive model forms and (b) delineate mechanisms of coupled materials physics. Our project develops the following tasks that unite artificial intelligence (AI) with the discovery of emergent physics: (1) Scale bridging from quantum mechanics to continuum PDEs. (2) Physics discovery via system identification and operator inference. (3) Bayesian inference and uncertainty quantification for learning from data and quantifying predictive quality. (4) Optimal experimental design for intelligent data acquisition and management to achieve efficient high-level learning. The first two tasks address important application-specific problems and the next two tasks provide the AI interface.

15. SUBJECT TERMS
Artificial Intelligence, Algorithms, Analytics, Data

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON *(monitor)*
a. REPORT	b. ABSTRACT	c. THIS PAGE			Jiangying Zhou
Unclassified	Unclassified	Unclassified			19b. TELEPHONE NUBER *(include area code)* (703) 526-2797

Contents

Acknowledgements

We would like to acknowledge funding from the DARPA Artificial Intelligence Exploration (AIE) Artificial Intelligence Research Associate (AIRA) program managed by Dr. Jiangying Zhou; insightful feedback to our milestone reports from Dr. Daniel Sparkman; support and organization from Drs. Zachary Lapin, Lee Pele, and Yardley Sands; and valuable discussions and interactions with many other researchers across DARPA, AFRL, and other teams on the AIRA program. The success of this project was made possible by the hard work from the students, postdocs, and research scientists on the team: Doruk Aksoy, James Duvall, Nicholas Galioto, Thomas Maulbeck, Wanggang Shen, Drs. Yaser Afshar, Krishnendu Ghosh, Bikash Kanungo, Greg Teichert, Zhenlin Wang, and Bobbie Wu.

1 SUMMARY

We develop a framework of new machine learning and operator inference methods to aid the discovery of physical phenomena and the prediction of material properties and responses. We specifically target the challenges in material physics associated with systematic attempts to *(a) abstract complexity from a hierarchy of scales into predictive model forms* and *(b) delineate mechanisms of coupled materials physics.*

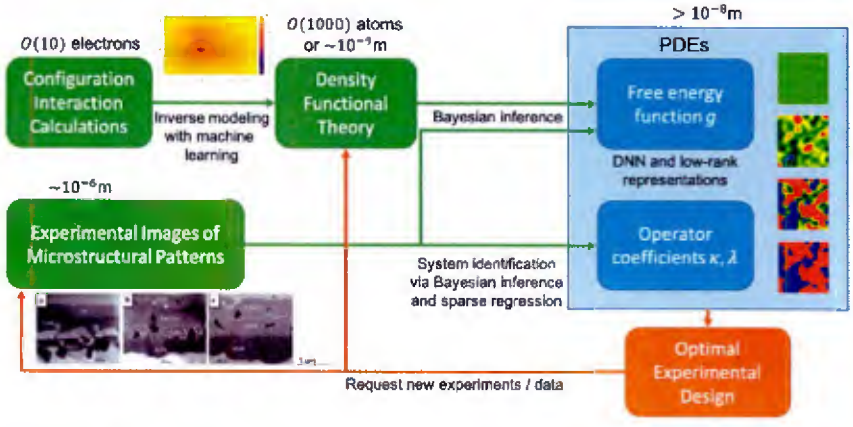

Figure 1: *Project overview, integration, and flow of information.*

We address problems characterized by the development of microstructural patterns in the distribution of atomic species, compounds and phases in materials during processing and operation. The formation and evolution of such patterns is often driven by phase transitions and can be studied with diverse microscopy techniques at scales of 10^{-9}–10^{-4} m. The microstructure plays a central role in governing a wide range of material properties. Currently no formal AI framework is used for material physics, and existing works are confined to single scale, low-dimensional, and non-probabilistic settings. We will address these, and the AI framework proposed here will not only catalyze systematic, efficient, and automated discovery of the physics of materials, but also have a far-reaching impact in the modeling of other multiscale/multiphysics problems.

Our project develops the following tasks that unite AI with the discovery of emergent physics: (1) Scale bridging from quantum mechanics to continuum PDEs. (2) Physics discovery via system identification and operator inference. (3) Bayesian inference and uncertainty quantification for learning from data and quantifying predictive quality. (4) Optimal experimental design for intelligent data acquisition and management to achieve efficient high-level learning. The first two tasks address important application-specific problems and the next two tasks provide the AI interface. Figure 1 shows the integration of our themes.

2 APPROACH, ARCHITECTURES, ALGORITHMS

2.1 Task 1: Data driven exchange-correlation functional in DFT

2.1.1 Significance

Density functional theory (DFT) has remained the most popular electronic structure method, owing to its remarkable balance of speed and accuracy. While in principle an exact theory, in practice DFT has remained far from exact due to the unavailability of the exact exchange-correlation

functional. To elaborate, the exchange-correlation functional (XC), described as an energy term (E_{xc}) or a potential term (v_{xc}) encapsulate the quantum many-electron interactions. Both v_{xc} and E_{xc} are known to be unique functionals of the electron density, $\rho(\boldsymbol{r})$, and are related to each other as $v_{\text{xc}}[\rho(\boldsymbol{r})] = \frac{\delta E_{\text{xc}}[\rho(\boldsymbol{r})]}{\delta \rho(\boldsymbol{r})}$. However, their exact forms are unknown, necessitating the use of approximations. Naturally, the development of accurate XC functionals has, for long, remained the grand challenge in DFT.

To this end, we propose to construct these functionals in a data-driven fashion, using accurate ground-state densities, $\rho(\boldsymbol{r})$, from wavefunction-based methods (e.g., quantum Monte-Carlo, full Configuration Interaction). The wavefunction-based methods, owing to their exponential computational complexity, remain untenable beyond a few tens of electrons. However, one can utilize the accurate ground-state densities from these small systems that are within the reach of the wavefunction-based methods to develop a machine-learned model for $v_{\text{xc}}[\rho(\boldsymbol{r})]$ and $E_{\text{xc}}[\rho(\boldsymbol{r})]$. This involves two distinct steps—(i) generation of training data set comprising of $\{\rho(\boldsymbol{r}), v_{\text{xc}}(\boldsymbol{r})\}$ pairs, (ii) use of machine learning algorithms to learn the functional form of $v_{\text{xc}}[\rho(\boldsymbol{r})]$ (and $E_{\text{xc}}[\rho(\boldsymbol{r})]$), conforming to the exact conditions.

2.1.2 Formulation for *inverse* DFT

As a first step towards the aforementioned goal of machine-learned exchange-correlation functional, we have successfully addressed the challenge of evaluating the exact $v_{\text{xc}}(\boldsymbol{r})$ corresponding to a given $\rho(\boldsymbol{r})$. This is otherwise known as the *inverse* DFT problem. We underline that, heretofore, *this had remained an open challenge owing lack of accurate and systematically convergent numerical techniques.* We have recently provided a robust and systematically convergent scheme to solve the inverse DFT problem, employing finite element basis [12]. In particular, we cast the inverse DFT problem of finding the $v_{\text{xc}}(\boldsymbol{r})$ that yields a given density $\rho_{\text{data}}(\boldsymbol{r})$ as the following partial differential equation (PDE) constrained optimization:

$$\arg\min_{v_{\text{xc}}(\boldsymbol{r})} \int w(\boldsymbol{r}) \left(\rho_{\text{data}}(\boldsymbol{r}) - \rho(\boldsymbol{r})\right)^2 d\boldsymbol{r}\,, \tag{1}$$

subject to

$$\left(-\frac{1}{2}\nabla^2 + v_{\text{ext}}(\boldsymbol{r}) + v_{\text{H}}(\boldsymbol{r}) + v_{\text{xc}}(\boldsymbol{r})\right)\psi_i = \epsilon_i \psi_i\,, \tag{2}$$

$$\int |\psi_i(\boldsymbol{r})|^2 \, d\boldsymbol{r} = 1\,. \tag{3}$$

In the above equation, $w(\boldsymbol{r})$ is an appropriately chosen weight to expedite convergence; $v_{\text{ext}}(\boldsymbol{r})$ represents the nuclear potential; $v_{\text{H}}(\boldsymbol{r})$ is the Hartree potential corresponding to $\rho_{\text{data}}(\boldsymbol{r})$; and ψ_i and ϵ_i denote the Kohn-Sham orbitals and eigenvalues, respectively. For simplicity, we restrict ourselves to only closed-shell systems, and hence, the Kohn-Sham density $\rho(\boldsymbol{r}) = 2\sum_{i=1}^{N_e/2} |\psi_i(\boldsymbol{r})|^2$. Equivalently, the above PDE-constrained optimization can be solved by minimizing the following Lagrangian,

$$\begin{aligned}\mathcal{L}\left(v_{\text{xc}}, \{\psi_i\}, \{p_i\}, \{\epsilon_i\}, \{\mu_i\}\right) = \int w(\boldsymbol{r}) \left(\rho_{\text{data}}(\boldsymbol{r}) - \rho(\boldsymbol{r})\right)^2 d\boldsymbol{r} + \sum_{i=1}^{N_e/2} \int p_i(\boldsymbol{r}) \left(\hat{H} - \epsilon_i\right)\psi_i \, d\boldsymbol{r} \\ + \sum_{i=1}^{N_e/2} \mu_i \left(\int |\psi_i(\boldsymbol{r})|^2 \, d\boldsymbol{r} - 1\right),\end{aligned} \tag{4}$$

with respect to all its constituent variables—p_i, μ_i, ψ_i, ϵ_i and $v_{\rm xc}$. In the above equation, $\hat{H} = -\frac{1}{2}\nabla^2 + v_{\rm ext}(\boldsymbol{r}) + v_{\rm H}(\boldsymbol{r}) + v_{\rm xc}(\boldsymbol{r})$ is the Kohn-Sham Hamiltonian; p_i is the adjoint function which enforces the Kohn-Sham eigenvalue equation corresponding to ψ_i; and μ_i is the Lagrange multiplier corresponding to the normality condition of ψ_i. Having solved the above optimality of $\mathcal{L}$ with respect to he p_i, μ_i, ψ_i, and ϵ_i, the variation (gradient) of $\mathcal{L}$ with respect to $v_{\rm xc}$ is given by

$$\frac{\delta \mathcal{L}}{\delta v_{\rm xc}} = \sum_{i=1}^{N_e/2} p_i \psi_i \ . \tag{5}$$

This constitutes the central equation for updating $v_{\rm xc}(\boldsymbol{r})$ via any gradient based optimization technique.

2.1.3 System identification for XC functional

As a first step towards the construction of a data driven XC functional, we have tried to model the exchange-correlation functional ($E_{\rm xc}[\rho]$) through system identification, using the exact $\left\{\rho^{(I)}, v_{\rm xc}^{(I)}\right\}$ pairs from our inverse DFT calculations as the training data. While an accurate modeling of $E_{\rm xc}[\rho]$ warrants more sophisticated machine-learning approached (e.g., deep neural networks), as a first step, we have attempted to model $E_{\rm xc}[\rho]$ through system identification for its good balance of simplicity and accuracy. Specifically, we have modeled $E_{\rm xc}$ as a range-separated local functional of ρ and $|\nabla\rho|$, consisting of a local density approximation (LDA) [18] near the nuclei and a generalized gradient approximation (GGA) away from the nuclei [16]. The choice of such a range-separated form is motivated by the fact that the LDA provides a more physically meaningful $v_{\rm xc}$ near the nuclei, as compared to the GGA. On the other hand, the GGA affords the flexibility to fit the $v_{\rm xc}$ well in the interstitial regions. We remark that there are several flavors of GGA to choose from. In our preliminary studies, we have employed an empirical version of GGA known as the second-order GGA (SOGGA), owing to its flexible form. Thus, our range-separated $E_{\rm xc}$ takes the following form

$$\begin{aligned} E_{\rm xc}\left[\rho, |\nabla\rho|\right] &= E_{\rm x}\left[\rho, |\nabla\rho|\right] + E_{\rm c}\left[\rho, |\nabla\rho|\right] \\ &:= \int d\boldsymbol{r} w(\boldsymbol{r})\rho(\boldsymbol{r})\epsilon_{\rm x}^{\rm LDA}[\rho] + \int d\boldsymbol{r}(1-w(\boldsymbol{r}))\rho(\boldsymbol{r})\epsilon_{\rm x}^{\rm LDA}[\rho]F_{\rm x}(s) + \\ &\int d\boldsymbol{r} w(\boldsymbol{r})\rho(\boldsymbol{r})\epsilon_{\rm c}^{\rm LDA}[\rho] + \int d\boldsymbol{r}(1-w(\boldsymbol{r}))\rho(\boldsymbol{r})\epsilon_{\rm c}^{\rm LDA}[\rho]F_{\rm c}(\rho, s) \end{aligned} \tag{6}$$

where $\epsilon_{\rm x}^{\rm LDA}[\rho]$ and $\epsilon_{\rm c}^{\rm LDA}[\rho]$ are the local density approximation (LDA) based exchange and correlation energy densities, respectively, and $w(\boldsymbol{r})$ is a smooth cutoff function that transitions from 1 to 0 as one proceeds from a nucleus. The variable $s = \frac{1}{2(3\pi^2)^{1/3}}\frac{|\nabla\rho|}{\rho^{4/3}}$ is the dimensionless (in atomic units) reduced density-gradient. The $F_{\rm x}(s)$ and $F_{\rm c}(\rho, s)$ are analytical functions of ρ and s. We refer to Report7a for the details of the system identification problem.

2.1.4 Continuous convolution neural network for XC functional in DFT

Motivation: The above system identification approach, although being an interesting direction to a data driven XC functional, failed to attain the desired chemical accuracy. We anticipate two plausible reasons for this inaccuracy:

1. Intuitively, the range-separated form should provide a more accurate description for the $v_{\rm xc}$. However, it can induce unphysical behavior in the $v_{\rm xc}$ transition region, thereby generating inaccurate ρ in a self-consistent field iteration.

2. The SOGGA form was originally constructed from the perspective of predicting better energy ($E_{\rm xc}$) rather than producing better potential ($v_{\rm xc}$). The LDA form, on the other hand, was designed to reproduce exact results in the uniform electron gas limit. Thus, there might be inherent deficiencies in the SOGGA and the LDA form to simultaneously maintain a good balance of accurate $v_{\rm xc}$ and $E_{\rm xc}$, especially for ρ that are far from uniform. Moreover, these deficiencies can, possibly, get compounded with the range-separated mixing of SOGGA and LDA.

Thus, to this end, we intend to model $E_{\rm xc}$ as a *continuous filter* convolution neural network (CNN), so as to seamlessly incorporate the desired non-locality. This is an ongoing effort. Below, we present its formulation.

We remark that, although CNNs are naturally suited to include non-local effects, the extension of the conventional CNN architectures to modeling $E_{\rm xc}[\rho]$ remains far from trivial. This non-triviality arises from the diversity in the nature of the ρ (input to the model) used in the training data as well as test data. To elaborate, the ρ, in its extent and salient features, are characterized by the position and nature of the atoms, which in turn vary significantly between systems. To complicate things further, the ρ is represented using a basis that is system dependent. The same diversity applies to $v_{\rm xc}$ (output of the model) as well. This is in sharp contrast with the usual framework of CNNs, wherein the input and the output vector have a common representation across samples (same size vectors or same basis functions). To this end, we intend to develop a continuous CNN framework (as opposed to the conventional discrete CNN), so as to treat the diverse forms and representations of ρ's and $v_{\rm xc}$'s in a unified manner.The key ideas of our formulation are (refer to Report9 for details):

1. Model ${E_{\rm xc}}^{\rm ML} = \int e_{\rm xc}^{\rm ML}[\rho]\, d\boldsymbol{r}$ with $e_{\rm xc}^{\rm ML}$ as a multi-layered CNN.
2. Use a continuous filter $K\left(\boldsymbol{r};\boldsymbol{\theta}^{(l)}\right)$ at each layer l, parameterized by the learnable parameters $\boldsymbol{\theta}^{(l)}$.
3. One possible parameterization of $K\left(\boldsymbol{r};\boldsymbol{\theta}^{(l)}\right)$ is through a linear combination of a basis set, i.e.,
$$K\left(\boldsymbol{r};\boldsymbol{\theta}^{(l)}\right) = \sum_{j=1}^{N} \theta_j^{(l)} u_j(\boldsymbol{r})\,. \tag{7}$$
4. Use a complete and preferably separable basis such as planewaves (PW) (i.e., $u_j(\boldsymbol{r}) = e^{i\boldsymbol{k}_j\cdot\boldsymbol{r}}$).
5. Use Tucker tensor decomposition of the integrands in any convolution integral so as to decompose the 3D integration into three independent 1D integral.

2.2 Task 2: Scale bridging and physics discovery

2.2.1 Integrable Deep Neural Networks and active learning for multiscale discovery of physics

The evolution of the microstructure in materials systems is primarily governed by the free energies of the various phases. The free energies of the various phases, also accounting for the concentration of the various constituent elements in the materials system, are governed by quantum mechanical interactions and can be computed using electronic structure calculations. Density functional theory (DFT), which is the workhorse of electronic structure calculations, while successful in predicting a

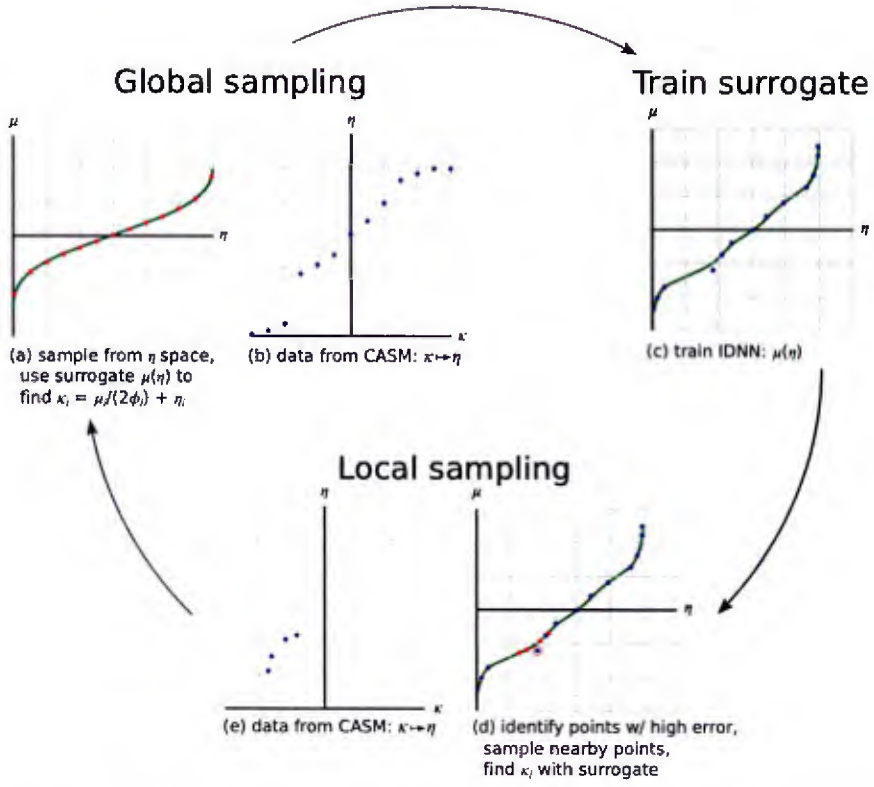

Figure 2: This active learning workflow combines local and global sampling, calculate data points using CASM, and trains an IDNN as the surrogate model.

wide range of material properties, is computationally expensive owing to the cubic computational complexity and restricts the accessible system sizes, and the number of calculations that can be performed.

We have developed a data driven framework, where we use DFT calculations on representative systems to construct cluster Hamiltonian expansions, which in turn are used to compute the free energy of a prescribed phase–order parameters defining the arrangements of atoms over a crystal– and concentration of the constituent elements. The DFT calculations that inform the cluster Hamiltonian expansions are computed using a generalized gradient approximation based exchange-correlation functional.

Continuum level phase field models rely on a free energy density function to model the evolution of phases in a material. DFT calculations and statistical mechanics can predict the chemical potential (i.e. the derivative of the free energy) of a system for particular values of material composition and order parameters. Specifically, we use the CASM code [1], a first-principles statistical mechanics code based on cluster Hamiltonian expansions, to compute the free energies of the multi-component systems. The data calculated using CASM are used to train a neural network description of the free energy. Since the chemical potential data describe the derivatives of the free energy, we train using an integrable deep neural network (IDNN), which can be analytically integrated to recover the free energy function. This neural network is initially used to guide the active learning algorithm to populate additional CASM runs in areas where more data are needed. After the workflow has converged, the free energy neural network is used is the continuum level phase field models.

For a system with n order parameters, the CASM framework takes inputs temperature T and bias parameters κ_i and ϕ_i, and returns the order parameters η_i, $i = 1, \ldots, n$. The values of

chemical potentials $\mu_i := \partial g/\partial \eta_i$ are then found through the relationship $\partial g/\partial \eta_i = -2\phi_i(\eta_i - \kappa_i)$ [15], and finally g can be recovered (corresponding to this $\boldsymbol{\eta}$) by performing an integration. Our initial strategy for learning $g(\boldsymbol{\eta})$ employs a "shotgun" approach, where we seek data that are well-sampled in the $\boldsymbol{\eta}$ space. This is done by fixing T and ϕ_i, while the η_i value are (initially) generated with a global sampling method. The κ_i values are initially obtained from the sampled η_i values by using the expression for the ideal solution chemical potential. However, the actual mapping between κ_i and η_i is more complex, and so the initial sampling of κ_i may yield η_i regions of poor coverage, particularly in regions of phase transition. To target additional data in those regions, a function (i.e. the IDNN) is trained on the current data to produce an *inverse* mapping $\boldsymbol{\eta} \to \boldsymbol{\kappa}$. Data points with the highest validation error are identified. These points likely lie in or near the unstable regions, and so additional points in their vicinity are used as input to CASM (local, error-based sampling). Even while this active learning progresses, global sampling (exploration) based on a Sobol′ sequence continues throughout the workflow (see Figure 2). This process is repeated until convergence is met.

2.2.2 Variational System Identification for physics discovery from data on evolving microstructures

The macroscopic experimental images produce quantities of interest (QoIs) that are dependent on the free energy through the governing PDEs. Hence it is particularly compelling to determine the governing partial differential equations (PDEs), because knowledge of the PDE directly translates to deep insights to the physics. With this approach to system identification, we are able to predict the dynamics under different initial and boundary conditions using the identified PDEs.

In Phase I we have presented our approach, called Varaitional System Identification (VSI), to identify systems of PDEs using full field data. However, the vast majority of microscopy techniques for evolving microstructure in a given material system deliver micrographs of pattern evolution over domains that bear no relation with each other at different time instants. The temporal resolution can rarely capture the fastest time scales that dominate the early dynamics, and noise abounds. Furthermore, data for evolution of the same phenomenon in a material system may well be obtained from different physical specimens. To address the challenges presented by image data on microstructures in materials physics, we extend our variational system identification methods in Phase II. Considering many statistical measurements, e.g. the moments of the concentration fields, can be well approximated using local data, moments can be recovered by specific choices of weighting function. It then allows a two-stage approach to separate the identification of algebraic and differential operators. VSI can thus be extended to incorporate spatially unrelated (and non-overlapping) and sparse information.

We have recently demonstrated that VSI recognizes that the evolution of material microstructure is not governed by second-order dynamics, but is first-order in time. This distinguishes pattern formation from Newtonian dynamics, which is second-order in time, and represents re-discovery of fundamental physical laws of nature. We have tested our approaches on synthetic data generated from a variety of first order dynamics.

In Phase II, we have focused on further developing and applying VSI to real data. The processing of raw imaging obtained previously on the Ag-Al alloy system [14] is discussed next in Section 2.2.3

2.2.3 Experimental Image Processing: Summary of Developed Tools

Preprocessing is required to leverage data obtained from experimental images in learning or system identification algorithms. Typical experimental images are contaminated with noise and contain

other imperfections which makes the extraction of statistical QOIs or other measures difficult. To address this, an image processing pipeline was developed using open source tools to prepare experimental images for image segmentation algorithms or other analysis.

Of particular interest here are images of two phase systems containing precipitates nucleated and grown in a background matrix. Several different systems are considered, including Ag-Al and Ti-20 at% Nb. These systems contain a wide variety of precipitate shapes, from elongated high aspect ratio precipitates to those nearly circular in shape. The preprocessing pipeline consists of denoising via Gaussian, average, and/or non-local means filtering, as well as adaptive histogram equalization to correct variations in brightness and contrast across the images. Image binarizations are obtained by applying Otsu's method to determine the brightness threshold which results in two classes with the greatest inter-class variance. The binarizations are used to generate Euclidean distance maps, which are used in generating precipitate markers and as input to a watershed transform image segmentation routine.

A watershed transform algorithm was selected as it can naturally handle the variety of precipitate shapes seen. However, watershed algorithms are marker controlled and open source methods for generating the markers proved to be deficient. They would either mark each precipitate multiple times or miss certain precipitates all together. To address this, a scheme which first finds a large number of markers and then eliminates redundant markers was developed. The method is based around searching for local maxima in the distance map. This occurs near the middle of the precipitates and the method searches for locations where the x or y component of the distance map gradient passes through zero. Each marker is compared against the others by examining the histogram of the pixels in a region connecting them. If the histogram is sufficiently bimodal, this indicates that they lie on separate precipitates and both are retained. If the histogram is not bimodal, this indicates that they lie on the same precipitate and the marker with the greater distance map value is retained. Using these preprocessing routines on the experimental images yielded reasonable segmentation results on some datasets and extraction of QOI's, such as precipitate size distribution, which can be used in system identification algorithms. These methods also allowed for the approximation of the spatial concentration field from experimental images which was used to identify the operators in the PDEs using the data of the Ag-Al system. An example of the concentration field and the extracted interface profiles are shown in Figure 3. Extraction of these operators then allowed the forward simulation of the identified Allen-Cahn equation with Fickian diffusion system, using the experimental data as the initial condition.

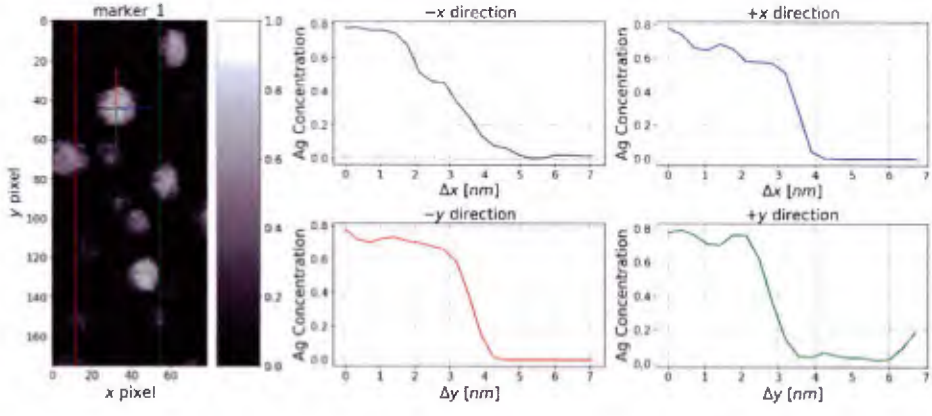

Figure 3: The approximate concentration field and the extracted profile in each direction as measured relative to the marker location for a time of 116 hours.

2.3 Task 3: Bayesian system identification

We have developed and evaluated a probabilistic formulation of system identification (ID) to sparse, noisy, and indirect data. We have then showed that the optimization functions derived from decision rules based on the full Bayesian solution to the system ID problem lead to significant system recovery benefits compared with many commonly used least-squares based learning objectives that include differentially constrained least squares and least squares reconstructions of discrete time steppers like dynamic mode decomposition (DMD) and sparse recovery of nonlinear dynamics.

Specifically, these other settings set up a least squares objective function that either assumes deterministic dynamics or learns a propagator. For example in [10], the authors learn a Hamiltonian system (denoted a Hamiltonian neural network) by minimizing the square difference between the derivative of the parameterized Hamiltonian and the numerical derivatives of the states

$$J(\theta) = \sum_{i=1}^{N} \left\| \frac{q_i - q_{i-1}}{\Delta t} - \frac{\partial \mathcal{H}_\theta}{\partial p} \right\|^2 + \left\| \frac{p_i - p_{i-1}}{\Delta t} + \frac{\partial \mathcal{H}_\theta}{\partial q} \right\|^2, \tag{8}$$

where q is the position, p the momentum, $\mathcal{H}_\theta$ the Hamiltonian dependent on parameters θ, and N the number of measurements. The authors used finite difference approximations indicated above to approximate true time derivatives — one facet of the existing approaches that we have found to significantly reduce the performance of these methods. We have found similar least squares formulations that are prevalent a number of other data-driven recovery techniques to be similarly deficient compared to the fully probabilistic formulation.

Instead we have developed methods to solve the following probabilistic formulation. Let $\mathbb{R}$ denote the set of reals and $\mathbb{Z}_+$ denote the set of positive integers. Let $d \in \mathbb{Z}_+$ denote the size of a state space, $m \in \mathbb{Z}_+$ denote the size of an observation space, and $k \in \mathbb{Z}_+$ denote a time index corresponding to a time $t_k \geq 0$. Sequential time indices will typically occur with a constant interval Δt so that $t_k = t_{k-1} + \Delta t$. A discrete-time dynamical system evolves hidden states $X_k \in \mathbb{R}^d$ at these time intervals. The hidden states are observed through a noisy and potentially indirect measurement operator providing us data $y_k \in \mathbb{R}^m$. These data are viewed as realizations of a state-dependent stochastic process Y_k.

All unknown parameters are denoted by $\theta \in \mathbb{R}^p$ for $p \in \mathbb{Z}_+$. These define a search space over models. These parameters are partitioned $\theta = (\theta_\Psi, \theta_h, \theta_\Sigma, \theta_\Gamma)$ across several uncertain sources including dynamics model parameters θ_Ψ, observation model parameters θ_h, process noise parameters θ_Σ, and observation noise parameters θ_Γ. Together these states, observations, and parameters are related through the following discrete-time stochastic process [19]

$$\begin{aligned} X_k &= \Psi(X_{k-1}, \theta_\Psi) + \xi_k; \quad \xi_k \sim \mathcal{N}(0, \Sigma(\theta_\Sigma)) \\ Y_k &= h(X_k, \theta_h) + \eta_k; \quad \eta_k \sim \mathcal{N}(0, \Gamma(\theta_\Gamma)), \end{aligned} \tag{9}$$

for $k = 1, \ldots, n$ where $\Psi : \mathbb{R}^d \times \mathbb{R}^p \to \mathbb{R}^d$ is the dynamics operator, ξ_k is the process noise with uncertain covariance $\Sigma(\theta_\Sigma)$, $h : \mathbb{R}^d \times \mathbb{R}^p \to \mathbb{R}^m$ is the observation/measurement operator, Y_k is the predictive stochastic process for the observable, and η_k is the observation noise with uncertain covariance $\Gamma(\theta_\Gamma)$.

Our system identification goal is to compute a posterior distribution over all possible systems. This posterior distribution is given by Bayes rule

$$p(\theta \mid \mathcal{Y}_n) = p(\theta) \frac{\mathcal{L}(\theta; \mathcal{Y}_n)}{p(\mathcal{Y}_n)}, \quad \text{where } \mathcal{Y}_n = (y_1, \ldots, y_n), \tag{10}$$

where the prior is denoted by $p(\theta)$ and the marginal likelihood is a function of the unknown parameter

$$\mathcal{L}(\theta; \mathcal{Y}_n) \equiv p(Y_1 = y_1, \ldots, Y_n = y_n \mid \theta). \tag{11}$$

In the context of optimization, the optimal solution would be found by optimizing the log posterior to obtain a maximum *a-posterior* estimate. However, in the Bayesian approach we sample from the posterior to account for the full uncertainty. Unlike the least squares objectives described in the introduction, this posterior accounts for parameter, model, and measurement uncertainty.

The Markovian properties of the system allow the likelihood to factorize as

$$\mathcal{L}(\theta; \mathcal{Y}_n) = \mathcal{L}_1(\theta; \mathcal{Y}_1) \prod_{k=2}^{n} \mathcal{L}_i(\theta; \mathcal{Y}_i), \tag{12}$$

where we have set $\mathcal{L}_1(\theta; \mathcal{Y}_1) \equiv p(Y_1 = y_1 \mid \theta)$ and $\mathcal{L}_k(\theta; \mathcal{Y}_k) \equiv p(Y_k = y_k \mid \theta, \mathcal{Y}_{k-1})$ for $k = 2, \ldots, n$.

Future predictions are then obtained by weighting the dynamics with respect to the dynamics model parameters

$$p(X_k \mid \mathcal{Y}_n) = \int p(X_k \mid \theta) p(\theta \mid \mathcal{Y}_n) d\theta, \tag{13}$$

for $k > 0$. As shown in [8], we use a stochastic process to enhance and robustify learning, even if our intent is really to learn a deterministic model $X_k = \Psi^k(X_0, \theta_\Psi)$, where Ψ^k refers to k applications of the dynamics. To this end, future predictions will therefore only incorporate the updated posterior of the dynamics model

$$p(X_k \mid \theta_\Psi) = \delta_{\Psi^k(x_0, \theta_\Psi)}(X_k), \tag{14}$$

so that the marginal will integrate over all possible parameters $p(X_k \mid \mathcal{Y}_n) = \int p(X_k \mid \theta_\Psi) p(\theta_\Psi \mid \mathcal{Y}_n) d\theta$.

The Bayesian inference problem (10) is solved using an approximate marginal delayed-rejection Markov Chain Monte Carlo method called UKF-MCMC [13]. This method involves using an unscented Kalman filter (UKF) to compute an approximation to the marginal likelihood. This marginal likelihood is represented recursively and given by the Theorem 2.3.

Theorem (Marginal likelihood (Th. 12.1 [19])). *Let $\mathcal{Y}_k \equiv \{y_i; i \leq k\}$ denote the set of all observations up to time k. Let the initial condition be uncertain with distribution $p(X_0 \mid \theta)$. Then the marginal likelihood (12) is defined recursively in three stages: prediction*

$$p(X_{k+1} \mid \theta, \mathcal{Y}_k) = \sqrt{2\pi}^{-d} |\Sigma(\theta_\Sigma)|^{-\frac{1}{2}} \times \int \exp\left(-\frac{1}{2} \|X_{k+1} - \Psi(X_k, \theta_\Psi)\|^2_{\Sigma(\theta_\Sigma)}\right) p(X_k \mid \theta, \mathcal{Y}_k) dX_k \tag{15}$$

update,

$$p(X_{k+1} \mid \theta, \mathcal{Y}_{k+1}) = p(X_{k+1} \mid \theta, \mathcal{Y}_k) \times \frac{\exp\left(-\frac{1}{2} \|Y_{k+1} - h(X_{k+1}, \theta_h)\|^2_{\Gamma(\theta_\Gamma)}\right)}{\sqrt{2\pi}^{m} |\Gamma(\theta_\Gamma)|^{\frac{1}{2}} p(Y_{k+1} \mid \theta, \mathcal{Y}_k)} \tag{16}$$

and marginalization,

$$\mathcal{L}_{k+1}(\theta \mid \mathcal{Y}_{k+1}) = \sqrt{2\pi}^{-m} |\Gamma(\theta_\Gamma)|^{-\frac{1}{2}} \times \int p(X_{k+1} \mid \theta, \mathcal{Y}_k) \exp\left(-\frac{1}{2} \|Y_{k+1} - h(X_{k+1}, \theta_h)\|^2_{\Gamma(\theta_\Gamma)}\right) dX_{k+1} \tag{17}$$

for $k = 1, 2, \ldots$.

The UKF-MCMC algorithm evaluates the above specified prediction, update, and marginalization steps by running a UKF every time the marginal likelihood needs to be evaluated within the Metropolis-Hastings step of the the MCMC procedure. In [8] we showed that the asymptotic cost of each log posterior evaluation scales cubically in the number of states and observations, and linearly in the number of data points.

2.3.1 Inverse design: the outerloop of physics-based learning

We now consider using some of the ML machinery that we have developed for the outer-loop application of design of materials. The chosen application is to design a self oscillating gel to obtain a target motion [3]. This problem is closely related to the other inversion procedures that we have considered as part of this project, however we are now inverting for fixed parameters to fit a time series of field data. The architecture of the inverse design structure can be seen in Fig. 4. Lattice coordinates and force components of simulated gels are used to train the neural networks. Specifically, data is obtained at 10 successive time steps, each $0.1s$ apart from each other are used to train the classifier networks.

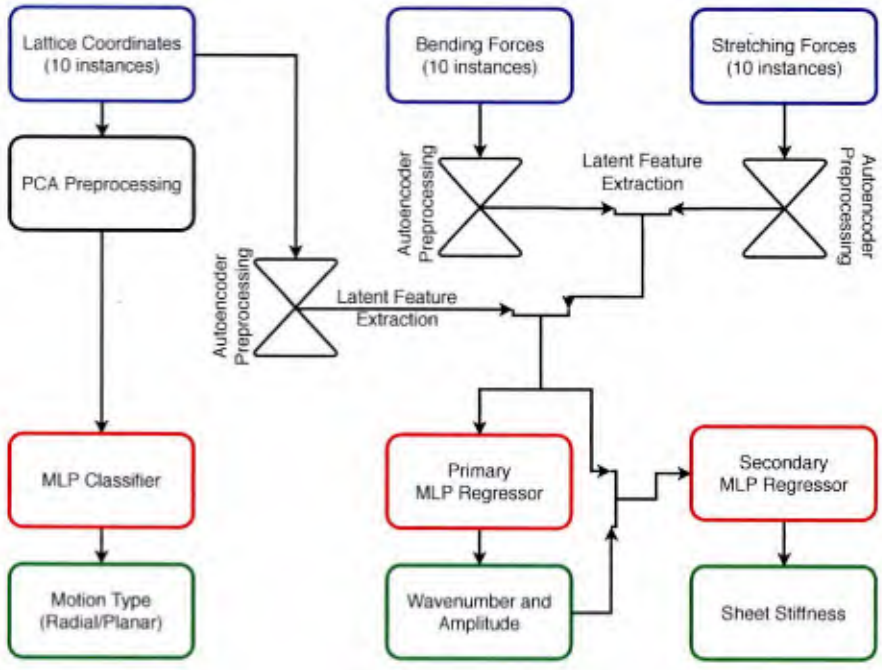

Figure 4: Our current best architecture for inverse design of materials application. The sheet motion is parametrized by the reaction type (motion type), wavenumber, amplitude, and sheet stiffness.

PCA Enhanced MLP Neural Network for Classification As it can be seen from the diagram in Fig. 4, the classification algorithm takes only the x coordinates of the lattice to classify the type of the motion on the gel sheet. PCA transformation is used to decrease the training time of the neural networks as well as increase the classification accuracy as shown in Fig. 5. With the introduction of PCA preprocessing to the architecture, training time of the best performing neural network combination (1 hidden layer, 30 hidden units) decreases from $13s$ to $1.3s$. However, the improvement in the classification accuracy is not as dramatic as in training time. The neural networks are tested over sets from different time instances and classification accuracy rises from from 95% to 99%.

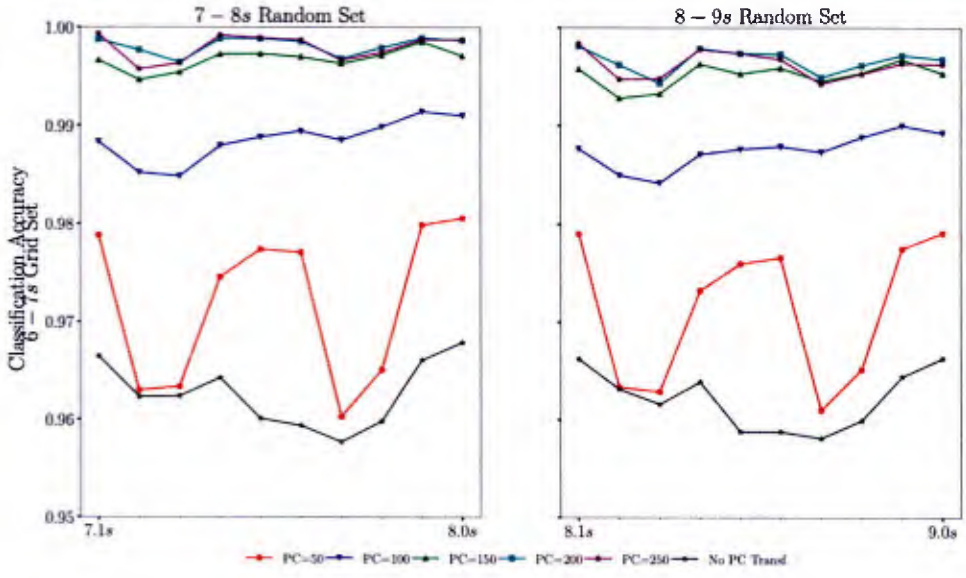

Figure 5: Results of the classification architecture

AE Enhanced 2-Layered MLP Architecture for Regression After classifying the motion on the sheet gel, the simulation outputs (displacement and force terms) are directed to the second stage of the inverse design architecture. In this stage there are two individual neural networks responsible for finding the parameters of the gel sheet, specialized in each one of the motion types.

To recover the sheet parameters, y and z coordinates from lattice positions are also used to leverage the full three dimensional motion. Similar to the classification case, 10 instances of gel motion are used to construct the training database. Instead of using PCA transformation for preprocessing the lattice coordinates, an auxiliary autoencoder network is trained to extract latent features from the coordinate dataset. In addition to the processed coordinate information, force components such as bending and stretching forces are also added to the training algorithm by transforming through separate autoencoder networks.

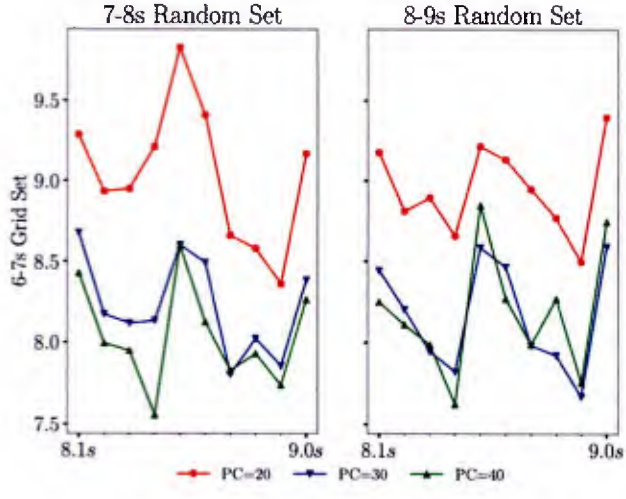

Figure 6: Mean Percentage Error for Wavenumber

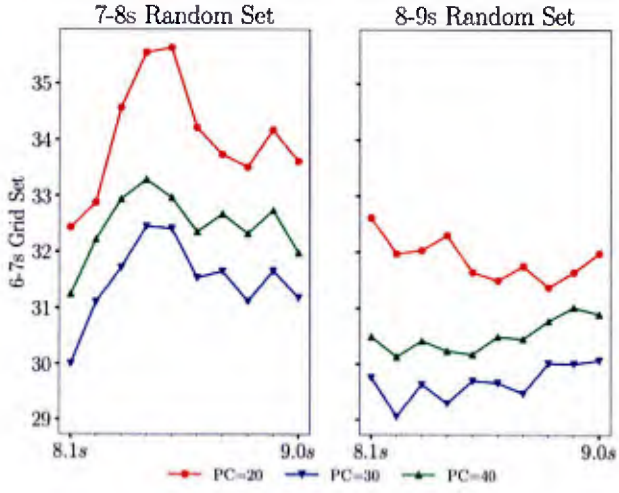

Figure 7: Mean Percentage Error for Sheet Stiffness

One of the main differences of the regression part is the introduction of the second layer MLP network. After processing the data through the first layer MLP, the trained first layer network recovers the number of the waves on a sheet and the amplitude of the waves traveling on the sheet with a mean estimation error $< 10\%$ (see Fig. 6). However, the sheet stiffness parameter can not be recovered with similar precision. Therefore, the recovered two parameters from the first layer of the neural network is added to the training dataset to train a secondary network that specializes in recovering the sheet stiffness. Currently, the sheet stiffness can be recovered with a mean estimation error $\approx 30\%$.

Other architectures are also currently investigated to further enhance the estimation performance of the continuous sheet parameters. The aim for the regression part is estimating the sheet parameters with a mean error 10% for all parameters and comparing the motion of the initial gel sheet with the motion reproduced using the recovered parameters by the inverse learning architecture.

2.4 Task 4: Optimal experimental design

In this section we present the formulation of the OED problems in our project. The corresponding results can be found in Section 4.2

2.4.1 Bayesian OED

We first apply Bayesian OED to system identification problems, where the forward model $G(\theta, d)$ outputs QoIs given parameters θ and design variables d. It involves solving the following reaction-diffusion equations for two species C_1 and C_2:

$$\begin{aligned} \frac{\partial C_1}{\partial t} &= e^1 \frac{\partial^2 C_1}{\partial x^2} + 0.1 - C_1 + 1C_1^2 C_2 \\ \frac{\partial C_2}{\partial t} &= e^1 \frac{\partial^2 C_2}{\partial x^2} C_2 + 0.9 - C_1^2 C_2 \end{aligned} \tag{18}$$

where the unknown parameters are highlighted in blue, and collectively denoted as θ. After obtaining the solution fields of C_1 and C_2, we still need to post-process them to calculate the QoIs,

which are the final outputs of the forward model. We derive four scalar QoIs (for each species):

- matrix composition: average of all local minima of the solution field;
- average of all local maxima of the solution field;
- the mean of size distributions of the peaks at the 25% threshold;
- the variance of size distributions of the peaks at the 25% threshold.

Therefore, we totally have 8 QoIs. Furthermore, we select design variables that we can control when conducting experiments to be

- average concentration in the initial condition of species 2

$$C_2(x, t = 0) = C_{2m} + \tilde{C}_2(x), \tag{19}$$

  where we design for $d_1 = \log_{10} C_{2m}$, and $\tilde{C}_2$ at each location x is a sample from uniform distribution $\mathcal{U}[-0.1, 0.1]$, which assumes that we know the noise on initial condition exactly. In principle, $C_1(t = 0)$ should also be a design variable, but we fix it to $C_1(x, t = 0) = 0.5 + \tilde{C}_1(x)$ here because we want to provide demonstrations in two-dimensional design space for visualization purposes.

- the time of concentration snapshot $d_2 = \tau$ at which QoIs are computed.

Collectively, we denote the design variables $d = [d_1, d_2]$.

The objective function of OED is selected as the information gain on the unknown parameters θ. The information gain is mathematically expressed as the Kullback-Leibler (KL) divergence from the prior(before-seeing-data) distribution to the posterior (updated-from-data) distribution, and a larger divergence indicates a greater information gain (larger uncertainty reduction).

$$\begin{aligned} U(d) &= \mathbb{E}_{y|d}[D_{\mathrm{KL}}\left(p(\theta|y,d)||p(\theta)\right)] \\ &= \int_{\mathcal{Y}} \int_{\Theta} \ln\left[\frac{p(\theta|y,d)}{p(\theta)}\right] p(\theta|y,d)\, d\theta\, p(y|d)\, dy. \end{aligned} \tag{20}$$

OED then seeks the design variable that maximizes this quantity:

$$d^* = \underset{d}{\operatorname{argmax}}\, U(d) \tag{21}$$

Equation 20 generally has no closed-form solution, thus we adopt a double-nested Monte Carlo (MC) estimator to approximate it [11]:

$$U(d) \approx \frac{1}{N} \sum_{i=1}^{N} \left\{ \ln[p(y^{(i)}|\theta^{(i)}, d)] - \ln[p(y^{(i)}|d)] \right\}, \tag{22}$$

where $\theta^{(i)} \sim p(\theta)$ are samples from the prior distribution, and $y^{(i)} \sim p(y|\theta^{(i)}, d)$ are samples from the likelihood model. Furthermore, the log-evidence term, $\ln\left[p(y^{(i)}|d\right]$, needs to be approximated separately still, using yet another MC estimator:

$$p(y^{(i)}|d) = \frac{1}{M} \sum_{j=1}^{M} p(y^{(i)}|\theta^{(i,j)}, d), \tag{23}$$

where now $\theta^{(i,j)} \sim p(\theta)$ are additional new samples from the prior. Equations (22) and (23) together complete the overall double-nested MC estimator.

2.4.2 Non-Bayesian OED

We implemented non-Bayesian OED on the variational system identification problem using two different criteria. The first criterion is used to find the optimal time t^* that maximizes the expected determinant of $\Xi^T\Xi$, hence minimize the uncertainty on unknown pre-factors w:

$$t^* = \underset{t}{\operatorname{argmax}}\, U(t) \equiv \mathbb{E}_\omega[\det(\Xi^T\Xi)] \tag{24}$$

where ω represents unknown true pre-factors, and Ξ is the operator matrix formed from the finite-element (FEM) results corresponding to ω. The second criterion is to minimize the expected mean squared error (MSE) between true pre-factors ω and estimation $\hat{\omega}$.

$$t^* = \underset{t}{\operatorname{argmin}}\, \mathbb{E}_\omega \mathbb{E}_\epsilon[\mathrm{MSE}(\hat{\omega}, \omega)] \tag{25}$$

where $\hat{\omega}$ is the stepwise regression results corresponding to the noisy coefficient matrix $\widetilde{\Xi}$ and RHS $\widetilde{y}$, which are obtained by adding additional noise $\epsilon \sim \mathcal{N}(0, 10^{-12})$ to the original operator matrix Ξ and observed time derivatives y:

$$\widetilde{\Xi}_{ij} = \Xi_{ij}(1+\epsilon)$$
$$\widetilde{y}_i = \widetilde{y}_i(1+\epsilon).$$

3 DATA SETS

3.1 Data driven exchange-correlation functional in DFT

3.1.1 Acquisition of accurate ground-state densities

The inverse DFT and the subsequent modeling of E_{xc} by system identification (or other means) requires high quality correlated *ab initio* ground-state densities. To this end, we obtain these densities using configuration interaction (CI) calculations. In particular, the densities are generated using the incremental full-CI approach presented in [26] and discretized using a large Gaussian basis-set (e.g., universal Gaussian basis set (UGBS) [6] or cc-PVTZ Gaussian basis set [7] based on the system). The incremental full-CI solves the electronic Schrdinger equation via a many-body expansion and asymptotically produces the exact electronic energy and density as the number of bodies in the expansion approaches the all-electron limit.

3.1.2 Removal of Gaussian basis-set artifacts

To begin with, the CI density obtained from a Gaussian basis has wrong decay characteristics away from the nuclei (i.e., Gaussian decay instead of exponential decay). This, in turn, results in incorrect asymptotics in the v_{xc} obtained from an inverse DFT calculation. Thus, to ensure the correct asymptotics in v_{xc}, we enforce approximate boundary conditions on the v_{xc} in the low density region ($\rho < 10^{-7}$), so as to ensure the correct $-1/r$ in decay in v_{xc}. In particular, we use the Fermi-Amaldi potential (v_{FA}) as the approximate boundary condition for v_{xc}. We remark that this approach is also crucial to obtaining an agreement between the highest occupied Kohn-Sham eigenvalue (ϵ_{H}) and the negative of the ionization potential (I_{p}), as mandated by the Koopmans' theorem [17, 5].

Additionally, the Gaussian basis-set based densities lack the cusp at the nuclei, which, in turn, leads to undesirable oscillations in the v_{xc} near the nuclei in any inverse DFT calculation (illustrated for the case of equilibrium H_2 in Fig.8(a)). We remedy these oscillations by adding a small

correction, $\Delta\rho(\boldsymbol{r})$ to $\rho_{\text{data}}(\boldsymbol{r})$, so as to correct for the missing cusp at the nuclei. The $\Delta\rho(\boldsymbol{r})$ is given by

$$\Delta\rho(\boldsymbol{r}) = \rho_{\text{FE}}^{\text{DFT}}(\boldsymbol{r}) - \rho_{\text{G}}^{\text{DFT}}(\boldsymbol{r})\,, \tag{26}$$

where $\rho_{\text{FE}}^{\text{DFT}}(\boldsymbol{r})$ is the ground-state density obtained from a forward DFT calculation using a known exchange-correlation functional (e.g., LDA, GGA) and discretized using the FE basis; and $\rho_{\text{G}}^{\text{DFT}}(\boldsymbol{r})$ denotes the same, albeit obtained using the Gaussian basis employed in the CI calculation. The key idea here is that $\rho_{\text{FE}}^{\text{DFT}}(\boldsymbol{r})$, obtained from the FE basis, contains the cusp. Thus, one can expect $\Delta\rho$ to reasonably capture the Gaussian basis set error near the nuclei. The efficacy of the $\Delta\rho$ correction is illustrate in Fig. 8(b) wherein the v_{xc}'s are obtained for the cusp-corrected density (i.e., $\rho_{\text{data}} + \Delta\rho$) for the equilibrium H_2, with two different $\Delta\rho$: $\Delta\rho_{\text{LDA}}$ evaluated using an LDA functional and $\Delta\rho_{\text{GGA}}$ evaluated using a GGA functional. As evident, both $\Delta\rho_{\text{LDA}}$ and $\Delta\rho_{GGA}$ based cusp-correction generate smooth v_{xc} profiles. More importantly, both the profiles are nearly identical, suggesting only weak sensitivity of the v_{xc} to the choice of the exchange-correlation functional used in $\Delta\rho$.

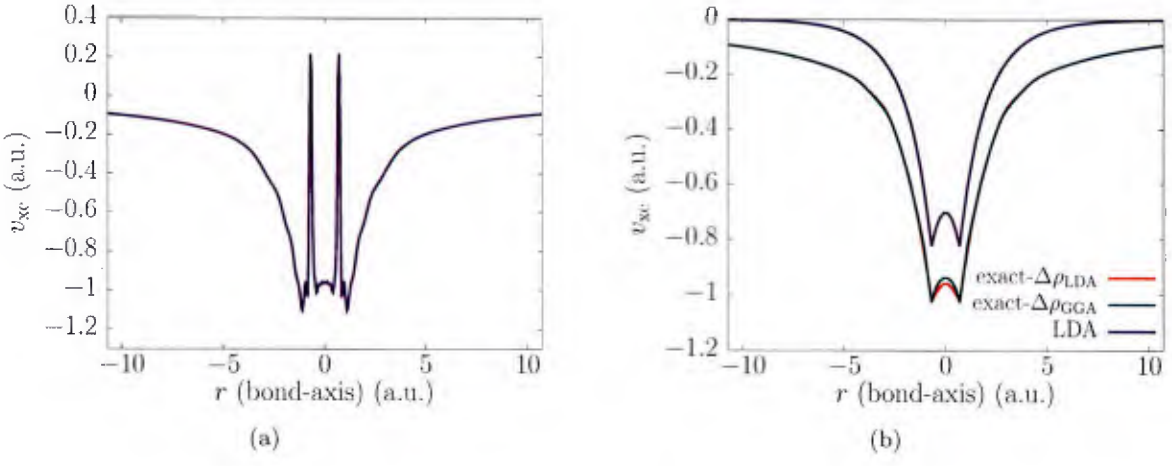

Figure 8: Comparison of the effect of $\Delta\rho$ correction on v_{xc}. In (a) the lack of cusp in ρ_{data} induces wild oscillations in the v_{xc} obtained through inversion. In (b) the exact exchange-correlation potential is evaluated using the cusp-corrected configuration interaction (CI) density. The effect of the choice of the functional used in evaluating the cusp correction is demonstrated using two different functionals—LDA (exact-$\Delta\rho_{LDA}$) and GGA (exact-$\Delta\rho_{GGA}$). Additionally, the profile of the widely used LDA-based v_{xc} is provided for a direct comparison against the exact v_{xc}.

3.1.3 Acquisition of exact v_{xc}

We combine the aforementioned strategy to remove the Gaussian basis-set artifacts, along with the finite element discretization, to evaluate the exact v_{xc} for several other systems, including H_2 at different bond-lengths, LiH, H_2O, benzyne (C_6H_4). For all the systems, the inverse problem is deemed to have converged when $||\rho_{\text{data}} - \rho||_{L^2} < 10^{-4}$, a stringent accuracy for inverse DFT calculations. Additionally, for all the systems we obtain remarkable agreement between ϵ_H and $-I_{\text{p}}$ (within 4 mHa), further validating the accuracy of our approach. Below we showcase the profile of the exact v_{xc} for few of the systems. Fig. 9 presents the exact v_{xc} for two H_2 molecule: one in the compressive regime (H_2^a, at 0.8× the equilibrium bond-length) and the other at the extension regime (H_2^b, at 1.35× the equilibrium bond-length). We emphasize that H_2^b is a prototypical system where all existing exchange-correlation approximations perform poorly, owing to their failure in

handling strong correlations. We could successfully solve the inverse DFT problem for it, thereby, suggesting that our approach works equally well for strongly correlated systems. Fig. 10 compares the exact and the LDA-based $v_{\rm xc}$ for LiH at equilibrium bond-length. In this case, for the exact $v_{\rm xc}$, we observe an atomic inter-shell structure-marked by the local maxima and minima at around $\pm = 0.4$ a.u.. This atomic inter-shell structure is a distinctive feature of the exact $v_{\rm xc}$ and is absent in the standard xc approximations, as evident from the LDA-based $v_{\rm xc}$. Moreover, we observe a distinct local maximum at the H atom, as opposed to a local minimum in the LDA-based $v_{\rm xc}$. Fig 11 compares the exact $v_{\rm xc}$ against $v_{\rm xc}^{\rm LDA}$ for the benzyne molecule (C_6H_4)—a strongly correlated species that has previously served as a test for accurate wavefunction methods. This example underscores the efficacy of our approach in handling both strongly correlated as well as large (by inverse DFT standards) systems. As expected for the exact $v_{\rm xc}$, we observe an atomic inter-shell structure—marked by a yellow ring around the C atoms, which, on the other hand, are absent in the case of $v_{\rm xc}^{\rm LDA}$. We also observe a deeper potential around the C atom, as compared to $v_{\rm xc}^{\rm LDA}$, suggesting a higher electronegativity on the C atom than that predicted by LDA.

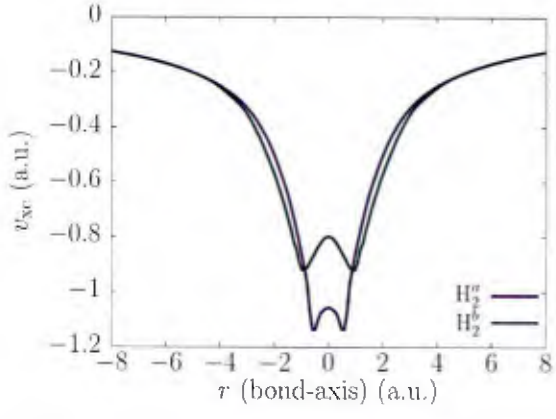

Figure 9: Exact $v_{\rm xc}$ for H_2 molecule at 0.8x and 1.35x the equilibrium bond-length, denoted as H_2^a and H_2^b, respectively.

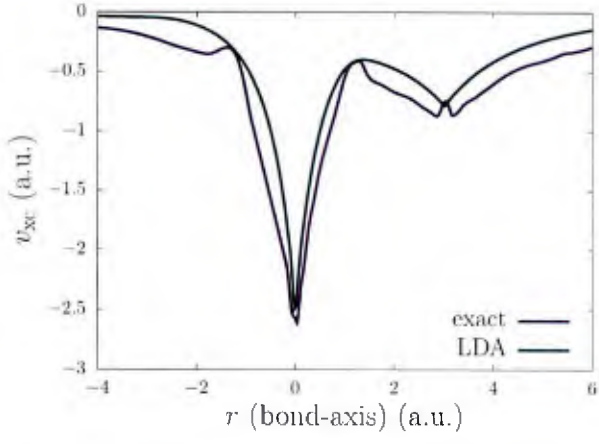

Figure 10: Comparison of the exact and LDA-based $v_{\rm xc}$ for LiH.

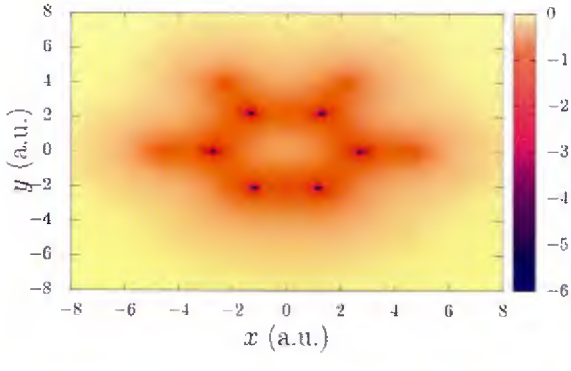

(a)

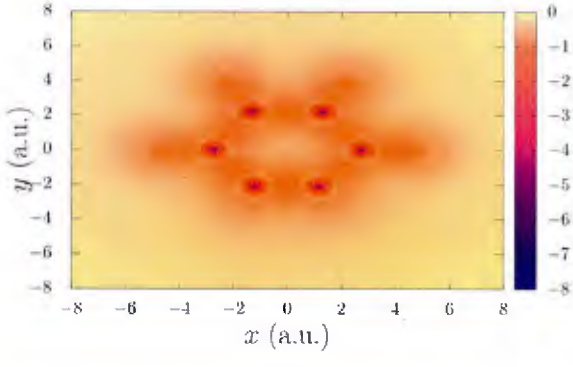

(b)

Figure 11: Comparison of LDA-based (a) and the exact (b) exchange-correlation potential for benzyne (C_6H_4) (in a.u.), presented on the plane of the molecule.

3.2 Bayesian system ID and inverse design

In the course of developing and testing algorithms for both learning dynamical systems as well as inverse design we have developed simulated datasets. For the inverse design set in particular, we have developed an extensive database of forward simulations which we will make available upon publication of the journal article. This dataset consists of the evolution of a self-oscillating gel over a fixed time period for various different parameter conditions. It uses the simulation developed in [3]. This simulation data set will be made available on `https://deepblue.lib.umich.edu/data`

4 RESULTS

4.1 Scale bridging and physics discovery

4.1.1 Scale bridging from atomistic to continuum modeling of Ni-Al through active learning of free energy density

Density functional theory calculations for the Ni-Al system and a subsequent cluster expansion Hamiltonian parameterized by Goiri and Van der Ven [9] were applied to the algorithm described in section 2.2.1. The evolution of the IDNN is presented in Figure 12 by plotting a slice of the predicted chemical potential μ_0 as a function of c and η_1, with $\eta_1 = \eta_2 = \eta_3$. Significant changes are seen in the first few iterations of the active learning workflow, with evident convergence in the later iterations. A slice of the final, analytically integrated free energy DNN, referenced to pure Ni and the perfectly ordered $L1_2$, is shown in Figure 13, again with $\eta_1 = \eta_2 = \eta_3$. An energy well is seen at about $c = 0.23$, corresponding to the γ' Ni-Al precipitates for the $L1_2$ variant with all positive valued order parameters. A well near $c = 0.045$ represents the γ solid solution phase. A few spurious regions of slight convexity exist in the DNN surface, but they do not seem to negatively affect the resulting precipitate formation in the phase field results (see Figure 14).

The following phase field equations, coupled with elasticity, were solved numerically using the finite element method and backward Euler time integration:

Cahn-Hilliard:

$$\frac{\partial c}{\partial t} = \nabla \cdot (M \nabla \mu_0) \tag{27}$$

$$\mu_0 = \frac{\partial f}{\partial c} + \frac{\partial \psi}{\partial c} + \boldsymbol{P} : \frac{\partial \boldsymbol{F}^e}{\partial c} - \chi_0 \nabla^2 c \tag{28}$$

Allen-Cahn:

$$\frac{\partial \eta_i}{\partial t} = -L \left(\frac{\partial f}{\partial \eta_i} - \chi_i \nabla^2 \eta_i \right), \qquad i = 1, 2, 3 \tag{29}$$

Elasticity:

$$0 = \nabla \cdot \left(\boldsymbol{P} \boldsymbol{F}^{\lambda^{-T}} \right) \tag{30}$$

The simulation results are plotted in Figure 14. The expected development of precipitates of multiple $L1_2$ variants are shown, demonstrating that the free energy DNN has captured the appropriate physics. The blue, orange, brown, and purple regions represent the four $L1_2$ variants that develop in the Ni-Al precipitates at about $c = 0.23$. The tan background shows the γ phase solid solution. As expected, each precipitate consists of a single $L1_2$ variant, without any antiphase boundaries forming within a single precipitate. This is due to the antiphase boundary energy being greater than twice the interfacial energy, as described by Wang, et al [22]. The precipitates are seen to grow over time, forming roughly rectangular shapes. This faceting is due to the interplay between the misfit strain, interfacial energy, and cubic elasticity. Note that the smaller precipitates remain circular, since the interfacial energy is dominating the anisotropic elastic response [21].

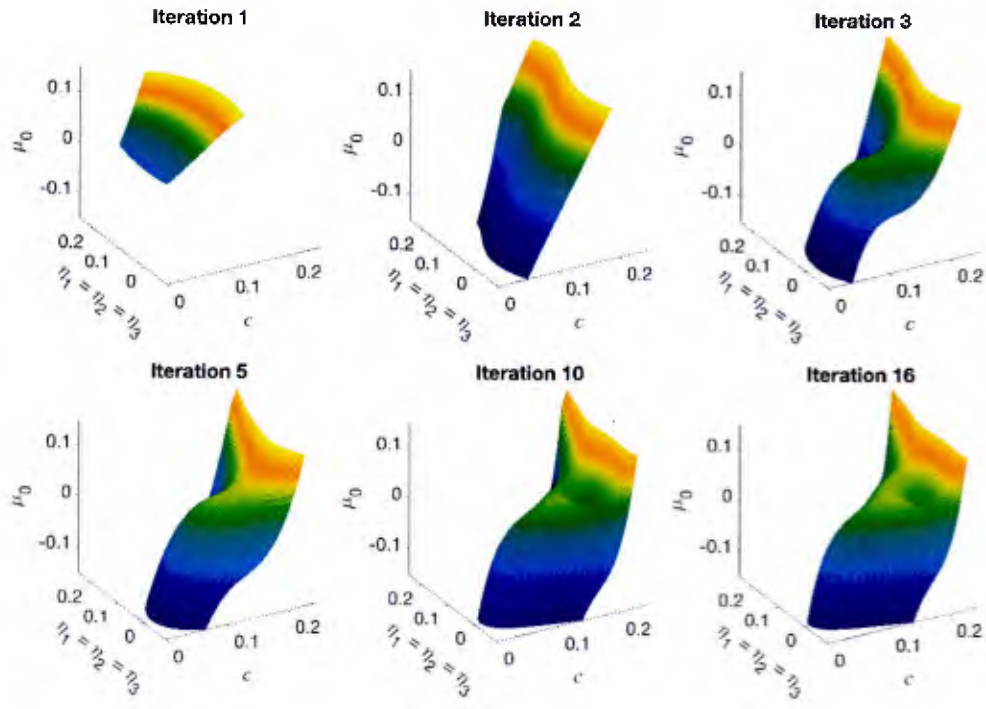

Figure 12: Evolution of the IDNN representing the chemical potential for a two-dimensional subspace, over iterations of the active learning workflow.

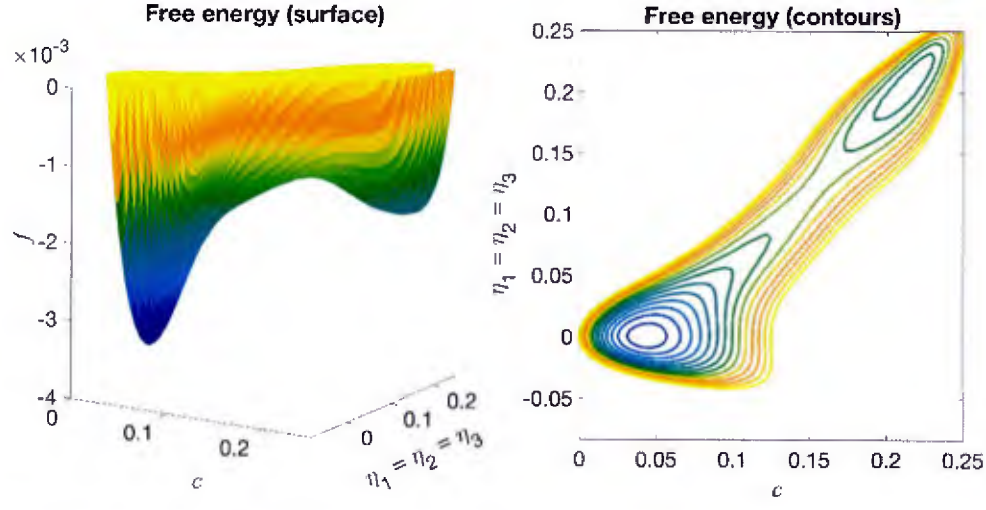

Figure 13: Surface and contour plots for a 2D subspace ($\eta_1 = \eta_2 = \eta_3$) of the converged DNN representation of the homogeneous free energy density.

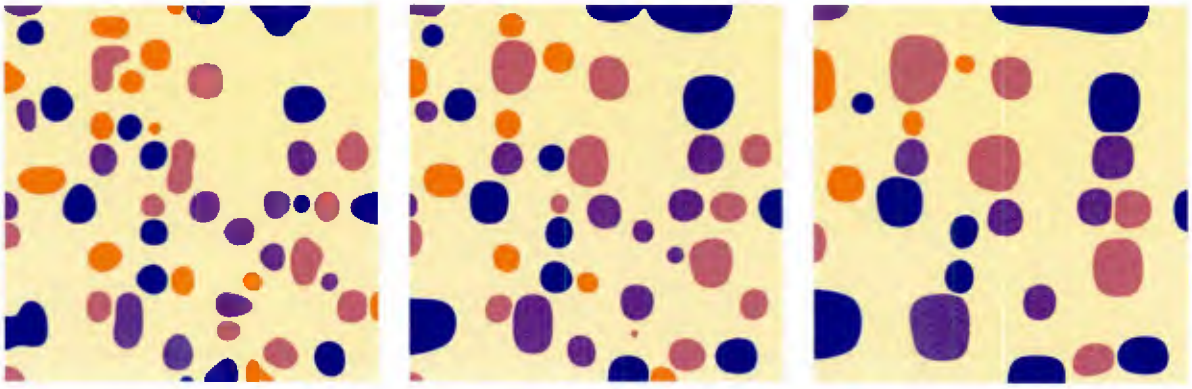

Time step 50 Time step 300 Time step 650

Figure 14: Evolution of the Ni_3Al precipitates in the phase field simulations, with four $L1_2$ variants shown as blue, orange, brown, and purple.

4.1.2 Identification of the physics of nucleation and growth in Ag-Al alloys

We obtained the experimental images of an Ag-Al system from early work by Co-PI Marquis [14]. Each image is a thin slice of the material, on the order of 10 nm thick, and the atoms visible are integrated through this depth. The resolution is on the order of a unit cell (~0.35 nm/pixel) so the intensity of the coloration may be interpreted as a concentration of Ag against the background Al matrix. The raw data are shown in Figure 15. At early times, the Ag atoms are distributed more uniformly when compared to later times. As time progresses, the Ag atoms move together and form precipitates which grow in size and color intensity.

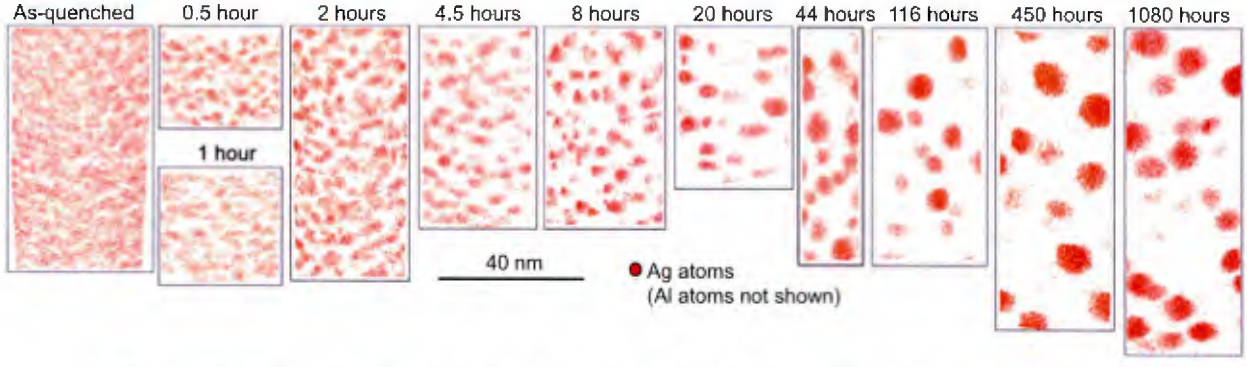

Figure 15: Raw APT data showing the evolution of GP zones in Ag-Al (2.7 atomic %).

Starting from the concentration field generated from the raw data, we can define the order parameter field η which takes on the value of 1 for the precipitate phase (GP zone) and 0 for the matrix phase with a smooth interface in between. We fitted a hyperbolic tangent function,

$$\widetilde{\eta} = \tanh(5(C - 0.42)), \tag{31}$$

to the concentration field, C, and then normalized it to the desired range:

$$\eta = \frac{\widetilde{\eta} - \min(\widetilde{\eta})}{\max(\widetilde{\eta}) - \min(\widetilde{\eta})} \tag{32}$$

get the η field as shown in Figure 16.

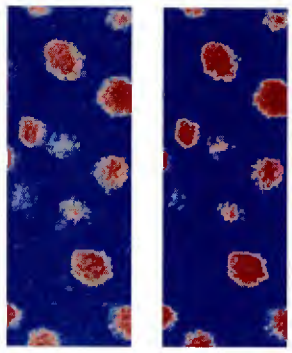

Figure 16: The concentration, C and order parameter, η fields at 450 hours. These fields are used to define nodal values in a grid-based variational approach, such as finite elements or isogeometric analysis.

We postulate that the underlying process is one of nucleation and growth, for which the Allen-Cahn equation coupled with conservative transport is the governing equation:

$$\begin{aligned} \frac{\partial C}{\partial t} &= -\nabla \cdot \boldsymbol{J} \\ \frac{\partial \eta}{\partial t} &= -L\mu_\eta \end{aligned} \tag{33}$$

where the flux is defined by $\boldsymbol{J} := -M\nabla\mu_C$, M is the mobility, and L is the kinetic coefficient. The chemical potentials $\mu_C = \delta g/\delta C$ and $\mu_\eta = \delta g/\delta\eta$ are found using standard variational methods. We aim to identify the free energy function that defines chemical potentials and the flux. We first proposed a general form of the free energy function:

$$\begin{aligned} g(C,\eta) =& \theta_1 C^4 + \theta_2 C^3\eta + \theta_3 C^2\eta^2 + \theta_4 C\eta^3 + \theta_5\eta^4 \\ &+ \theta_6 C^3 + \theta_7 C^2\eta + \theta_8 C\eta^2 + \theta_9\eta^3 + \theta_{10} C^2 + \theta_{11} C\eta + \theta_{12}\eta^2 + \theta_{13} C + \theta_{14}\eta \end{aligned} \tag{34}$$

The prefactors, $\theta_1 ... \theta_{14}$, in the free energy function are also shown in the PDEs associated with different operators which are the targets for system identification.

Leveraging the near-steady state data from growth of Ag-rich precipitates at 450 hours or later in our data set [14], we are able to identify the free energy function:

$$\begin{aligned} g(C,\eta) =& 8.9046C^4 - 34.061C^3\eta + 42.6123C^2\eta^2 - 21.7848C\eta^3 + 4.7169\eta^4 \\ &+ 1.5053C^3 + 1.4634C^2\eta - 3.9779C\eta^2 - 0.2706C\eta + 0.9861\eta^2 \end{aligned} \tag{35}$$

Figure 17 shows the double well structure of the identified free energy function. The transportation parameters can then be inferred using dynamic data. The identified Allen-Cahn equation allow us to simulate the growth of alloy precipitates from nuclei as shown in Figure 18.

4.1.3 Identification of the physics of opaque materials from 3D MRI

Working with collaborators J. Estrada and E.M. Arruda, we have access to three-dimensional, full-field displacement data measured by 3D MRI for opaque materials under different loading conditions as shown in Figure 19. Given these full-field data we have applied the Variational System Identification technique to identify the mechanisms of deformation by pinpointing the coefficients of terms in the strain energy function. We consider volumetric, isochoric and orthotropic mechanisms:

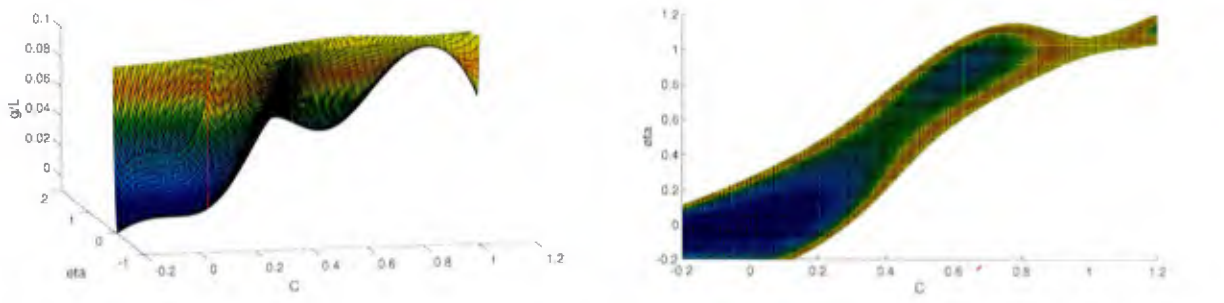

Figure 17: Identified free energy function. Note there are two wells in the free energy functions at $C = 0$ and $C = 60\%$.

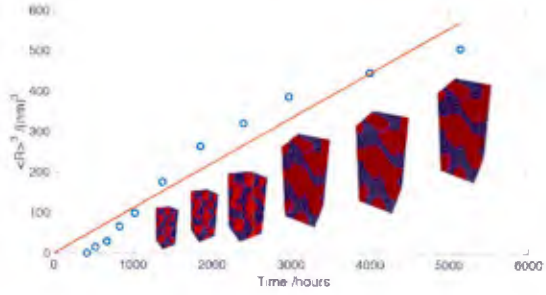

Figure 18: 3D simulated results generated by the identified Allen-Cahn equation. The cubed average radius of particles increases linearly during growth.

$$W(\boldsymbol{C}) = \underbrace{\kappa(J-1)^2}_{\text{volumetric}} + \underbrace{\mu(\overline{I}_1 - 3) + \theta_1(\overline{I}_1 - 3)^2 + \theta_2(\overline{I}_2 - 3) + \theta_3(\overline{I}_2 - 3)^2}_{\text{isochoric}} + \underbrace{\theta_4(\overline{I}_a - 1) + \theta_5(\overline{I}_a - 1)^2}_{\text{orthotropic}} \tag{36}$$

The first Piola-Kirchhoff stress tensor is $\boldsymbol{P} = \partial W/\partial \boldsymbol{F}$ and leads to the form

$$\boldsymbol{P} = \kappa \boldsymbol{P}_0(J, \boldsymbol{F}) + \mu \boldsymbol{P}_s(\overline{I}_1, \boldsymbol{F}) + \theta_1 \boldsymbol{P}_1(\overline{I}_1, \boldsymbol{F}) + \theta_2 \boldsymbol{P}_2(\overline{I}_2, \boldsymbol{F}) + \theta_3 \boldsymbol{P}_3(\overline{I}_3, \boldsymbol{F}) + \theta_4 \boldsymbol{P}_4(\overline{I}_a, \boldsymbol{F}) + \theta_5 \boldsymbol{P}_5(\overline{I}_a, \boldsymbol{F}) \tag{37}$$

where $\boldsymbol{P}_i$ is the ith basis of $\boldsymbol{P}$. The weak form of the stress equilibrium equation is:

$$\int_\Omega \nabla \boldsymbol{w} \colon \boldsymbol{P} \mathrm{d}V - \int_{\partial\Omega} \boldsymbol{w} \cdot \boldsymbol{T} = 0, \tag{38}$$

where $\boldsymbol{T}$ is the applied load, which is known. Using Variational System Identification with stepwise regression, we are able to find the best model that describes the data, thereby identifying the mechanisms of volumetric and deviatoric deformation that determine the material response. After the deformation mechanisms are identified out of Equation (36), the values of coefficients correspond identified bases could be further improved by minimizing the error on $\boldsymbol{u}$ subject to the identified physics:

$$\kappa, \mu, \boldsymbol{\theta} = \arg \min_{\widetilde{\kappa}, \widetilde{\mu}, \widetilde{\boldsymbol{\theta}}} \left\{ \sum_i^N \int_\Omega (\boldsymbol{u}^i - \boldsymbol{u}^i_{\text{ref}})^2 \mathrm{d}v \right\} \tag{39}$$

$$\text{subject to: } \nabla \cdot \boldsymbol{P} = \boldsymbol{0} \tag{40}$$

Figure 19: Reconstructed deformation from data on a soft polymer showing the reference (left), first deformed (middle) and ninth deformed (right) configurations. VSI discovers deformation mechanisms from these data.

where N is the total loading cases, $\boldsymbol{u}^i$ is the simulated displacement field under the ith loading, and $\boldsymbol{u}^i_{\text{ref}}$ is the corresponding MR-u data. This is the so called PDE-constrained optimization approach. The identified model and the coefficients show good agreement with the model chosen by expert knowledge combined with pure parameter inference. The shear and bulk moduli respectively were found to be 46.3 kPa and 1430 kPa

$$W = 46.3(J^{-2/3}tr(C) - 3) + 1430(J-1)^2 \tag{41}$$

4.2 Optimal experimental design

4.2.1 Bayesian OED

We first specify the settings of the OED problem:

- box design space $d_1 \in [-1.3, 0.7]$ and $d_2 \in [2.5, 15]$;
- prior $\theta_1 \sim \mathcal{N}(\log(0.1), 2^2), \theta_2 \sim \mathcal{N}(\log(4), 2^2), \log(-\theta_3) \sim \mathcal{N}(0, 2^2)$;
- likelihood $y|\theta, d \sim \mathcal{N}(G(\theta, d), (0.1|G| + 0.01)^2)$ reflecting a 10% noise (i.e., measurement noise is 10% of the signal plus a base term, where the base term is used to avoid zero noise).

The expected utility in the design space and the joint posterior distribution of θ_1 and θ_2 given θ_3 is known have been shown in Figure 20 and 21. It can be found that the expected utility will be higher if the snapshot time is earlier, and performing experiments at $\log_{10} C_{2m} = 0.5$ would be slightly better. From the second row in Figure 21, we can find that the posterior is narrower at $\tau = 2.5$ (earlier snapshot) compared to that at $\tau = 15$ (later snapshot), which is consistent to the utility plot. However, the first row indicates that when θ_3 is fixed at -1, the posterior at a later snapshot is even narrower, which is contradictory to the utility plot. This is because $\theta_3 = -1$ is only one possible scenario, and if we consider all the possible scenarios suggested by the prior of θ_3, we are still expected to gain more information at $\tau = 2.5$.

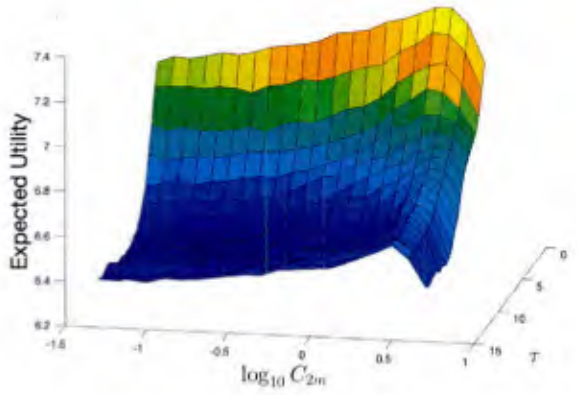

Figure 20: Expected utility surface

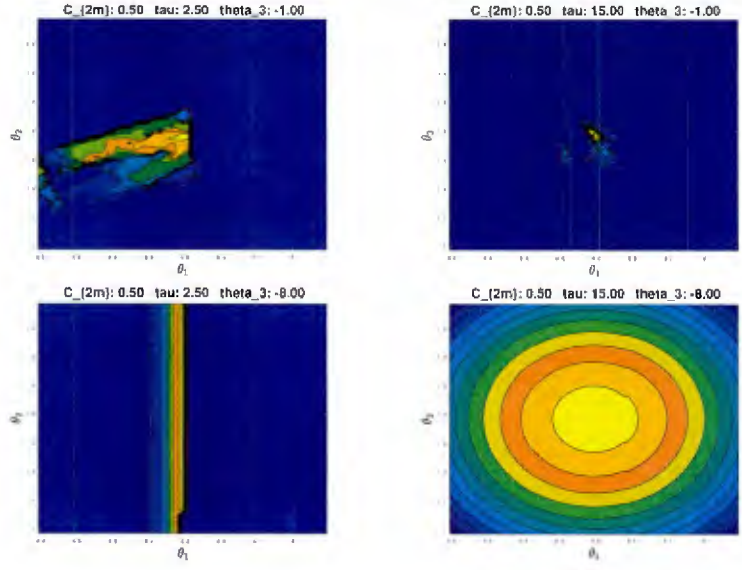

Figure 21: Joint posterior distribution of θ_1 and θ_2 given that θ_3 is known, and initial concentration of species 2 is $C_{2m} = 0.5$ (top-left: $\tau = 2.5$ and $\theta_3 = -1$; top-right: $\tau = 15$ and $\theta_3 = -1$; bottom-left: $\tau = 2.5$ and $\theta_3 = -8$; bottom-right: $\tau = 15$ and $\theta_3 = -8$).

4.2.2 Non-Bayesian OED

Figure 22 and 23 show the expected determinant and the expected MSE w.r.t. snapshot time for 6 cases, where those cases are generated by running FEM simulation. It turns out that the determinants of all cases are super small when $t < 20$, have a sudden increase between $t = 20$ and $t = 50$, and then stablize for $t > 60$. From this figure, we can conclude that we should perform the experiment after $t = 60$ for a higher determinant value and correspondingly a lower uncertainty. Moreover, the MSE curves are consistent to the determinant curves. Before $t = 20$, the MSE is high, and then sharply decreases between $t = 20$ and $t = 50$. Therefore, we can also conclude that $t > 60$ is the optimal time to perform experiments.

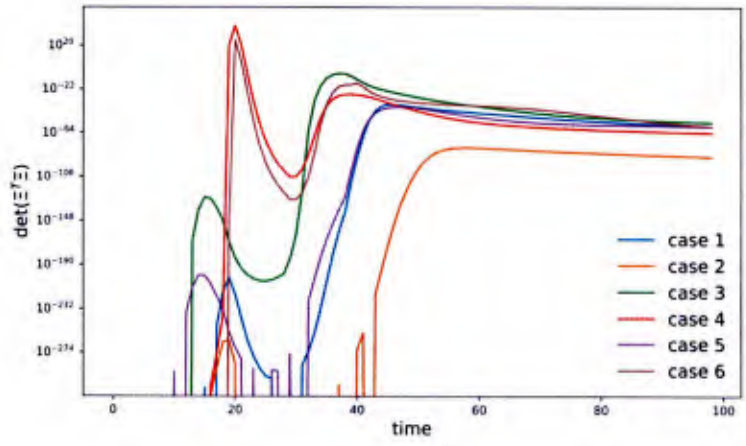

Figure 22: Determinant of $\Xi^T\Xi$ w.r.t. time for 6 cases with different true ω.

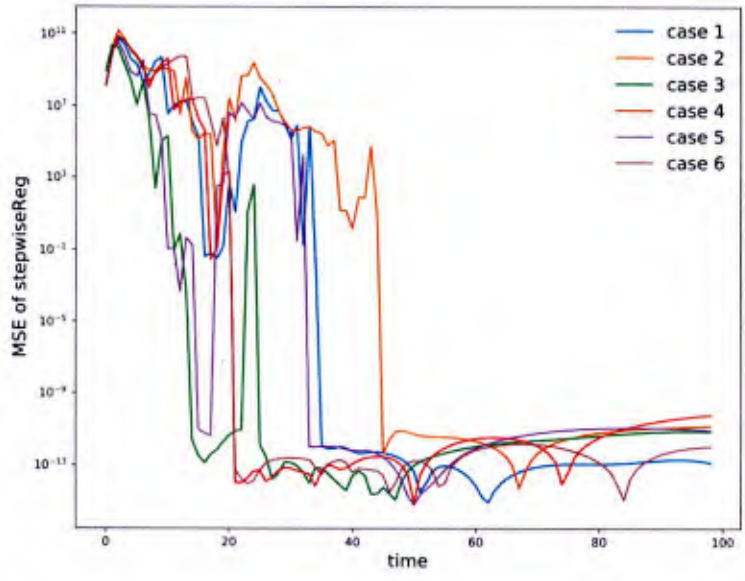

Figure 23: Expected MSE $\mathbb{E}_\epsilon[\text{MSE}(\hat{\omega}, \omega)]$ of stepwise regression w.r.t. time for 6 cases with different true ω.

5 HACKATHON – PHASE 2 SUMMARY

We focus on **Challenge Problem #2:** Ultrasound Microstructure Interaction, with the overarching **goal** to develop methods for microstructure characterization from Ultrasonic Wavefield Imaging techniques.

We approach our hackathon challenge as an inverse problem, illustrated in Figure 24. Here, the left image is one of the Euler angles obtained from electron backscattering diffraction (EBSD) measurements, and the right image is a snapshot of the ultrasound wavefield imaging measurements, of the same material specimen.

- **Forward problem/model (deduction):** given the full microstructure information (i.e., grain structures, sizes, orientations), solve the governing physics (e.g., elastodynamic wave equations) to obtain a prediction of the wave propagation.

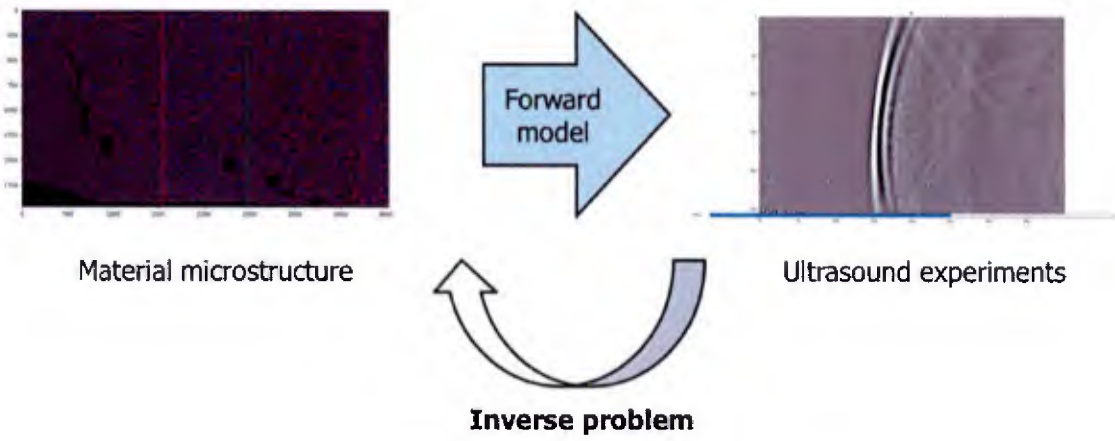

Figure 24: We approach our hackathon challenge as an inverse problem.

- **Inverse problem (inference):** given observations of the wave propagation, can we say something about the properties of the underlying microstructure?

The inverse problem is important because (1) learning the material microstructure is often of high practical interest, and (2) EBSD is expensive and destructive so ideally we would like to learn the microstructure from the non-destructive ultrasound measurements. We proceed in two parts below to analyze the different components for establishing the forward model, which is needed to solve the inverse problem.

5.1 Part 1: Forward Model Analysis from EBSD Data

The EBSD data for the specimen is available as the left image in Figure 24. Ideally, we would like to set this entire field of Euler angles as the inverse problem target. However, that poses an extremely difficult and ill-posed inverse problem, thus we begin by seeking inversion of grain *statistics* of the microstructurel. To obtain these statistics from the EBSD data, we perform reconstruction of the grains using MTEX software [4] as depicted in Figure 25 (left), and from which we can extract the grain size distribution shown in the histogram (right). Similarly, MTEX also provides visualization of the Euler angles, through the orientation distribution functions (ODFs) shown in Figure 26. ODFs with values close to unity everywhere would imply weak texture (close to isotropic material), and the hot and cool spots in for our specimen thus suggests the existence of some anisotropy/texture.

Given the local elasticity tensor for the nickel specimen, the Euler angles for each grain (and thus the rotation matrices), we can compute the elastic constants for each grain in the global coordinate system. Given the relationships of longitudinal and shear wave speeds $v_l = \sqrt{c_{11}/\rho}$, $v_{sl} = \sqrt{c_{44}/\rho}$, $v_{s2} = \sqrt{(c_{11} - c_{12})/(2\rho)}$, we can obtain the local wave velocity speeds corresponding to each grain, as well as their distributions, shown in Figure 27.

5.2 Part 2: Ultrasound Experimental Measurements

In Part 2 we extract the Rayleigh wave speeds that is observed directly from the ultrasound measurements. Figure 28 presents such an example with snapshot at different times, where a 2.25 MHz ultrasound wave is propagated from the right. Here we compute the wavefront at any given time to be identified by the points which undergo their maximum displacement at that time. This allows

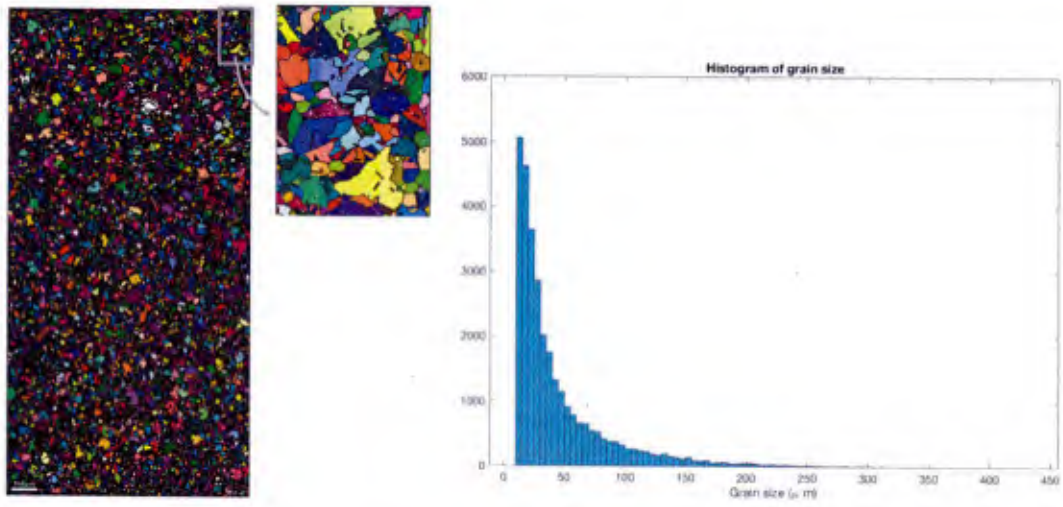

Figure 25: MTEX reconstruction of grains from EBSD data (left) and the corresponding grain size distribution (right).

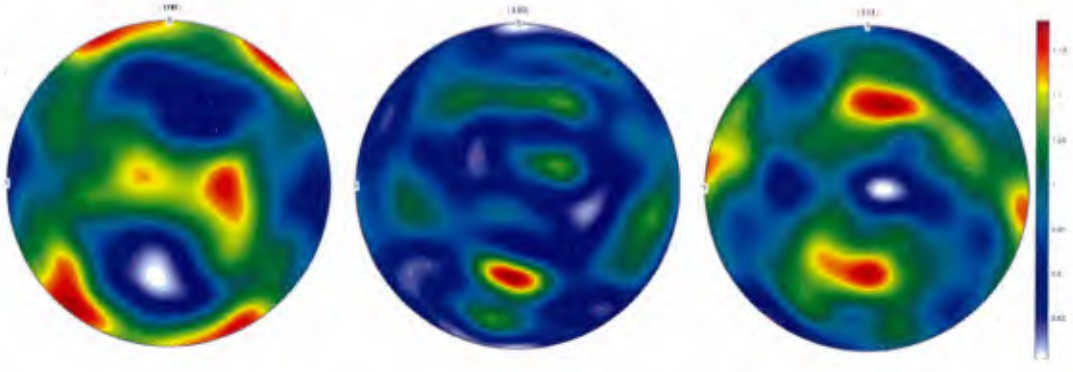

Figure 26: Orientation distribution functions from EBSD data.

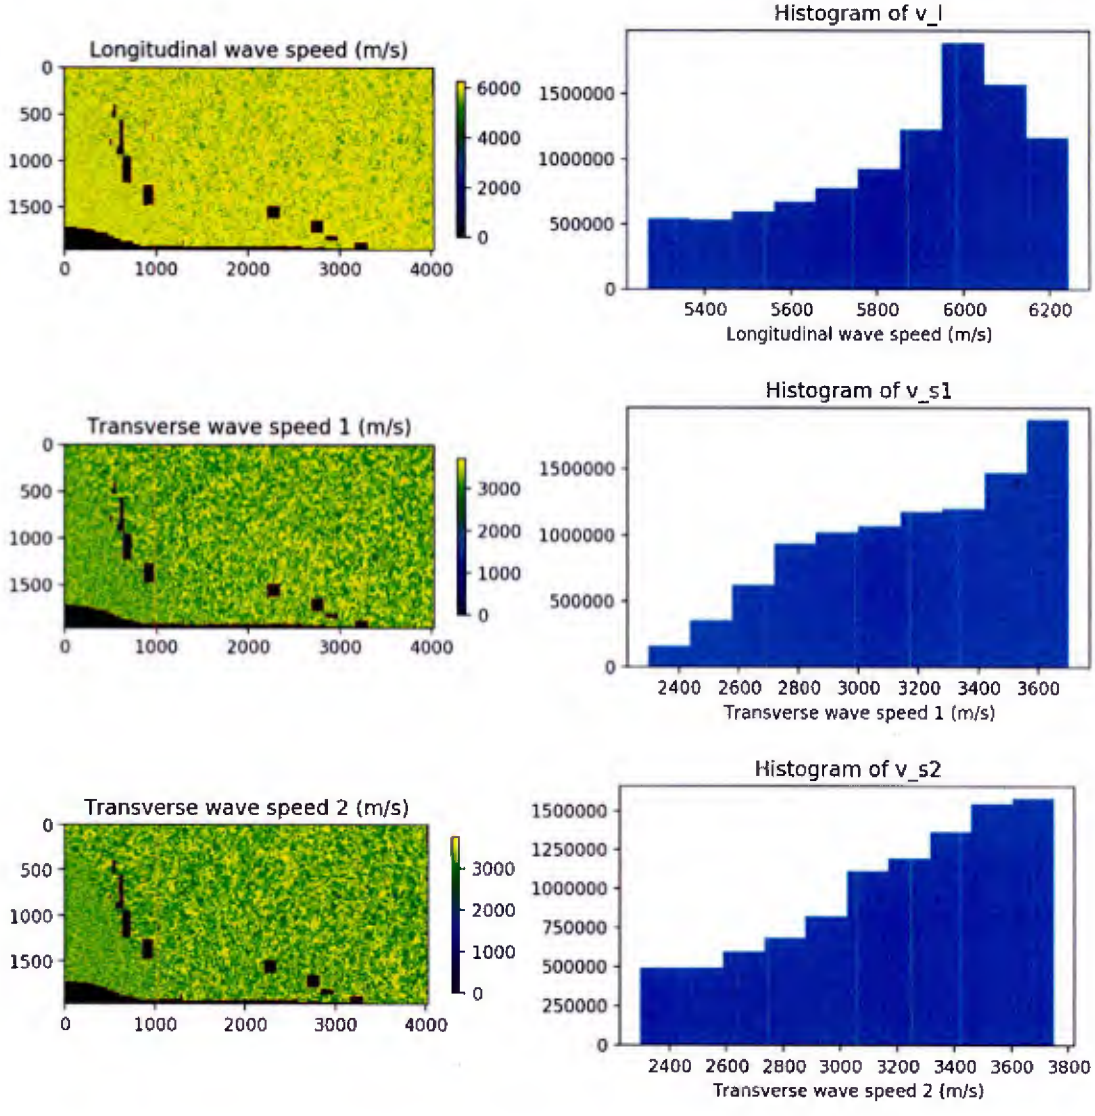

Figure 27: Local wave speeds inside each grain, and their distributions, computed from EBSD data.

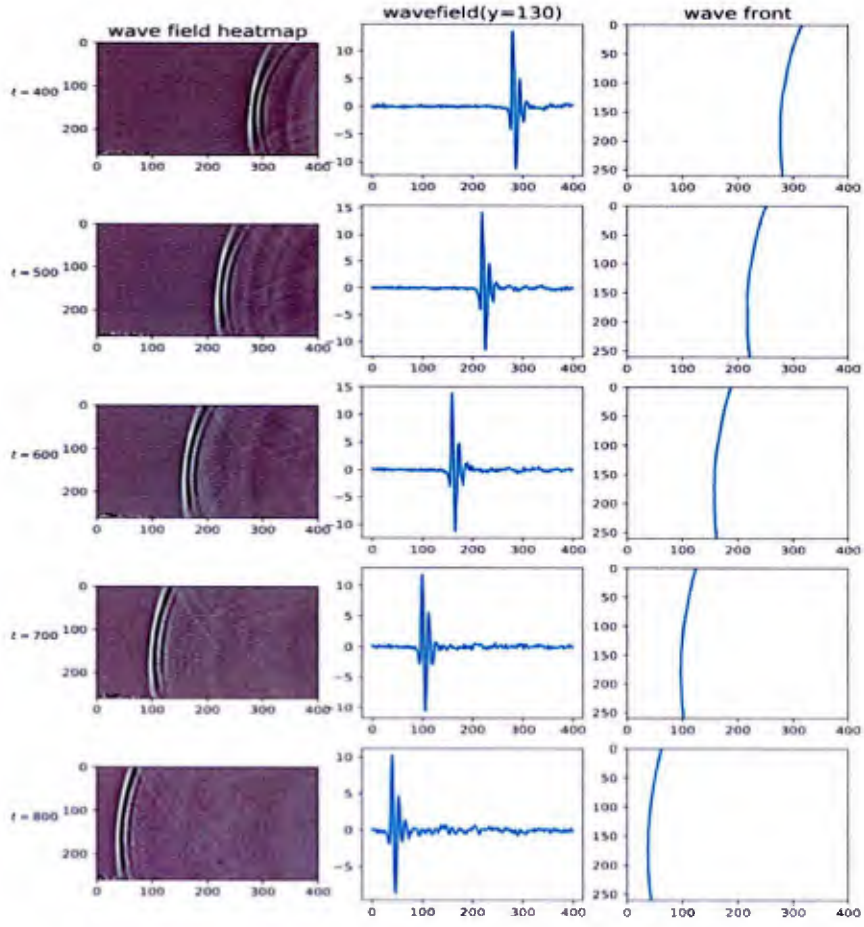

Figure 28: An example of Rayleigh wavefront from the ultrasound experiments. Each row is a different snapshot in time. The left column is the original ultrasound data; middle column is the displacement as seen at a slice of fixed y; right column is the digitized wavefront.

us to "digitize" the wavefront as shown by the blue curves in the third column of the figure. Next, we estimate the wavefront speed by performing a finite difference of the weavefront location along the normal direction of the front. Figure 29 illustrates the spatial field of this Rayleigh wavespeed, along with its distribution.

5.3 Comparisons and Conclusions

We can first compare the shear waves predicted by theory from EBSD data: focusing on the histogram for v_{s2}, it has a mean of 3192 m/s, standard deviation 389 m/s, and range 2298–3750 m/s. In comparison, the Rayleigh wavespeeds estimated from the ultrasound data has mean 2800 m/s, standard deviation 235 m/s, and range 2043–3483 m/s. Furthermore, the theory relating Rayleigh and shear wave speeds for *isotropic* matrial is

$$v_R = v_s \frac{0.862 + 1.14\nu}{1 + \nu}, \tag{42}$$

which translates to about $v_R = 0.92 v_s$ if assuming a $\nu = 0.3$. Our results above present a ratio of 0.88, which is quite close to 0.92. Overall, we see that the theory and observations are fairly consistent.

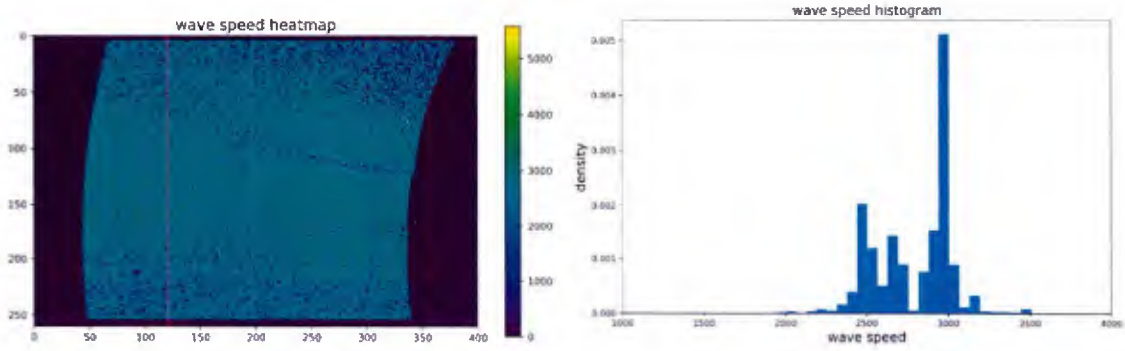

Figure 29: Rayleigh wavefront speeds computed from ultrasound measurement data, presented on the original specimen field (left) and as a histogram distribution (right).

However, the current ultrasound measurements do not give us enough information about grain size and orientation. We hypothesize is is due to the ultrasound wavelengths ($\sim 700\ \mu m$ from our estimate) is greater than the typical grain sizes in this specimen $\sim 50 - 250\ \mu m$. As a result, the ultrasound waves are not able to "see" the much smaller grains, and thus presents an information limit of the currently available data. If additional data are available with ultrasound wavelengths smaller (sub-grain size), then we anticipate to observe

- spatial extent of velocity fluctuation to indicate grain sizes; and
- values of fluctuation would indicate grain orientations.

With these new data, we can then build a forward model that allows us to learn more about the underlying microstructure, where we plan to use system identification to find ignored/missing physics from the microstructure to the Rayleigh waves, Bayesian inference to estimate uncertainty of these inverse problem solutions, and optimal experimental design to suggest additional experiments that can provide the greatest information gain.

6 CONCLUSIONS

The AIMM-Discover AIRA project has made considerable progress when viewed in the context of our original "Challenge goals".

The main aim of gaining generalizable prediction and subsequent control of microstructure with quantifiable uncertainty has been achieved. Following discovery, prediction was achieved as demonstrated by the results on the evolution of alloy microstructure in Section 4.1.2, and in the publication on system identification in a Bayesian setting [25]. The form of equations discovered could guide experimental control. While such control has not been demonstrated, we have made a step toward it with our PCA-guided auto-encoder approach to classifying mechanisms and inverting for the deformation of polymer gels, as we describe in the following paragraphs. This is a step toward ultimate design of chemically driven swimming motion to propel polymer gels. Thus, our framework of Bayesian inversion, and deterministic Variational System Identification allow discovery of materials physics with subsequent prediction of their evolution beyond the regimes attainable in experiments. Design will follow.

In section 2.3 we described a Bayesian system identification that is robustly able to identify dynamical systems from sparse and noisy data. The primary benefit of this approach is that it

rigorously derives an appropriate learning objective, rather than heuristically choosing a least-squares misfit objective and a suitable regularization term. As such, decision making rules could be used to extract optimal estimates of the true underlying system.

However, this rigorous specification of the learning problem is significantly more challenging to solve computationally. In the context of PDEs, it requires an acknowledgement that *any model parameterization is wrong* and one should not trust the evolution to match the data at all times. To account for this uncertainty, it is required to specify a stochastic PDE and then marginalize out all possible states of the system. In this project, we have employed simplifications of this computational expensive procedure by using linearization-based extended Kalman filters for this task. However, more work needs to be done to develop "discretization-free" approaches — those that recognize and discover the smallest possible dimensionality of the system.

In section 2.3 we also described an inverse-design outer loop. Future work will couple this outer loop task with the system identification. Indeed, in this projct we obtained data from a *known* forward model to then feed into the inverse design process. When we couple this outer loop to the inverse learning process, the data will be used to first identify a forward model and then to use it generating further simulated data. In this way both simulated and experimental data could be used to train the inverse-design procedure in cases where the forward model is uncertain.

We have achieved automated multiscale learning between quantum and statistical mechanics, allowing larger scale computations of evolving materials microstructure at the continuum scale (Section 4.1.1 and Refs [2, 20]). To the best of our knowledge, this is the most complete demonstration of scale bridging using thermodynamic information in materials from quantum mechanical scales, through statistical mechanics to the continuum. This is an important step toward the long-standing challenge of bridging scales in materials physics, and brings us closer to first principles based prediction of materials behavior of a broad spectrum of materials relevant to energy applications (batteries and solar power, ferro-electrics), biology (soft tissue, cellular processes, bacteria).

In the context of physics discovery, we also mention that, in very recent results (June, 2020) we have been able to show that Variational System Identification recognizes that the evolution of materials microstructure is not governed by second-order dynamics, such as Newtonian mechanics, but is first-order in time. This demonstrates re-discovery of fundamental laws of nature. This result will appear soon in a paper under preparation.

The above successes notwithstanding, completely unguided discovery remains a long way off; entirely hidden physics cannot be discovered. For example, we need to know how many components are in an alloy. If the data reveal nothing about the dynamics by which certain components of an alloy evolve–not even their presence in the system–then the methods developed in AIMM-Discover will not discover the related internal mechanisms. However, phenomenology bounded by coarser-grained physics will still be identified. This approach to systematic, efficient and automated discovery of physics of materials [24, 23] was originally developed in the context of the problem of microstructural evolution. It has been extended to other systems as shown by our application to deformation mechanisms of soft materials (Section 4.1.2). Thus, our methods generalize well.

Using our recently developed algorithms for inverse DFT we obtained the exact exchange-correlation (XC) potentials for a range of molecules. We used the data on exact XC potentials to evaluate some of the widely used models for XC functionals, which revealed consistent deficiencies, both qualitative and quantitative, of these XC models. This revealed the deficiency and insufficiency of the functional forms employed in these models.

The lessons learnt over the coarse of the AIMM-Discover are worthwhile repeating here. Numerically generated data make discovery easy. Smooth and complete data are certainly ideal; however sometimes noise can help condition systems. For example, in experimental design, where we may use the information matrix generated by the mathematical operators being considered for system

identification, noise can prevent the appearance of linear dependence and avoid indeterminacy (zero determinants of the information matrix).

Real experiments are messy. In the field of experimental materials characterization, it is extremely difficult to obtain spatially registered, high temporal resolution data. Nearly all microstructure data is measured on different experiments on what is nominally the same system. This rules out the methods that work, for instance, to discover the physics of fluids in the classical Eulerian experiments of tracking quantities at fixed positions in space. We therefore abandoned spatio-temporally registered data early on in the AIMM-Discover project. However, we developed a mathematically rich and general framework in response to this challenge. It now makes the discovery of the physics of material microstructure evolution possible within the reality of spatio-temporally unregistered data.

During the development of multiscale learning for scale bridging, we were faced with the challenge that the identification of quantities as inputs and outputs is not preserved by theoretical frameworks across scales. In statistical mechanics computations, an order parameter is the result of a (usually) Monte Carlo computation in which the chemical potential is the controlled quantity. However, the continuum thermodynamic treatment uses a formalism in which chemical potentials are functions of order parameters, and invertibility is not guaranteed. This presents a challenge to the theoretical framework of scale bridging. We have developed AI methods, centered on active learning to meet this challenge (Section 2.2.1 and Ref [20])

In general, we contended with the lack of statistical mechanics data for complex systems, relevant and quality experimental data (as alluded to above). However, this was surmounted in the case of materials microstructure by developing Variational System Identification for sparse, spatio-temporally unregistered data.

The choice of criteria for optimal experimental design is non-obvious. As we carried out OED using the *classical* setup seeking to minimize the covariance determinant, we have come to realize the inconsistency between this measure and the objective of performing system identification—to reveal active PDE operators. As such, we proceeded to construct *goal-oriented* OED criteria that better reflect this object, such as with the number of correctly identified operators. However, there are numerical challenges, e.g., difficulty to optimize these discrete-valued quantities. These investigations spurred broader interest in setting up the problems for AI: we need pose the right questions for our AIRA in order to obtain results that are meaningful and that practitioners will actually adopt.

As we move forward we intend to apply AIMM-Discover methods to a multitude of experimental systems. In addition to pattern formation in material, which was the setting for much of the data in this project, we will continue to apply these methods to the mechanics of soft polymers, including biological tissue. Other work is now starting to apply Variational System Identification to pattern formation that is crucial to the functioning of cellular processes in bacteria.

The deficiencies revealed by the exact XC potentials points to an important direction data-driven modeling can provide in improving the XC functional description in DFT, which remains the holy grail of ab-initio modeling. To this end, a tightly integrated on-the-fly data-driven framework, using exact XC potentials and the known exact conditions on XC functionals, can be employed in addressing one of the most outstanding problems in materials modeling. The insights obtained from this DARPA project provides the foundation for such an effort.

6.1 Next steps for AIRA: AI for Discontinuous Scientific Discovery

We propose that the community direct attention toward *AI for Discontinuous Scientific Discovery*. While this iteration of the program has made progress toward various aspects of discovery in

physics, the developments have all been in the framework of continuous learning. Loss functions, regression, Bayes' Theorem and many other instruments do not address the question of how "phase transitions" in understanding happen. Can we build on the structured discovery demonstrated in AIRA to make leaps of intuition? This will need to include *de novo* cognition at the least, from the wider range of AI methods that go beyond machine learning alone.

References

[1] CASM: *First-principles statistical mechanical software for the study of multi-component crystalline solids*, https://github.com/prisms-center/CASMcode.

[2] *Machine learning materials physics: Integrable deep neural networks enable scale bridging by learning free energy functions*, Computer Methods in Applied Mechanics and Engineering, 353 (2019), pp. 201 – 216, https://doi.org/https://doi.org/10.1016/j.cma.2019.05.019, http://www.sciencedirect.com/science/article/pii/S0045782519302889.

[3] S. Alben, A. A. Gorodetsky, D. Kim, and R. D. Deegan, *Semi-implicit methods for the dynamics of elastic sheets*, Journal of Computational Physics, 399 (2019), p. 108952.

[4] F. Bachmann, R. Hielscher, and H. Schaeben, *Grain detection from 2d and 3d EBSD dataSpecification of the MTEX algorithm*, Ultramicroscopy, 111 (2011), pp. 1720–1733, https://doi.org/10.1016/j.ultramic.2011.08.002.

[5] M. E. Casida, *Correlated optimized effective-potential treatment of the derivative discontinuity and of the highest occupied Kohn-Sham eigenvalue: A Janak-type theorem for the optimized effective-potential model*, Phys. Rev. B, 59 (1999), pp. 4694–4698.

[6] E. V. R. de Castro and F. E. Jorge, *Accurate universal Gaussian basis set for all atoms of the periodic table*, J. Chem. Phys., 108 (1998), pp. 5225–5229.

[7] T. H. Dunning, *Gaussian basis sets for use in correlated molecular calculations. i. the atoms boron through neon and hydrogen*, J. Chem. Phys., 90 (1989), pp. 1007–1023.

[8] N. Galioto and A. Gorodetsky, *Bayesian system id: Optimal management of parameter, model, and measurement uncertainty*, arXiv preprint arXiv:2003.02359, (2020).

[9] J. G. Goiri and A. Van der Ven, *Phase and structural stability in Ni-Al systems from first principles*, Physical Review B, 94 (2016), https://doi.org/10.1103/PhysRevB.94.094111.

[10] S. Greydanus, M. Dzamba, and J. Yosinski, *Hamiltonian neural networks*, in Advances in Neural Information Processing Systems, 2019, pp. 15353–15363.

[11] X. Huan and Y. M. Marzouk, *Simulation-based optimal Bayesian experimental design for nonlinear systems*, Journal of Computational Physics, 232 (2013), pp. 288–317, https://doi.org/10.1016/j.jcp.2012.08.013.

[12] B. Kanungo, P. M. Zimmerman, and V. Gavini, *Exact exchange-correlation potentials from ground-state electron densities*, Nat. Commun., 10 (2019), pp. 1–9.

[13] M. Khalil, A. Sarkar, S. Adhikari, and D. Poirel, *The estimation of time-invariant parameters of noisy nonlinear oscillatory systems*, Journal of Sound and Vibration, 344 (2015), pp. 81 – 100.

[14] E. Marquis, *A reassessment of the metastable miscibility gap in Al-Ag alloys by atom probe tomography*, Microscopy and Microanalysis, 13 (2007), pp. 484–492, https://doi.org/10.1017/S1431927607070870.

[15] A. R. Natarajan, J. C. Thomas, B. Puchala, and A. Van der Ven, *Symmetry-adapted order parameters and free energies for solids undergoing order-disorder phase transitions*, Phys. Rev. B, 96 (2017), p. 134204, https://doi.org/10.1103/PhysRevB.96.134204, https://link.aps.org/doi/10.1103/PhysRevB.96.134204.

[16] J. P. Perdew, K. Burke, and M. Ernzerhof, *Generalized gradient approximation made simple*, Phys. Rev. Lett., 77 (1996), pp. 3865–3868. https://doi.org/10.1103/PhysRevLett.77.3865, https://link.aps.org/doi/10.1103/PhysRevLett.77.3865.

[17] J. P. Perdew and M. Levy, *Comment on "Significance of the highest occupied Kohn-Sham eigenvalue"*, Phys. Rev. B, 56 (1997), pp. 16021–16028.

[18] J. P. Perdew and Y. Wang, *Accurate and simple analytic representation of the electron-gas correlation energy*, Phys. Rev. B, 45 (1992), pp. 13244–13249, https://doi.org/10.1103/PhysRevB.45.13244, https://link.aps.org/doi/10.1103/PhysRevB.45.13244.

[19] S. Särkkä, *Bayesian Filtering and Smoothing*, Institute of Mathematical Statistics Textbooks, Cambridge University Press, 2013.

[20] G. Teichert, A. Natarajan, A. V. der Ven, and K. Garikipati, *Scale bridging materials physics: Active learning workflows and integrable deep neural networks for free energy function representations in alloys*, arXiv:2001.05646, (2020).

[21] P. Voorhees and W. C. Johnson, *The thermodynamics of elastically stressed crystals*, vol. 59 of Solid State Physics, Academic Press, 2004, pp. 1 – 201, https://doi.org/https://doi.org/10.1016/S0081-1947(04)80003-1, http://www.sciencedirect.com/science/article/pii/S0081194704800031.

[22] Y. Wang, D. Banerjee, C. Su, and A. Khachaturyan, *Field kinetic model and computer simulation of precipitation of l12 ordered intermetallics from f.c.c. solid solution*, Acta Materialia, 46 (1998), pp. 2983 – 3001, https://doi.org/https://doi.org/10.1016/S1359-6454(98)00015-9, http://www.sciencedirect.com/science/article/pii/S1359645498000159.

[23] Z. Wang, X. Huan, and K. Garikipati, *Identification of the partial differential equations governing microstructure evolution in materials: Inference over incomplete, sparse and spatially non-overlapping data*, arXiv:2001.04816.

[24] Z. Wang, X. Huan, and K. Garikipati, *Variational system identification of the partial differential equations governing the physics of pattern-formation: Inference under varying fidelity and noise*, Computer Methods in Applied Mechanics and Engineering, 356 (2019), pp. 44 – 74, https://doi.org/https://doi.org/10.1016/j.cma.2019.07.007.

[25] Z. Wang, B. Wu, K. Garikipati, and X. Huan, *A perspective on regression and bayesian approaches for system identification of pattern formation dynamics*. arXiv:2001.05646.

[26] P. M. Zimmerman, *Incremental full configuration interaction*, J. Chem. Phys., 146 (2017), p. 104102.