

Enabling in-time Prognostics with Surrogate Modeling through Physics-enhanced Dynamic Mode Decomposition Method

Katelyn Jarvis¹, Matteo Corbetta², Christopher Teubert³, and Stefan Schuet⁴

^{1,3,4} NASA Ames Research Center, Moffett Field, CA 94035

katelyn.j.jarvis@nasa.gov

christopher.a.teubert@nasa.gov

stefan.r.schuet@nasa.gov

² KBR Inc., NASA Ames Research Center, Moffett Field, CA 94035

matteo.corbetta@nasa.gov

ABSTRACT

Computational models provide essential quantitative tools for assessing and predicting the health and performance of physical systems. However, high-fidelity models are rarely used in real-time operations or large optimization loops, due to their time-intensive nature. A common approach to improving computational efficiency of prognosis is to employ surrogate models. Such models can significantly decrease computation time for some accuracy loss. In this context, use of Dynamic Mode Decomposition (DMD) is proposed to generate surrogate models for lithium-ion (Li-ion) battery discharge. DMD has been suggested and used successfully in the area of fluid dynamics for over a decade, but it has not been applied to the PHM domain, where far-ahead prediction of nonlinear behavior is crucial to propagate faults or predict Remaining Useful Life (RUL). For Li-ion battery health management, the standard application of DMD using only the observable quantities of interest was unable to capture the nonlinear discharge of batteries exhibited in lab testing. The Koopman theory, however, provides a mechanism to trade-off low dimensional nonlinear models with high-dimensional linear ones in a DMD framework, by augmenting nonlinear state variables into the system representation. In this way, DMD allows for configurable simulation accuracy dependent on the dimensionality of the Koopman operator. For battery health management, we augmented the observable variables with the hidden states of a higher-fidelity physics model to build the DMD surrogate. In comparison to a high-fidelity model, the surrogate improved computational efficiency with only a minimal loss of accuracy, and enabled long-term prog-

nostics horizons. A generalized method for this was implemented in the `prog_models` python package.

1. MOTIVATION

The field of prognostics and health management involves developing quantitative tools to estimate the health of physical systems and predict how system degradation will evolve with use. To aid in risk monitoring and mitigation, prognostics algorithms must be both accurate and timely. Often, high accuracy is attained by representing physical systems with high-fidelity models, which can be computationally inefficient. For example, in some resource-constrained scenarios, such as small Uncrewed Aerial Systems (sUAS), small satellites, or spacecraft, the computational cost of high-fidelity models is prohibitive to achieve real-time or quasi real-time prognostics. This is especially true in cases where prognostics is performed inside an optimization loop, such as an application where system loading is optimized to maximize useful life.

Surrogate, or reduced-order, modeling is a computationally efficient alternative to more detailed physical models. In surrogate modeling, a lower-fidelity, faster version of the model (i.e., the surrogate) is generated and used in place of the original high-fidelity model. This approach can be used to speed up complex simulations when the need for faster prediction outweighs the need for high accuracy. For example, in simulations of advanced air mobility vehicles, surrogate models of the battery and powertrain can provide an alternative to physical models that describe the full electro-chemistry of the battery and the high-frequency electronic speed controller, leading to faster simulations when highly detailed modeling is not necessary to assess the system performance.

One such surrogate modeling technique is Dynamic Mode Decomposition (DMD), a powerful data-driven tool used to

Katelyn Jarvis et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

identify approximate linear dynamics from high-dimensional data (Kutz, Brunton, Brunton, & Proctor, 2016). Developed originally within the computational fluid dynamics community by Schmid (Schmid, 2010), DMD and its variants have been successfully used for more than a decade. The method is remarkably simple and requires minimal assumptions about the underlying physical system. In its standard form, DMD aims at approximating the nonlinear dynamics of a system through a linearized model expressed as

$$X' = AX,$$

where X is a “snapshot” matrix, with each column consisting of a vector of observable data from the system at successive points in time (i.e. column one includes all observable data at t_0 , column two includes data at t_1 , etc.). Similarly, X' contains the same observable data one step forward in time (i.e. from t_1 onward). Once constructed, the linear approximation can be solved for by a simple matrix inversion, and produces accurate look-ahead predictions for short time horizons (Kutz et al., 2016). Its simplicity is also supported by the computational cost, limited to the computation time of a single Singular-Value Decomposition (SVD) for the regression stage. However, while proven to be a valuable tool in specific applications for short time scales, DMD has yet to be explored extensively in a prognostics setting where the system state dimension is not always large, fast computations can be crucial in certain engineering applications, and importantly, long prediction horizons are necessary.

As motivation for using DMD in a prognostics scenario, we performed some exploratory work using DMD to predict the voltage discharge curve of a lithium-ion (Li-ion) battery subject to step-wise randomized current input. The discharge curve was generated using a well-established, nonlinear electro-chemistry model (Daigle & Kulkarni, 2013). The relationship between input current and output voltage is highly non-linear, and thus provides an ideal opportunity for DMD to approximate this complex system with a linear model. As in a standard DMD application, we construct a snapshot matrix with only output quantities. However, in the proposed battery case study, there is just one output quantity (i.e. voltage) at each point in time. As in previous work with limited state dimensions, we attempted to adjust for this by composing each column of the snapshot matrix with a set of lagged time series of voltage values, similarly to an autoregressive model (i.e. for a configurable model order, j , the first column of the matrix would include the voltage values from time t_0 to t_j , the second column from t_1 to t_{j+1} , and so on; see (Tirunagari, Kouchaki, Poh, Bober, & Windridge, 2017)). However, applying DMD in this way, the model is unable to capture the nonlinear dynamics of voltage discharge, thus failing to produce acceptable predictions, Fig. 1.

This poor linearization forced us to explore alternatives to the




Figure 1. Prediction of a voltage discharge curve using DMD with one-dimensional, lagged time series of voltage values.

use of observable states only. An extension of DMD was presented in (Williams, Kevrekidis, & Rowley, 2015), where the authors proposed to augment the set of observables with a set of basis functions that would aid the regression. By doing so, the increased number of rows in the snapshot matrices increases the dimensionality of the model and its ability to reproduce the reduced-order dynamics (this approach connects DMD with Koopman theory (Koopman, 1931)). However, the selection of the basis functions is arbitrary, and there is no guarantee that the functions selected for one application would work seamlessly for others. For the same battery case study, we tried using radial basis functions of type “thin plate spline” following the approach in (Korda & Mezić, 2018), where one hundred basis functions with random control points were used to augment the dimension of the observable states. Korda and Mezić successfully found a finite-dimensional approximation of the Koopman operator for the nonlinear systems they studied, and applied it to model a predictive control problem. Their use-case required a limited number of predictions steps, given the high-frequency nature of feedback controls. In our case, however, the success encountered in short-term state forecasts did not translate to ac-

curate long-term predictions of voltage discharge curves, regardless of our many attempts to tune the model by changing the type and number of basis functions.

In this work, we explore an alternative extension of DMD using the system internal states in the snapshot matrices. Having the ability to access the internal states and equations of the high-fidelity model (as in the Li-ion battery example) provides the opportunity to include physics-based information, instead of basis functions that are not driven by physical processes. We find that a physics-enhanced DMD allowed us to accurately capture battery voltage discharge predictions over a number of blind tests. Thus, this work highlights the potential for using physics-enhanced DMD for long-term predictions.

The rest of the paper is structured as follows: Section 2 summarizes DMD and the basic equations necessary to implement the algorithm. Section 3 shows the application of the physics-enhanced DMD as a surrogate for a simple toy problem (Sec. 3.1) and voltage discharge predictions generated through the electro-chemistry battery model (Secs. 3.2 - 3.5). Section 4 concludes the paper with a critical analysis and suggestions for future work.

2. BACKGROUND

This section summarizes the fundamentals of the DMD algorithm, its connection with the Koopman operator, the SVD-based approach to find A in large scale problems, and the application of DMD with control inputs. Most of the material was retrieved from (Kutz et al., 2016; Williams et al., 2015), and the interested reader is referred to those publications for more details.

We define the discrete dynamical system similarly to (Kutz et al., 2016) as follows:

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{F}(\mathbf{x}_k), \\ \mathbf{y}_{k+1} &= \mathbf{g}(\mathbf{x}_{k+1}), \end{aligned}$$

where $\mathbf{x}_k \in \mathbb{R}^{n \times 1}$ is the vector of states representing the dynamical system at time $t_k \in \mathbb{R}^{\geq 0}$, $k \in \mathbb{N}$. Vector $\mathbf{y}_k \in \mathbb{R}^{q \times 1}$ represents the measurement vector. Symbols $\mathbf{F}(\cdot) : \mathbb{R}^{n \times 1} \rightarrow \mathbb{R}^{n \times 1}$ and $\mathbf{g}(\cdot) : \mathbb{R}^{n \times 1} \rightarrow \mathbb{R}^{q \times 1}$ represent the state-space functions.

It should be noticed that, in the DMD context, the measurement function $\mathbf{g}(\cdot)$ may be composed of equations derived from physical relationships between the hidden states and the “observable” $\mathbf{y}$ values, and/or data-driven kernels (such as basis functions) that increase the state vector to a higher dimensionality (Korda & Mezić, 2018), so that DMD can be applied to this higher-dimensional vector $\mathbf{y}$ to find a higher-dimensional linearized system (see Section 2.2 for more details).

2.1. Basics of DMD: Estimating Matrix A

As described in (Kutz et al., 2016), DMD is “an equation-free method capable of providing a decomposition of a complex system into spatiotemporal coherent structures that may be used for short-time future state prediction and control” (chapter 1, page 1). One of the key advantages of DMD is its simplicity. The algorithm starts by using the observable variables of the system to construct snapshot matrices X and X' , with X' lagged by one time step, Δt :

$$\begin{aligned} X(t) &= \begin{bmatrix} x_1(t_0) & \dots & x_1(t_i) & \dots & x_1(t_{m-1}) \\ x_2(t_0) & \dots & x_2(t_i) & \dots & x_2(t_{m-1}) \\ \dots & \dots & \dots & \dots & \dots \\ x_n(t_0) & \dots & x_n(t_i) & \dots & x_n(t_{m-1}) \end{bmatrix}, \\ X'(t) &= \begin{bmatrix} x_1(t_1) & \dots & x_1(t_{i+1}) & \dots & x_1(t_m) \\ x_2(t_1) & \dots & x_2(t_{i+1}) & \dots & x_2(t_m) \\ \dots & \dots & \dots & \dots & \dots \\ x_n(t_1) & \dots & x_n(t_{i+1}) & \dots & x_n(t_m) \end{bmatrix}, \end{aligned}$$

Thus $X, X' \in \mathbb{R}^{n \times m}$. For simplicity of notation, we will describe the dynamical system in its discrete form, omitting the time dependency of the matrices. In the case of unforced systems, DMD aims at finding the constant matrix A to satisfy equality $X' \approx AX$, and the best fit can intuitively be obtained by:

$$A = X' X^+. \quad (1)$$

where $^+$ denotes the inverse (we use the Moore-Penrose pseudo-inverse), and $A \in \mathbb{R}^{n \times n}$. This is equivalent to minimizing the Frobenius norm of the residuals $X' - AX$, (Kutz et al., 2016). Once A is computed, the dynamics can then be approximated with discrete first order integration $\mathbf{x}_{k+1} \approx A \mathbf{x}_k$.

2.2. Connection with Koopman Operator

DMD attempts to approximate a nonlinear system with a linear one. It can be viewed as a finite-dimensional approximation of the Koopman operator, $\mathcal{K}$ (Koopman, 1931; Kutz et al., 2016), which is an infinite-dimensional, composition operator for $\mathbf{g}$ and $\mathbf{F}$:

$$\begin{aligned} \mathcal{K} g &= g \circ \mathbf{F}, \\ \mathcal{K} g(\mathbf{x}_k) &= g(\mathbf{F}(\mathbf{x}_k)) = g(\mathbf{x}_{k+1}), \end{aligned}$$

The challenge of the approach mainly lies in the selection (or definition) of function $\mathbf{g}(\cdot)$ such that DMD matrix A built upon the higher-dimensional observables $\mathbf{y}_k$ is a satisfactory finite-approximation of $\mathcal{K}$. As stated in (Kutz et al., 2016), this is different from linearizing the system dynamics, where it is necessary to define an equilibrium point, $\mathbf{x}_0$, around which the linearization is valid. In the next section, we will utilize this principle by introducing hidden states of the physical model as observables to build matrix A .

2.3. DMD for Large-Scale Problems

Large state dimensions n cause the inversion of X to be numerically intractable, thus requiring rank reduction to build the matrix A , such as singular-value decomposition (SVD) (Stewart, 1993). The problems we present in the next sections do not belong to this category, but we provide a brief summary of the SVD-based algorithm for completeness. The interested reader is referred to (Kutz et al., 2016) for more details.

Matrix X is first split into U, Σ, V^* according to SVD, where U and V are the left and right eigenvectors of X , respectively, Σ is a square diagonal matrix of eigenvalues, and superscript $*$ indicates the complex conjugate. Matrix A is calculated using the pseudo-inverse of X through SVD:

$$A = X' V \Sigma^{-1} U^*. \quad (2)$$

To reduce dimensionality, matrices U, Σ, V are truncated at the desired reduced rank r such that $\tilde{U} \in \mathbb{C}^{n \times r}$, $\tilde{\Sigma} \in \mathbb{C}^{r \times r}$, $\tilde{V} \in \mathbb{C}^{m \times r}$. The rank-reduced representation of A , $\tilde{A}$, can be obtained from Eq. (2) by pre- and post-multiplying A by U^* and U , respectively, which leads to:

$$\tilde{A} = \tilde{U}^* X' \tilde{V} \tilde{\Sigma}^{-1},$$

thus obtaining the reduced-rank dynamical system: $\tilde{x}_{k+1} = \tilde{A} \tilde{x}_k$. The full state vector can be retrieved from $x_k = U \tilde{x}_k$.

2.4. DMD with Control

The goal of DMD with control is to estimate matrices A and B that approximate the nonlinear system with a linear one, written in canonical form:

$$x_{k+1} = A x_k + B u_k.$$

In this case, in addition to X and X' , DMD requires the construction of the input snapshot matrix Υ :

$$\Upsilon = \begin{bmatrix} u_1(t_0) & \dots & u_1(t_i) & \dots & u_1(t_{m-1}) \\ u_2(t_0) & \dots & u_2(t_i) & \dots & u_2(t_{m-1}) \\ \dots & \dots & \dots & \dots & \dots \\ u_q(t_0) & \dots & u_q(t_i) & \dots & u_q(t_{m-1}) \end{bmatrix},$$

which leads to the state equation in matrix form:

$$X' \approx A X + B \Upsilon.$$

To solve, we can simply stack matrices X and Υ along the state dimension, i.e., $\Omega = \begin{bmatrix} X \\ \Upsilon \end{bmatrix}$, and then run DMD by substituting Ω to the original X . Matrices A and B then come from the splitting of the left eigenvectors U into two separate components (Kutz et al., 2016):

$$\begin{aligned} A &\approx X' V \Sigma^{-1} U_1^*, \\ B &\approx X' V \Sigma^{-1} U_2^*, \end{aligned}$$

where U_1 and U_2 are the components of U such that $U_1 \in \mathbb{C}^{n \times r}$, $U_2 \in \mathbb{C}^{q \times r}$, and $U = \begin{bmatrix} U_1 \\ U_2 \end{bmatrix}$.

It should be noticed that if a reduced-rank approximation of A, B is desired (i.e., $r < \text{rank}(X)$), the user should replace the full-rank matrices with the reduced rank ones and follow the steps in Subsection 2.3.

3. PHYSICS-ENHANCED DMD: APPLICATION AND RESULTS

3.1. Toy Example: DMD and Thrown Object

As a preliminary test as to whether a physics-enhanced DMD method is effective for long-term prediction horizons, we applied DMD to the simple “thrown object” model from the *prog.models* python package (Teubert, Corbetta, Kulkarni, Jarvis, & Daigle, 2022). The *ThrownObject* model is a simple nonlinear model that describes the 1-D dynamics of a mass thrown vertically into the air, using a standard drag equation with coefficient C_D producing the resistance force. The model is written in state-space canonical form using state vector $x = [x, \dot{x}]^\top$, and measurement vector $y = y = x$,

$$\begin{aligned} \frac{dx}{dt} &= \dot{x} \\ \frac{d\dot{x}}{dt} &= -g - \frac{\rho C_D A}{2m} |\dot{x}| \dot{x} \end{aligned} \quad (3)$$

where g is gravity, m is the object’s mass, A is the object reference area, and ρ is the air density. The reference frame is a unit vertical axis pointing upwards normal to the ground. For this case, we ignore the effect of the small altitude changes to air density.

Beginning with a standard DMD approach, we constructed a snapshot matrix consisting of a lagged time series of only the observable values, which in this case is the position of the object (as in (Tirunagari et al., 2017)). Next, we considered the addition of physics-based information into the snapshot matrix. To do so, we construct a snapshot matrix that includes internal states at each point in time, such that both position and velocity are included in X and X' , giving

$$X = \begin{bmatrix} x(t_0) & \dots & x(t_i) & \dots & x(t_{m-1}) \\ \dot{x}(t_0) & \dots & \dot{x}(t_i) & \dots & \dot{x}(t_{m-1}) \end{bmatrix} \quad (4)$$

and the corresponding X' lagged by one time step.

Further, we explore the configurability of a “physics-enhanced” DMD by augmenting the state vector with $\dot{x}^2$ (i.e. same as Eq. 4 with additional row of $\dot{x}^2$ throughout time). This addition is obvious for this specific toy problem, as $\dot{x}^2$ appears in the state-space model in Eq. (3). However, which additional terms to include may not be clear when creating surrogates of complex models. It should be noticed, nonetheless, that gravity should be introduced in the matrices




Figure 2. Prediction of a thrown object trajectory using DMD with (i) one-dimensional, lagged time series of position, (ii) position and velocity, (iii) position, velocity, and velocity-squared. Velocity (bottom) was not predicted with the lagged time series approach as we only introduced x in the algorithm.

X and X' as a constant value, i.e., augmenting the state vector with g . By so doing, the linear case with $C_D = 0$ can be reconstructed even with no additional physics-related variables. This is strictly related to the addition of a “constant” to be estimated in linear regression, corresponding to the zero-order parameter or intercept term. Regardless of the introduction of g , DMD cannot reconstruct the model perfectly when $C_D > 0$, because this system has no “closure” with respect to its nonlinearity; it cannot be perfectly reconstructed with a finite-dimensional linear system (Kutz et al., 2016).

We compared the outcomes of the three scenarios: (i) lagged time series, (ii) x , $\dot{x}$, and (iii) x , $\dot{x}$, $\dot{x}^2$, in the snapshot matrices, against a ground truth test case. The initial condition of the object was set to $x_0 = 1.83$ m and $\dot{x}_0 = 20$ m/s, and we used the following parameters:

$$\begin{aligned} C_D &= 1.1, \\ g &= 9.81 \text{ m/s}^2, \\ m &= 0.02 \text{ kg}, \\ A &= 0.01 \text{ m}^2, \text{ and} \\ \rho &= 1.225 \text{ kg/m}^3 \end{aligned}$$

After estimating matrix A , and setting the appropriate initial conditions, we then applied A iteratively to the augmented state vector to approximate the dynamics throughout time. The results illustrate that physics-enhanced DMD performs better than a standard DMD approximation, and that the introduction of $\dot{x}^2$ allows for even more accurate predictions (Fig. 2). The significant improvement of the DMD approximation upon the addition of internal states in the snapshot




Figure 3. Root mean square error of the reconstructed position (top) and velocity (bottom) using DMD, as a function of the drag coefficient C_D . The introduction of $\dot{x}^2$ in matrices X , X' improves predictions.

matrix highlights the benefit of including physics-based information in DMD.

As a final consideration, we explored the approximation accuracy as the thrown object model became increasingly nonlinear, by sweeping through C_D values from 0 (linear) up to an unreasonably high drag coefficient of 7.0. Training and testing the model for various drag coefficients, we find that the accuracy with which DMD captures the system is nearly constant, but the root mean square error obtained by including $\dot{x}^2$ in the state vector is much lower than the surrogate model with only x and $\dot{x}$ (Fig. 3). As expected, when the model is linear (i.e. $C_D = 0$), both DMD with and without $\dot{x}^2$ can reconstruct the A matrix up to numerical precision, so the error is null.

3.2. Application of DMD to Battery Discharge Prediction

Next we applied DMD to an electro-chemistry based state-of-charge model for a Li-ion battery. The model used as a high-fidelity reference was built upon first principles derived from the work (Karthikeyan, Sikha, & White, 2008), as well as empirical evidence, and the interested reader is referred to (Daigle & Kulkarni, 2013) where the model was first proposed. The model used here is implemented in *prog.models* (Teubert, Corbetta, Kulkarni, Jarvis, & Daigle, 2022) as `prog.models.models.BatteryElectroChemEOD`.

The core model equations are: (i) Nernst’s equation for the equilibrium potential (Eq. (5)), and (ii) Butler-Volmer equation from which the surface overpotential (Eq. (6)) derives.

$$V_{U,i} = U_0 + \frac{RT}{m'F} \ln \frac{1 - \tilde{x}_i}{\tilde{x}_i} + V_{ni,i} \quad (5)$$

$$V_{\eta,i} = \frac{RT}{F\alpha} \operatorname{arcsinh} \left(\frac{J_i}{2J_{i0}} \right) \quad (6)$$

The subscript i indicates the negative or positive electrode, $i = \{n, p\}$. Term U_0 is the reference potential, R is the universal gas constant (and should not be confused with the total resistance of the model R_0), T is the electrode temperature, m' is the number of electrons transferred in the reaction ($m' = 1$ for Li-ion batteries), F is the Faraday constant, $\tilde{x}$ is the mole fraction for the Lithium-intercalated host material, and V_{ni} is the activity correction term, defined non-ideal (ni) voltage, null in ideal conditions. In Eq. (6), α is a symmetry factor, J_i is the current density and J_{i0} is the exchange current density, all defined in (Daigle & Kulkarni, 2013). The battery output voltage V can then be approximated using:

$$V = V_{U,p} - V_{U,n} - V_0 - V_{\eta,p} - V_{\eta,n},$$

where V_0 is the nominal voltage drop defined by the required current, i_{app} , multiplied by the lumped internal resistance, R_0 . Mole fractions, amount of available Li-ions, as well as concentration gradients can be retrieved from the original paper (Daigle & Kulkarni, 2013), and are omitted for the sake of brevity.

Similar to the thrown-object example, the high-fidelity model provides physics-based information about the system dynamics in the form of internal system states. We leverage this known physical information about the system by building out the snapshot matrix to include the internal state information, as well as the voltage (i.e., the observable output) and the State-of-Charge (SOC) of the battery throughout time. Thus, the snapshot matrix X is of the form

$$X = \begin{bmatrix} x_1(t_0) & x_1(t_1) & \dots \\ x_2(t_0) & x_2(t_1) & \dots \\ \dots & \dots & \dots \\ x_8(t_0) & x_8(t_1) & \dots \\ V(t_0) & V(t_1) & \dots \\ SOC(t_0) & SOC(t_1) & \dots \\ u(t_0) & u(t_1) & \dots \end{bmatrix},$$

where x_j are the eight internal states of the electrochemistry battery model, V is the voltage, SOC is the state-of-charge, and u is the current. The battery hidden state vector $\mathbf{x} \in \mathbb{R}^{8 \times 1}$ is composed of temperature T , nominal voltage drop V_0 , surface overpotential $V_{\eta,n}$, $V_{\eta,p}$, and charges for bulk $q_{b,i}$ and surface $q_{s,i}$:

$$\mathbf{x} = [T, V_0, V_{\eta,n}, V_{\eta,p}, q_{b,n}, q_{s,n}, q_{b,p}, q_{s,p}]^\top.$$

With this structure, we constructed snapshots using a collection of training data with varying loading profiles, each sequentially stacked to create one large training dataset, and solved for the linear approximation as described in Section 2.

3.3. Introduction of Artificial Noise to Aid Regression

One of the challenges related to the estimation of matrices A and B is the potential for linear dependencies among rows of X and X' . In fact, (Kutz et al., 2016) touches on the subject when presenting DMD with control (chapter 6, page 98). Linear dependencies in the input variables prevent the correct identification of A and B because the solution is not unique. Notably, we observed empirically that quasi-linearly dependent rows can also prevent DMD from finding an accurate solution, even if matrices the X and X' have full rank.

In the battery use case presented here, some hidden states show an almost-linearly dependent behavior, even if not exact. We find that with this quasi-linear dependence, DMD fails to generate a reliable model. To overcome the problem, we introduced random perturbations into X and X' , which helps the stability of DMD by “breaking” the dependency among the state variables. By adding small Gaussian noise with a standard deviation of 0.01 when building the snapshot matrices, we observed more robust matrices A and B and a more reliable DMD model.

It is important to note that the introduction of noise into the input data should be used carefully, as the standard DMD algorithm introduces a bias in the estimation process when noise is present. Loosely, DMD can be considered a simple linear regression approach where X is the independent variable and X' is the dependent variable. In linear regression, noise is supposed to be present in the dependent variable, but not in the independent one. However in our case, noise is present in both X and X' . Thus, for a more rigorous approach, one should rely on one of the more advanced DMD techniques presented in (Dawson, Hemati, Williams, & Rowley, 2016) to correct for sensor noise. However, the methods in (Dawson et al., 2016) are not tailored for DMD with control, which is necessary in our battery use-case. At this stage of the research, our attempt to use methods similar to (Dawson et al., 2016) failed at reproducing a more robust linear model. Future works will investigate the use of bias-correction methods for the case studies where B is necessary.

3.4. Prediction Accuracy

We performed a number of blind tests for a diverse set of loading profiles (three of which are illustrated in Fig. 4). In each, the initial condition of the internal states, voltage, and SOC was known, as well as the expected load throughout time. The linearized model was applied iteratively to perform a multi-step look-ahead prediction, and the results were compared to the original high-fidelity electro-chemistry model, from which the surrogate model was generated. The resulting physics-enhanced DMD model provides approximations to the high-fidelity model and a significant improvement over DMD with only observable values (Fig. 4).

Notably, the approximations are accurate for long time horizons (e.g. until SOC $\approx 20\%$ compared to an SOC $\approx 85\%$ for the standard, observable only surrogate), a new feature not achievable with a classic DMD method. However, the DMD approximation fails to capture the dynamics near End-of-Discharge (EOD) (e.g., ≥ 2000 s in Fig. 4). We find that this final stage of discharge is too nonlinear to be well approximated with our DMD model, resulting in a poor approximation of the time of final EOD. Although the sudden drop in state-of-charge and voltage are not captured, the model performs well up to a State-of-Charge (SOC) of 20%. High accuracy until $\sim 20\%$ SOC is particularly relevant because applications involving Li-ion batteries, such as UAV operations, are typically not allowed to continue below such small SOC values. Thus, using a conservative 30% minimum SOC, the DMD predictions are satisfactory.

In future work, we will explore further extensions of DMD to improve this final fit.

3.5. Prognostic Performance

One utility of a surrogate model is its ability to be used in prognostics. So, as a final exploration into the benefits of a physics-enhanced DMD, we considered the use of a DMD approximation within a prognostics framework. It was expected that DMD would provide a reduction in computation time at the cost of prediction accuracy.

To evaluate the prognostic performance, we calculate the Mean Squared Error (MSE) of predicted voltages at future time-points (a metric of interest in many prognostic applications) when compared to a ground truth. For the ground truth, we use the results from the high-fidelity model with a time step $\Delta t = 0.1$ s. This is appropriate, as the goal of a surrogate model is to adequately represent the behavior of the high-fidelity model.

To evaluate the computational performance, we measure the CPU time (using `time.process_time` in python, average of 50 runs in Table 1) per second of simulation. Standardizing per second of simulation will prevent an overestimation or underestimation of EOD from affecting the metric, allowing for a fair comparison of the computational efficiency of the methods. The results of this investigation are illustrated in Fig. 5 and Table 1.

Results illustrate that prognostics performance is significantly better for the high-fidelity model at low step sizes (Δt). At comparable step sizes, the surrogate model produces only modest or no performance improvements (e.g. compare $\Delta t=5$ for the two models in Table 1). However, the performance of the high-fidelity model decreases dramatically as the step size increases (Jarvis, Teubert, Okolo, & Kulkarni, 2022). At larger step sizes, the DMD model outperforms the high-fidelity model significantly (Fig. 5 and $\Delta t=28$ in Table 1).





Figure 4. Prediction of battery voltage (middle) and battery SOC (bottom) given loading profiles (top) using the electrochemistry model (solid lines) and a physics-enhanced DMD method (dashed lines).

Applications using prognostics generally define a prognostics performance requirement. For example, for an application with an acceptable MSE of $8 \cdot 10^{-4}$, this can only be achieved for the high-fidelity model with a Δt of 10 seconds or less. Using the surrogate model the same accuracy can be achieved with a Δt of 28 seconds. For the same accuracy, the surrogate

Table 1. Prognostic and Computational Performance at Different Step Sizes

Model	Δt (s)	CPU-time per sim second (mean $\pm$ SD, sim to EOD)	Voltage MSE (volts, sim to EOD)	Voltage MSE (volts, sim to 20% SOC)
High-fidelity	$\Delta t = 5$	$(9.53 \pm 0.37) \cdot 10^{-5}$	$4.10 \cdot 10^{-5}$	$4.55 \cdot 10^{-5}$
High-fidelity	$\Delta t = 10$	$(4.83 \pm 0.23) \cdot 10^{-5}$	$7.60 \cdot 10^{-4}$	$8.44 \cdot 10^{-4}$
High-fidelity	$\Delta t = 28$	$(2.08 \pm 0.19) \cdot 10^{-5}$	18.47	1.94
Physics-enhanced DMD	$\Delta t = 5$	$(6.35 \pm 1.96) \cdot 10^{-5}$	$1.04 \cdot 10^{-3}$	$8.92 \cdot 10^{-4}$
Physics-enhanced DMD	$\Delta t = 10$	$(3.26 \pm 2.03) \cdot 10^{-5}$	$9.80 \cdot 10^{-4}$	$7.99 \cdot 10^{-4}$
Physics-enhanced DMD	$\Delta t = 28$	$(1.09 \pm 0.09) \cdot 10^{-5}$	$7.26 \cdot 10^{-4}$	$5.90 \cdot 10^{-4}$

Figure 5. Simulation Benchmarking Results: A) Voltage Curves at different Δt , and B) MSE vs Time Step

model would perform the simulation in almost 1/5 of the time required by the equivalent higher-fidelity model.

The identification of proper Δt is very application-dependent. Some applications have rapidly varying inputs or require a precise end-of-discharge estimate, and may not be able to use a Δt above 10 seconds, in which case the surrogate model would only offer small performance improvement. For other applications, a Δt above 28 seconds is acceptable.

4. DISCUSSION AND FUTURE WORK

This work explored the use of DMD as a surrogate modeling tool for long-term predictions in the context of battery

discharge prognostics. We made use of the extended DMD approach to lift the state vector to a higher dimension in the search for a finite approximation of the Koopman operator. Instead of using basis functions or other traditional machine learning tools, we leveraged the hidden states of the higher-fidelity Li-ion battery model to construct the snapshot matrices X , X' , thus including state variables inherently linked to the physics of the system.

The introduction of physics-related quantities enabled the creation of a successful DMD-based surrogate model. While we observed some accuracy loss when predicting voltage discharge (as expected with a surrogate model), the physics-

enhanced DMD method was able to make accurate predictions, even for long time horizons. Other techniques such as lagged time-series (Tirunagari et al., 2017) and lifting through radial basis functions (Korda & Mezić, 2018) did not produce a robust model. While these methods and others can still approximate voltage and SOC appropriately for few time steps (i.e. small time horizon), they failed at predicting voltage discharge for many time steps ahead, as required in prognostics.

Our error and computational cost analysis (Section 3.5) shows performance in a prognostic setting, and highlights the advantage of using the DMD-based linear model for multi-step predictions. The time step size (Δt) used to propagate the linear dynamics can be increased drastically for physics-enhanced DMD when compared against the requirement of the full-order model, without reduction in performance. In specific applications, this can provide significant improvement in computational efficiency.

In the context of prognostics, the linearization of nonlinear dynamical systems also opens the door to faster uncertainty propagation methods, which are crucial for system diagnostics, remaining-useful-life predictions, and more in general PHM. Uncertainty in initial conditions, model parameters, and equation error (process noise) can be propagated more easily in linear systems, while full-order models often require sampling-based methods (e.g., Monte Carlo) for a characterization of future state uncertainty, which can be computationally costly.

Overall, a physics-enhanced DMD method shows potential as a surrogate model that can provide accurate long-term predictions and feasible real-time computation times in a prognostics setting.

4.1. Limitations

The scope of this work is limited to substituting a fast surrogate model in place of a higher-fidelity model. By using internal state information from the full order model, we preclude the use of field data, from which internal states of the system cannot be measured, except in rare cases. DMD methods using only observables and/or data-driven extensions do not suffer from the same limitation, and they could be applied to experimental data as well.

The selection of the lifting variables, either as hidden states from the full order dynamics or other physics-related quantities, remains an open problem. There is no guarantee that using the full state vector, as we did in the prognostic of battery voltage and SOC, will produce an acceptable reduced order model. At this stage of the research, we observed, empirically, that introducing all hidden states allowed us to build a battery reduced order model robust to different initial conditions and current input profiles. However, we cannot point to a rigorous approach to infer which variables should be in-

cluded in X and X' , and determining such a method is left to future work.

4.2. Implementation in prog_models

We have implemented a generic version of this physics-enhanced DMD framework within the open-source *prog_models* python package (Teubert, Corbetta, Kulkarni, Jarvis, & Daigle, 2022), allowing the surrogate method to be applied to any model created using that package. This is done using the `generate_surrogate` method. For example:

```
1 from prog_models import models
2
3 # Model using default parameters
4 m = models.BatteryElectroChemEOD()
5
6 # Define future loading profiles
7 # Here, we define 2 simple profiles
8 def profile1(t, x=None):
9     if t > 500: # After 500 seconds
10         return m.InputContainer({'i': 5})
11     return m.InputContainer({'i': 3})
12
13 def profile2(t, x=None):
14     if t > 500: # After 500 seconds
15         return m.InputContainer({'i': 3})
16     return m.InputContainer({'i': 5})
17
18 # Generate surrogate model using DMD and profiles.
19 # save_freq specifies the DMD step_size
20 # dt specified the data generation resolution to
   build the surrogate
21 s = m.generate_surrogate([profile1, profile2],
   method = 'dmd', dt = 0.1, save_freq = 1)
```

Listing 1. Surrogate Generation

The resulting surrogate model can be used interchangeably for simulation or prognostics (using the *prog_algs* python package (Teubert, Corbetta, & Kulkarni, 2022)). Future versions of the software are expected to implement additional methods for generating surrogate models.

4.3. Future Work

In future work, we plan to explore further extensions to the physics-enhanced DMD method and apply it to a wider range of applications. In the context of battery degradation, we hope to improve the DMD approximation to more accurately capture behavior near end of discharge. Accurately predicting battery EOD is crucial for safety in many in-time prognostic scenarios. As one approach, we plan to explore the potential of fitting an entire voltage curve by stitching together linear approximations made from DMD matrices trained on different sections of the training data.

Additionally, it will be worthwhile to further explore the configurability of a DMD approximation by including different subsets of the known physical information. A rigorous

method for determining what physical information to include would lead to a more robust physics-enhanced DMD technique.

Finally, we intend to explore alternative surrogate model generation methods, such as those emerging from the physics-informed machine learning communities, and compare with the benefits of the current DMD approach.

REFERENCES

- Daigle, M., & Kulkarni, C. S. (2013). Electrochemistry-based battery modeling for prognostics. In *Annual conference of the phm society* (Vol. 5).
- Dawson, S., Hemati, M. S., Williams, M. O., & Rowley, C. W. (2016). Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition. *Experiments in Fluids*, 57(3), 1–19.
- Jarvis, K., Teubert, C., Okolo, W., & Kulkarni, C. (2022). Improving computational efficiency of prognostics algorithms in resource-constrained settings. In *2022 ieee aerospace conference*.
- Karthikeyan, D. K., Sikha, G., & White, R. E. (2008). Thermodynamic model development for lithium intercalation electrodes. *Journal of Power Sources*, 185(2), 1398–1407.
- Koopman, B. O. (1931). Hamiltonian systems and transformation in hilbert space. *Proceedings of the united states national academy of sciences*, 17(5), 315.
- Korda, M., & Mezić, I. (2018). Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93, 149–160.
- Kutz, J. N., Brunton, S. L., Brunton, B. W., & Proctor, J. L. (2016). *Dynamic mode decomposition: data-driven modeling of complex systems*. SIAM.
- Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656, 5–28.
- Stewart, G. W. (1993). On the early history of the singular value decomposition. *SIAM review*, 35(4), 551–566.
- Teubert, C., Corbetta, M., & Kulkarni, C. (2022, May). *Prognostics algorithm python package*. Retrieved from https://github.com/nasa/prog_algs
- Teubert, C., Corbetta, M., Kulkarni, C., Jarvis, K., & Daigle, M. (2022, May). *Prognostics models python package*. Retrieved from https://github.com/nasa/prog_models
- Tirunagari, S., Kouchaki, S., Poh, N., Bober, M., & Windridge, D. (2017). Dynamic mode decomposition for univariate time series: analysing trends & forecasting. Williams, M. O., Kevrekidis, I. G., & Rowley, C. W. (2015). A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6), 1307–1346.

BIOGRAPHIES



Katelyn Jarvis is an applied mathematician and research engineer at NASA Ames Research Center. Katy received her B.S. in Mathematics and M.S. and Ph.D. in Applied Mathematics from the University of California, Davis. At NASA, Katy is a researcher within the System-Wide-Safety and Data and Reasoning Fabric projects, as well as a member of the Diagnostics and Prognostics group. Her research interests include development of novel algorithmic approaches to prognostics, improvement of computational efficiency of these methods, and application of prognostics in human-health related fields.



Matteo Corbetta is a Research Engineer with KBR Inc. at NASA Ames Research Center, CA. He is a member of the Diagnostics & Prognostics Group at Ames. He works at the intersection of physics-based models, machine learning and physics-informed learning, and uncertainty quantification methods for diagnostics and prognostics of aerospace systems. He is a member of the Editorial Board of the PHM Society.



Stefan Schuet received the B.S. and M.S. degrees in electrical engineering from Santa Clara University, Santa Clara, CA, in 2001 and 2004, respectively. He joined NASA Ames Research Center in 2001, and currently leads research that combines applied-physics with machine learning and optimization to address challenging aerospace engineering problems. Mr. Schuet holds a U.S. patent and several distinguished awards including the 2008 NASA Government Invention of the Year Award, the 2009 R&D Magazine Top 100 Award, and the 2018 NASA Exceptional Engineering Achievement Medal.



Christopher Teubert is the Diagnostics and Prognostics group lead at NASA Ames Research Center. Christopher received his B.S. in Aerospace Engineering from Iowa State University and his M.S. in Computer Science and Engineering from Santa Clara University. At NASA, Christopher performs research in software architectures for prognostics, is the lead of NASA's Prognostics CoE, and is the lead developer for the Prognostics Python Packages (prog_models, prog_algs, and prog_server).