

Dynamic mode decomposition for particle-in-cell simulations of a Hall thruster and plume

IEPC-2024-412

*Presented at the 38th International Electric Propulsion Conference, Toulouse, France
June 23-28, 2024*

Joshua D. Eckels*, Thomas A. Marks†, Doruk Aksoy‡, Sruti Vutukury§, and Alex A. Gorodetsky¶
University of Michigan, Ann Arbor, MI 48105, United States of America

In this work, we target the acceleration of particle-in-cell (PIC) methods for Hall thrusters by way of data-driven reduced-order modeling through the dynamic mode decomposition (DMD). The primary goal of data-driven reduced-order modeling is to learn a system’s dynamics from observations of its state over time. Once obtained, the reduced-order model can be used in place of the original system for cheap forecasting of its state. In this work, we build reduced-order models for the dynamics of kinetic PIC simulations of a Hall thruster plasma. Specifically, we apply the linear DMD approximation to a fully-kinetic simulation of a Hall thruster discharge channel and separately to a 3d hybrid fluid/PIC simulation of a Hall thruster plume. We found that a rank-truncated linear fit is sufficient to qualitatively, and sometimes quantitatively, predict the primary modes of the plasma dynamics, including the azimuthal wave propagations in the discharge channel that are indicative of the electron drift instability. We also observed that high noise, high-dimensional, and multi-scale behavior presents a computational challenge for efficiently and accurately fitting PIC simulations. With further refinement, we envision that “kinetic” reduced-order models of a Hall thruster have potential use in downstream applications such as online control and closure modeling.

I. Introduction

The high efficiency and specific impulse of Hall thrusters has led to their increasing deployment for satellite station-keeping, orbit-raising, and deep-space missions.¹ As missions scale to longer durations and higher powers, there is an increasing need for modeling and simulation efforts to aid experiment for design, test, and qualification. State-of-the-art (SOTA) Hall thruster models use fluid or hybrid fluid/particle descriptions of the discharge plasma coupled to the applied electric and magnetic fields to predict key performance metrics such as thrust, erosion rates, and efficiency. SOTA models are typically tuned to match available experimental data for a given thruster and operating condition, and then function as a “soft sensor” to query new or inaccessible properties of the plasma; this approach has led to the design of magnetically-shielded thrusters,² among other innovations, and continues to play a role in design and qualification today.

Alternatively, SOTA models do not readily generalize across thrusters and operating conditions and must be tuned to fit experiments on a case-by-case basis. Demonstrating a “predictive”, generalizable Hall thruster model has been a topic of wide consideration³ with many promising benefits for Hall thruster development, but to date has been precluded by the unknown physics of anomalous electron transport.^{4,5} It is widely believed that kinetically-driven instabilities result in turbulent transport of electrons from cathode to anode,⁶ a phenomenon which is critical in the overall operation of the device but is not captured by a fluid description of the plasma. Instead, a fully-kinetic description through the solution of the Boltzmann equation for each plasma species is required to resolve the underlying physics.⁷

*PhD candidate, Department of Aerospace Engineering, eckelsjd@umich.edu.

†Postdoc, Department of Aerospace Engineering

‡PhD candidate, Department of Aerospace Engineering

§Masters student, Department of Aerospace Engineering

¶Associate Professor, Department of Aerospace Engineering, goroda@umich.edu.

The particle-in-cell (PIC) method provides a straightforward and intuitive approach for the solution of the kinetic Boltzmann equation.⁸ The distribution function of each species is approximated by a finite number of computational “macro-particles” that represent a much larger number of real physical particles. The macro-particles obey Newton’s laws and are time-advanced by the electric and magnetic fields. The particles may undergo collisional processes with surfaces or other species. The charge and current density carried by the particles are then deposited onto a finite-element mesh, where Maxwell’s equations are applied to update the electric and magnetic fields (typically the electrostatic assumption is made in Hall thrusters to update an electrostatic potential through Poisson’s equation). The fields are then interpolated to the particle positions and the cycle repeats.

In principle, a 3d full-PIC simulation of a thruster with sufficient spatial/temporal resolution of the turbulent electron transport could serve as a “numerical experiment” for studying complex thruster behavior (akin to direct numerical simulation in the CFD community). Indeed, such attempts at increasing 3d resolution have been recently performed.^{9–11} However, the memory and computational costs for increasing resolution and device scale quickly exceed modern computing capabilities. Furthermore, such large-scale simulations are not practical for many-query engineering tasks such as design, optimization, and uncertainty quantification. It is therefore desirable to distill the pertinent kinetic information available from PIC into practicable closures for SOTA fluid Hall thruster models.

To this end, we target the acceleration of existing PIC methods by way of building data-driven reduced-order models (ROMs). The general principle behind a ROM is that the dynamics of high-dimensional physical systems (e.g. the evolution of field quantities in a PIC simulation) can be represented in a low-dimensional space. A ROM can furthermore be learned non-intrusively directly from numerical or experimental data of the system, which is especially pertinent for PIC simulations where the field quantities evolve from complex interactions with macro-particles and have unknown governing equations. Similar strategies have been applied to PIC for other systems^{12–14} and more recently to similar Hall thruster configurations;^{15,16} we extend these works to new Hall thruster test cases and conditions. Specifically, we apply the dynamic mode decomposition¹⁷ to predict the time-evolution of field quantities in a fully-kinetic PIC simulation of the Hall thruster discharge channel, and of a separate hybrid fluid/PIC simulation of a Hall thruster plume. We compare the results for varying rank truncations of the dynamic mode decomposition.

Our contributions include the analysis of two verification PIC simulations of a Hall thruster plasma and a preliminary test of data-driven reduced-order modeling for Hall thruster PIC data. In Section II, we outline the PIC simulations and the ROM methods. In Section III, we summarize the performance of the ROMs and compare across varying rank truncations. Finally, we provide concluding remarks and provide paths for continuing work. We envision that the acceleration of PIC methods by reduced-order modeling can inform and improve kinetic closures in SOTA fluid Hall thruster models to ultimately enable their wider usage as a predictive engineering tool.

II. Methods

In this section, we outline the two PIC simulations under consideration, namely a full-PIC study of a Hall thruster discharge channel and a hybrid fluid/PIC study of a Hall thruster plume. Then, we describe the basis of the dynamic mode decomposition and the specific methods we apply for reduced-order modeling of the PIC simulations. The performance of the ROMs are presented and discussed in Section III.

A. Hall thruster PIC simulations

We first study a fully-kinetic (i.e. particle description of all species) simulation of the discharge channel using the open-source Warp-X code¹⁸ on the axial-azimuthal benchmark from Charoy et al.¹⁹ Then, we study a hybrid simulation of a Hall thruster plume based on that of Boyd et al.,²⁰ but modified to make use of the 3d Air Force Research Laboratory in-house Thermophysics Universal Research Framework (TURF) code.²¹ The simulation settings are summarized in the following subsections.

1. Fully-kinetic discharge channel (Warp-X¹⁸)

The highly-scalable, open-source Warp-X PIC code was adapted in a parallel study²² for fully-kinetic simulation of the Hall thruster discharge channel and was benchmarked against the 2d axial-azimuthal case of Charoy et al.¹⁹ This benchmark was developed to study the growth of the azimuthal electron drift instability,

which plays a key role in turbulent axial electron transport. Critically, the benchmark treats both electrons and ions with a particle description; this is a necessary approach for resolving the turbulent transport of interest to fluid closure models, which makes this a useful test case for our application. The computational domain is a structured grid of 512×256 cells of total length $2.5 \text{ cm} \times 1.28 \text{ cm}$ in the axial-azimuthal directions, respectively, as shown in Figure 1. We study three field quantities on the computational mesh, namely the ion density, and the axial and azimuthal current densities. The simulation runs with a time step of $\Delta t = 5 \text{ ps}$ for a total duration of $t_f = 20 \text{ } \mu\text{s}$. For our analysis, we only use data after transient startup ($t = 10 \dots 20 \text{ } \mu\text{s}$).

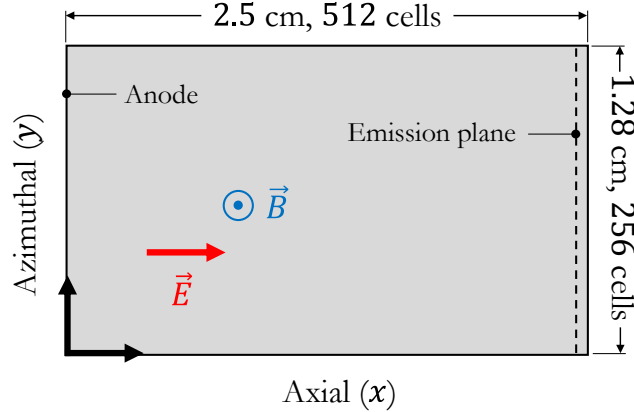


Figure 1: Computational domain of the 2d axial-azimuthal PIC benchmark of a Hall thruster discharge channel from Charoy et al.¹⁹

Ultimately, saving once every 2000 time steps, this results in a $(512, 256, 3, 1000)$ array of PIC data for 1000 total snapshots of the three field quantities (ion density and axial/azimuthal current densities) on the structured mesh grid. We leave the rest of the simulation details in the original work.¹⁹

2. Hybrid fluid/particle Hall thruster plume (TURF²¹)

TURF²¹ is a modular computational framework enabling the study of thruster, plume, and spacecraft interaction. In a parallel study,²³ we adapt TURF to compare with the simulation results of Boyd et al²⁰ in predicting plume properties of an in-flight SPT-100.²⁴ The model treats electrons as an isothermal, unmagnetized, collisionless fluid and heavier species (ions and neutrals) as particles. The plasma potential is obtained by assuming quasineutrality. Typical beam properties of the SPT-100 are used as injection conditions at the computational origin. The neutral atoms are injected with a number density of 10^{18} m^{-3} , face-normal velocity of 300 m/s , and a temperature of 1000 K , and the ions with a density of 10^{17} m^{-3} , a normal velocity of 17000 m/s , and a temperature of 4 eV . We assume both species have an effective divergence angle of 15° and are injected with radially-uniform properties. The Boltzmann relation is used with an electron temperature of 3 eV , fixed over the computational domain. Normal electric fields at the outflow boundaries are set to zero and particles exiting the domain are removed from the simulation.

The domain consists of a $120 \times 100 \times 100$ structured 3d grid over a $12 \times 10 \times 10 \text{ m}^{-3}$ volume representing one quadrant of an exhaust plume, as shown in Figure 2. We study two field quantities on the computational mesh, namely the ion and neutral densities.

The thruster is oriented along the positive x -axis of an (x, y, z) Cartesian coordinate system, with the range $x \in (-2, 10)$ allowing backward scattering of ions in the plume due to collisions. Charge-exchange collisions are treated with the direct simulation Monte Carlo method (as in the original work) and elastic collisions are treated with the Monte Carlo collision method provided in TURF. The simulation runs for $t_f = 62.5 \text{ ms}$ with a time step of $\Delta t = 5 \text{ } \mu\text{s}$. Removing startup transients before $t = 2.5 \text{ ms}$ and saving once every 50 time steps, we result in a $(120, 100, 100, 2, 240)$ array of PIC data for 240 total snapshots of the two field quantities (ion and neutral density) on the 3d structured grid.

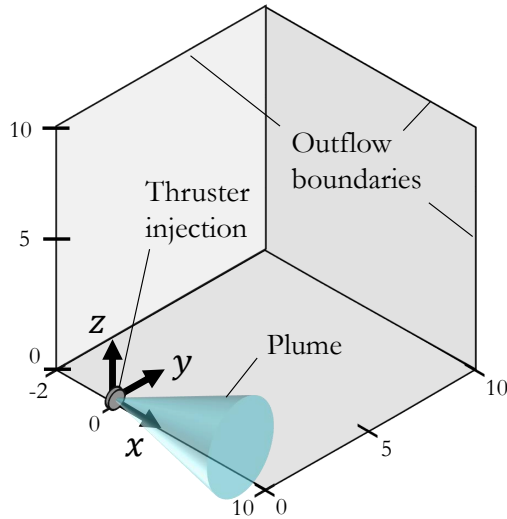


Figure 2: Computational domain of the 3d hybrid fluid/PIC simulation of a Hall thruster plume using the TURF code.²¹

B. Dynamic mode decomposition

In this section, we outline the basis of dynamic mode decomposition (DMD) for reduced-order modeling and describe the specific methods we apply to the PIC simulations. The time-evolution of a dynamical system with n states $\mathbf{x}(t) \in \mathbb{R}^n$ can be represented with the initial value problem:

$$\dot{\mathbf{x}} = f(\mathbf{x}), \quad \mathbf{x}_0 \triangleq \mathbf{x}(t=0), \quad (1)$$

which is solved by standard time-stepping numerical tools to obtain the system state at some future time $\mathbf{x}(t = t_f)$ for the known dynamics f and initial condition \mathbf{x}_0 . The guiding principle of DMD is to approximate the dynamics of \mathbf{x} as linear via:

$$\dot{\mathbf{x}} \approx A\mathbf{x}, \quad \text{for which } \mathbf{x}(t) = e^{At}\mathbf{x}_0. \quad (2)$$

This linear assumption permits the analytical solution $\mathbf{x}(t)$ given in Eq. (2) for any future time t . The main task of DMD then is to solve for the best-fit matrix A from numerical or experimental data. Given the states $\mathbf{x}_1 \dots \mathbf{x}_m$ observed at $k = 1 \dots m$ distinct time steps, and given the discrete-time form of Eq. (2) via $\mathbf{x}_{k+1} = A_d\mathbf{x}_k$ with $\mathbf{x}_k \triangleq \mathbf{x}(t = k\Delta t)$, the discrete-time matrix $A_d = e^{A\Delta t}$ can be obtained through the least-squares fit:

$$X = \begin{bmatrix} | & & | \\ \mathbf{x}_1 & \dots & \mathbf{x}_{m-1} \\ | & & | \end{bmatrix}, \quad X' = \begin{bmatrix} | & & | \\ \mathbf{x}_2 & \dots & \mathbf{x}_m \\ | & & | \end{bmatrix}, \quad (3)$$

$$X' \approx A_d X, \quad \text{with} \quad (4)$$

$$A_d = \arg \min_B \|X' - BX\|_2 = X'X^\dagger. \quad (5)$$

The primary challenge of this approach applied to engineering simulations (e.g. PIC) is the storage and computational requirements of $A_d \in \mathbb{R}^{n \times n}$, typically with $n \gg m$ due to high spatial mesh resolution in the computational domain. As such, we explore several truncation levels of the standard exact-DMD.²⁵ The exact-DMD approach uses a rank- r truncation of the SVD modes of the data matrix via:

$$X = U\Sigma V^* = \begin{bmatrix} | & & | \\ u_1 & \dots & u_n \\ | & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_{m-1} \end{bmatrix} \begin{bmatrix} - & v_1 & - \\ & \vdots & \\ - & v_{m-1} & - \end{bmatrix} \quad (6)$$

$$\approx U_r \Sigma_r V_r^* = \sum_{i=1}^r \sigma_i u_i v_i^*, \quad (7)$$

where the singular values σ_i are arranged in descending order with corresponding left- and right-singular vectors u_i and v_i , respectively. The discrete-time dynamics A_d are then projected onto the reduced basis U_r with the similarity transform $\tilde{A}_d = U_r^* A_d U_r$. An eigen-decomposition in the reduced space of \tilde{A}_d can be related²⁵ to a reduced set of eigenvectors ϕ_i and eigenvalues λ_i , $i = 1 \dots r$ of the full discrete-time dynamics A_d . Finally, a diagonalization of the continuous-time dynamics $A = \Phi \Omega \Phi^{-1}$ permits the modal-decomposition form of the solution in Eq. (2) via:

$$\mathbf{x}(t) = \Phi e^{\Omega t} \Phi^{-1} \mathbf{x}_0 \approx \sum_{i=1}^r \phi_i e^{\omega_i t} b_i, \quad (8)$$

where $\Omega = \text{diag}(\omega_i)$, $\omega_i = \ln \lambda_i / \Delta t$ resulting from $A = \ln A_d / \Delta t$, and b_i is an amplitude from the i -th component of $\Phi^{-1} \mathbf{x}_0$. As given by Eq. (8), the approximation of $\mathbf{x}(t)$ ultimately results from the rank- r truncation of the full eigenbasis $\Phi = [\phi_1 \dots \phi_n]$, with higher accuracy for larger ranks. In this work, we compare three truncation ranks $r = (10, 50, 100)$ on the accuracy of DMD applied to the PIC datasets. We note that log normalization was applied for ion and neutral density field quantities (due to orders of magnitude ranges in their raw data values) and (0, 1) min-max scaling for all other field quantities.

III. Results

1. Fully-kinetic discharge channel (Warp-X¹⁸)

We first test the performance of exact-DMD in predicting the steady-state dynamics of the fully-kinetic Warp-X PIC simulation of the Hall thruster discharge channel. The state vector \mathbf{x} in Eq. (1) is composed in this case of the field quantities on the spatial PIC grid. We use the first 50% of time snapshots to train the DMD fit and then compare future-time DMD predictions to the ground truth PIC simulation. Figure 3 compares the DMD prediction of ion density at the final prediction time step to the ground truth for truncation ranks of $r = 10, 50$, and 100. We additionally report the absolute error $|\mathbf{x} - \hat{\mathbf{x}}|$ between the ground truth \mathbf{x} and the DMD prediction $\hat{\mathbf{x}}$.

The primary steady-state modes of ion density in this test case include the strong ionization in the upstream axial region and the azimuthally propagating waves that are indicative of the electron drift instability. We observe in Figure 3 that a rank-10 fit is sufficient to capture this phenomena, with higher rank approximations not significantly improving the fit further. Regions of high absolute error indicate that the DMD fit may be out of phase with the ground truth.

Figure 4 summarizes the relative L_2 error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ over time, with the gray shaded region indicating the training period. Note that $t = 0$ corresponds to the start of steady-state operation, with a final time of $t_f = 10 \mu\text{s}$. We observe in Figure 4 that using higher rank approximations results in improved training accuracy, but does not lead to significant improvement in future-time prediction. This may indicate again that the rank-10 approximation is sufficient to capture the majority of the steady-state dynamics, with a relative error of around 7% at the end of the prediction interval. We note that a rank- r SVD truncation of a data matrix $X \in \mathbb{R}^{n \times m}$ for n states and m time steps provides a compression ratio of $nm / (r(n + m + 1))$, which is ≈ 100 for the rank-10 approximation of the Warp-X dataset presented here.

In addition to ion density, we also study the axial and azimuthal current densities in the Warp-X PIC simulation. Figure 5 compares the axial profiles of the two current densities for the same three rank truncations of $r = 10, 50$, and 100. Due to inherent small-scale fluctuations and statistical noise, we average the current density profiles over the time-prediction window $t = 5 \dots 10 \mu\text{s}$ and spatially over the azimuthal direction. The relative errors ($\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$) for each case are provided in Table 1.

We observe in Figure 5 that exact-DMD is able to capture the average profile of axial current density in the discharge channel, as well as the strong azimuthal Hall current indicated by the large negative spike in azimuthal current density in the upstream discharge channel. We note the spikes in current density at the end of the domain are due to plasma sheath boundary conditions in the original PIC simulation. Overall, we observe that exact-DMD captures the steady-state profiles of axial and azimuthal current densities well for a rank-truncation of $r = 10$, with higher ranks not significantly improving the fits.

2. Hybrid fluid/particle Hall thruster plume (TURF²¹)

In this section, we report the performance of exact-DMD in predicting the dynamics of the hybrid fluid/particle TURF simulation of a Hall thruster plume. We again use the first 50% of time snapshots to train the DMD




Figure 3: Comparison of exact-DMD ion density predictions to the ground truth Warp-X simulation at the final prediction time for truncation ranks of $r = 10, 50,$ and 100 . The rank-10 approximation captures the dominant steady-state modes of ion density. Higher ranks do not significantly improve the fit.

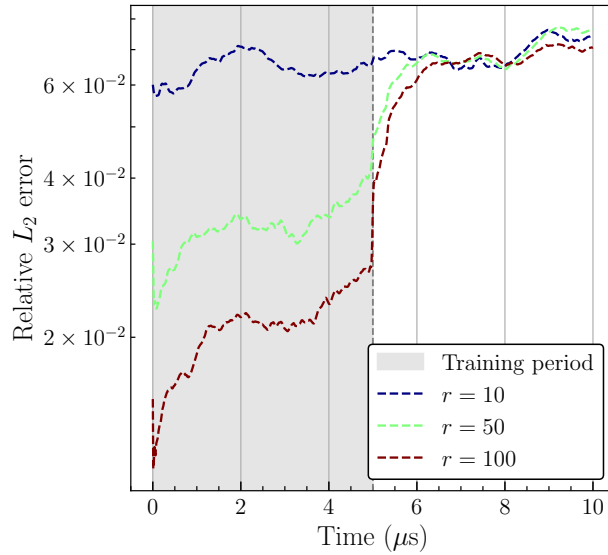


Figure 4: Exact-DMD relative errors in ion density over time for the Warp-X PIC simulation for truncation ranks of $r = 10, 50,$ and 100 . The shaded gray region indicates the training period. Higher rank approximations improve the fit during training, but do not significantly improve future-time prediction accuracy.

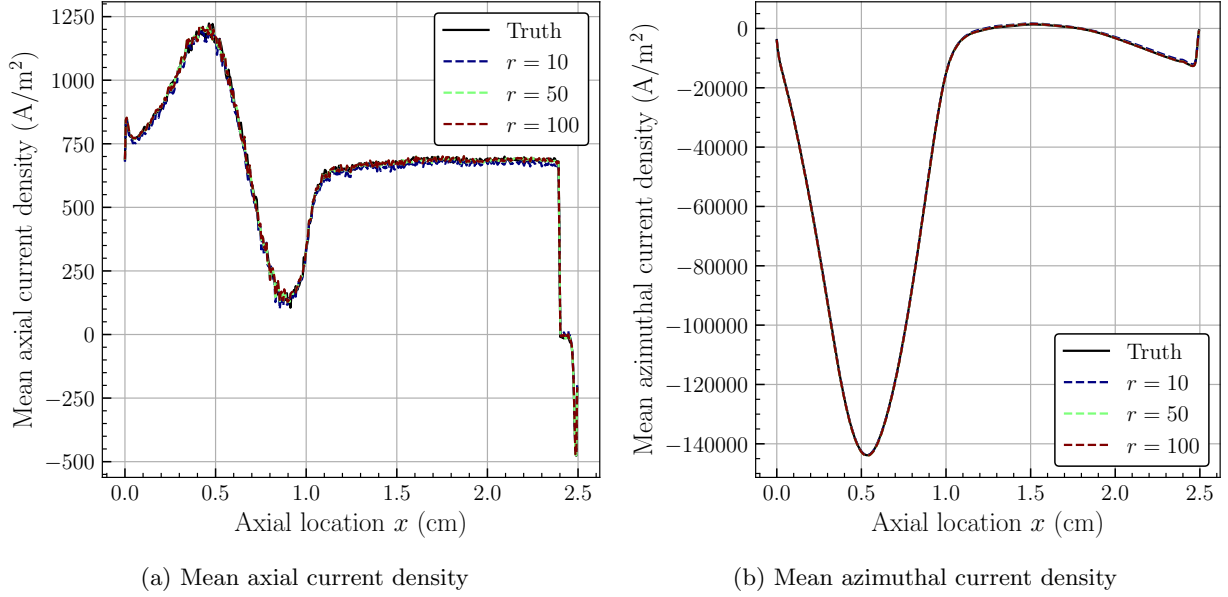


Figure 5: Comparison of exact-DMD current density predictions to the ground truth Warp-X simulation for truncation ranks of $r = 10, 50, \text{ and } 100$. The current densities are averaged over the azimuthal direction and over the time-prediction window $t = 5 \dots 10 \mu\text{s}$. The ground truth steady-state profiles are well-captured by the DMD predictions.

Table 1: Exact-DMD relative errors of axial and azimuthal current densities for varying rank truncations.

Rank	Relative error	
	Axial current density	Azimuthal current density
10	0.0282	0.0050
50	0.0205	0.0033
100	0.0250	0.0031

fit and compare prediction accuracy to the ground truth PIC simulation for future time steps. Figure 6 compares the exact-DMD predictions of plume ion density at the final prediction time step for truncation ranks of $r = 10, 50, \text{ and } 100$, with the absolute error $|\mathbf{x} - \hat{\mathbf{x}}|$ between the ground truth \mathbf{x} and DMD prediction $\hat{\mathbf{x}}$. Figure 7 makes a similar comparison at the final time step for the prediction of neutral density in the plume. For simplicity, we only show an (x, y) cross-section at $z = 0$ of the full 3d simulation results here (see Figure 2), but we include 3d comparisons in Appendix A.

The plume ion density is characterized by the main beam expanding along the positive x -axis and a scattered beam that expands radially away from the main beam due to collisions with slow-moving neutrals. The neutrals diffuse radially away from the injection source at $x = 0$ with a divergence angle of less than 90 deg (i.e. the sharp cutoff region in Figure 7). In contrast to the Warp-X simulation, a rank-10 approximation is not sufficient to capture the steady-state dynamics in the TURF simulation, as observed by the absence of the main beam in the DMD predictions in Figures 6 and 7. As higher ranks are included, the primary characteristics of the main and scattered beams are qualitatively predicted, while the absolute error is still high.

The relative error of the DMD ion and neutral density predictions over time are summarized in Figure 8. The 50% training period is indicated by the shaded gray region. We again compare the DMD fit for




Figure 6: Comparison of exact-DMD ion density predictions to the ground truth TURF simulation at the final prediction time for truncation ranks of $r = 10, 50,$ and 100 . Predictions are shown as (x, y) cross-sections at $z = 0$ of the full 3d domain (see Figure 2). The DMD prediction improves for higher rank approximations, with the main plume characteristics not captured until $r = 100$.




Figure 7: Comparison of exact-DMD neutral density predictions to the ground truth TURF simulation at the final prediction time for truncation ranks of $r = 10, 50$, and 100 . Predictions are shown as (x, y) cross-sections at $z = 0$ of the full 3d domain (see Figure 2). The DMD prediction improves for higher rank approximations, with the main plume characteristics not captured until $r = 100$.

rank truncations of $r = 10, 50$, and 100 . We observe in Figure 8, similar to the Warp-X results, that higher rank approximations serve to improve training accuracy but do not generalize to future-time prediction. However, even though the rank-100 relative error is still increasing at the end of the prediction window, it produces qualitatively better results. This may indicate that higher ranks or more training data are required for improved long-time prediction accuracy.

Overall, we obtained a qualitatively good fit to the TURF plume dynamics using a rank-100 DMD approximation, while the magnitude of error in future-time prediction is still high. This result illustrates some of the challenges of reduced-order modeling for PIC simulations. In this case, higher ranks and more training data are likely required simply because the 3d problem has more states (i.e. mesh points) to capture than a similar 2d simulation (roughly 6x more between the TURF and Warp-X cases in this study). As more data and higher ranks are included, the computational and memory requirements can make the solution infeasible. This problem for larger datasets may be mitigated by using incremental techniques.²⁶

Another challenge with building reduced-order models for PIC simulations is inherent noise in the field quantities. We observed high noise, for example, in the current densities from the Warp-X simulation – these quantities may be predicted in an average sense as we demonstrated, but it is difficult for exact-DMD to capture the dynamics of this behavior. We also observed large ranges in ion and current density values in the TURF simulation, often spanning several orders of magnitude between the main and scattered plume regions. Along with appropriate normalization, higher ranks are likely required to capture the dynamics of both regions adequately.

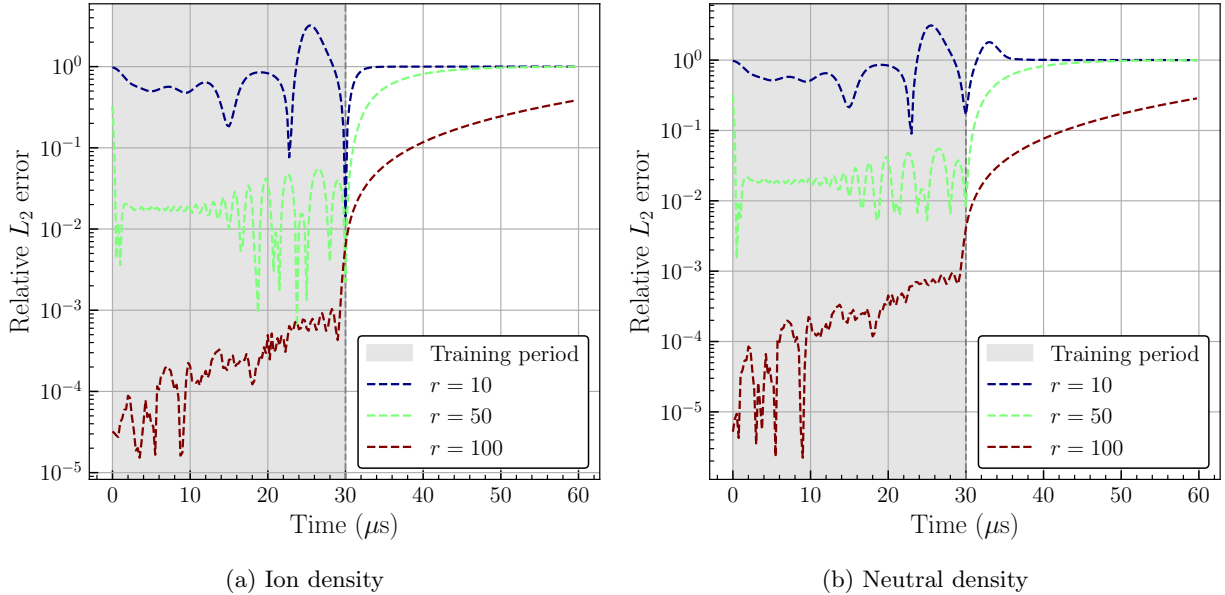


Figure 8: Exact-DMD relative errors in ion and neutral densities over time for the TURF PIC simulation using rank truncations of $r = 10, 50$, and 100 . The shaded gray region indicates the training period. Higher rank approximations improve training accuracy but do not generalize to improved future-time prediction accuracy.

IV. Conclusion

In this work, we considered the data-driven, non-intrusive reduced-order modeling of Hall thruster PIC simulations through the dynamic mode decomposition (DMD). We first applied DMD for future-time prediction of the fully-kinetic dynamics in an axial-azimuthal Hall thruster discharge channel simulation – we found that a low-rank approximation was sufficient to capture the primary modes of the plasma, including the azimuthal waves characteristic of the electron drift instability. Then, we applied DMD to predict the dynamics of a 3d hybrid fluid/PIC simulation of a Hall thruster plume – while we obtained a qualitative fit with the data, we found the high-dimensionality and multi-scale nature of the 3d simulation difficult to capture with DMD, with higher ranks and more training data likely required to obtain a better fit, along with the potential use of iterative methods to mitigate increased memory and computational requirements.

Overall, we demonstrated that acceleration of Hall thruster PIC simulations is achievable through data-driven reduced-order modeling. In this study, a simple rank-truncated linear fit of the plasma dynamics was sufficient to obtain qualitative, and in some regions even quantitative, agreement with the ground truth PIC simulations. Furthermore, the primary modes of interest were captured by the linear fit, namely the azimuthal wave propagation of interest in studying anomalous electron transport. We envision that the ability to cheaply run “kinetic” reduced-order models may have potential application in areas such as online control, parameter/state estimation, and even in use as a direct information source for fluid closure models. Future work includes refinement of the current linear methods, analysis of more expressive nonlinear methods, expansion to more parameter and operating regimes, and integration into downstream applications such as online control and closure modeling.

Appendix

A. 3d TURF comparison plots

In this section, we supplement the results of Section III.2 with 3d comparison plots between the ground truth TURF PIC simulation and the DMD prediction. We show the 3d profiles of ion and neutral densities in the plume in Figures 9 and 10, respectively, for the highest accuracy rank-100 truncation only. The field quantities are shown as three perpendicular cross-sections of the primary (x, y, z) coordinate planes at the

final prediction time step.

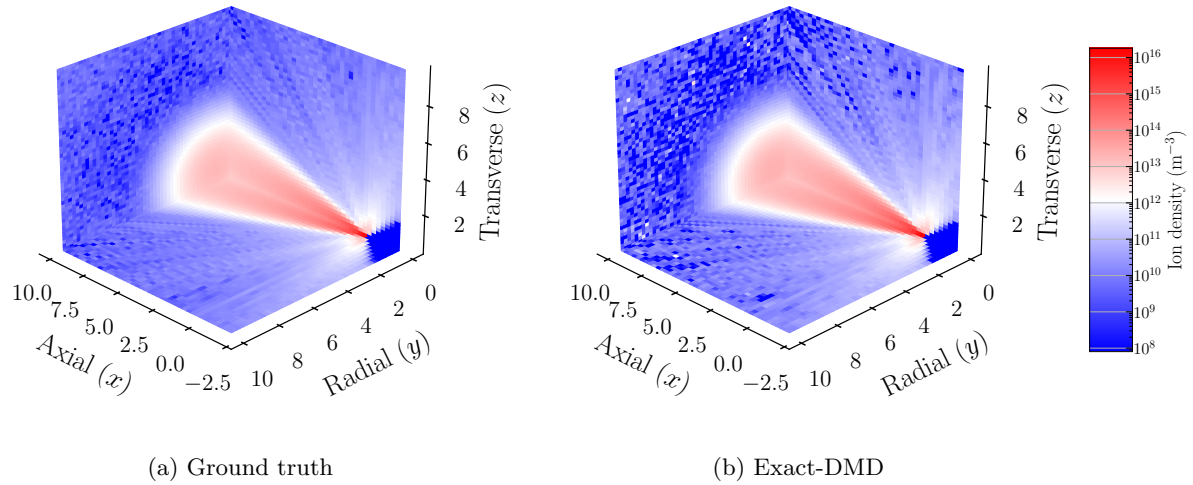


Figure 9: Comparison of exact-DMD ($r = 100$) ion density prediction to the ground truth TURF simulation at the final time step. Ion density is shown on three perpendicular cross-sections of the primary (x, y, z) coordinate planes. Exact-DMD qualitatively predicts the full 3d profile of the main and scattered beams in the computational domain.

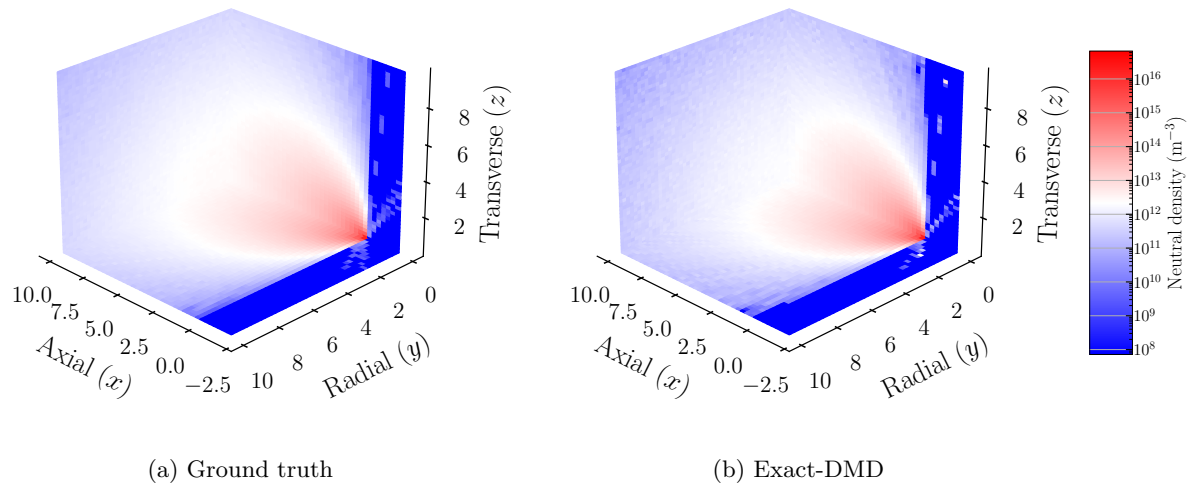


Figure 10: Comparison of exact-DMD ($r = 100$) neutral density prediction to the ground truth TURF simulation at the final time step. Neutral density is shown on three perpendicular cross-sections of the primary (x, y, z) coordinate planes. Exact-DMD qualitatively predicts the full 3d profile of the neutral diffusion from the thruster injection source.

We observe qualitative agreement between the exact-DMD prediction and the ground truth PIC simulation of the 3d Hall thruster plume. The main and scattered beam profiles are captured in Figure 9 and the radial neutral diffusion is captured in Figure 10. Overall, the simulation displays axisymmetry around the x -axis which is well-captured by the DMD approximation. It is evident that the majority of the compu-

tational domain is filled with low density or vacuum conditions as opposed to the high density of the main beam plume – we observe noisier DMD predictions in these regions due to lack of any interesting dynamic behavior. We note this limitation especially when the object of the simulation is to accurately predict ion energy in these far-field regions for studying spacecraft interaction and erosion processes.

Acknowledgements

This research was supported in part through computational resources provided by Advanced Research Computing at the University of Michigan and additionally by an AFOSR DURIP under Program Manager Dr. Fariba Fahroo, grant number FA9550-23-1-006. Funding for this work was provided by NASA in part through the Joint Advanced Propulsion Institute (JANUS) under grant number 80NSSC21K1118, as well as in part through a NASA Space Technology Graduate Research Opportunity grant 80NSSC23K1181. Additionally, we would like to acknowledge support by Los Alamos National Laboratory under the project “Algorithm/Software/Hardware Co-design for High Energy Density applications.”

References

- ¹Jason D. Frieman, Hani Kamhawi, Jon Mackey, Peter Y. Peterson, James H. Gilland, Richard R. Hofer, Derek Inaba, Hoang Dao, Nicholas A. Branch, and Benjamin Welander. Extended Wear Testing of the 12-kW Advanced Electric Propulsion System Engineering Test Unit Hall Thruster. In *37th International Electric Propulsion Conference*, Boston, MA USA, June 2022.
- ²Ioannis G. Mikellides and Ira Katz. Numerical simulations of Hall-effect plasma accelerators on a magnetic-field-aligned mesh. *Physical Review E*, 86(4):046703, October 2012.
- ³Jean-Pierre Boeuf. Tutorial: Physics and modeling of Hall thrusters. *Journal of Applied Physics*, 121(1):011101, January 2017.
- ⁴Thomas A. Marks and Benjamin A. Jorns. Challenges with the self-consistent implementation of closure models for anomalous electron transport in fluid simulations of Hall thrusters. *Plasma Sources Science and Technology*, 32(4):045016, April 2023.
- ⁵Thomas A. Marks and Benjamin A. Jorns. Evaluation of algebraic models of anomalous transport in a multi-fluid Hall thruster code. *Journal of Applied Physics*, 134(15):153301, October 2023.
- ⁶Igor D. Kaganovich, Andrei Smolyakov, Yevgeny Raitsev, Eduardo Ahedo, Ioannis G. Mikellides, Benjamin Jorns, Francesco Taccogna, Renaud Guerout, Sedina Tsikata, Anne Bourdon, Jean-Pierre Boeuf, Michael Keidar, Andrew Tasman Powis, Mario Merino, Mark Cappelli, Kentaro Hara, Johan A. Carlsson, Nathaniel J. Fisch, Pascal Chabert, Irina Schweigert, Trevor Laffeur, Konstantin Matyash, Alexander V. Khrabrov, Rod W. Boswell, and Amnon Fruchtman. Physics of $E \times B$ discharges relevant to plasma propulsion and similar technologies. *Physics of Plasmas*, 27(12):120601, section 5, December 2020.
- ⁷Richard Michael Jack Kramer, Eric C. Cyr, Sean T. Miller, Edward Geoffrey Phillips, Gregg Arthur Radtke, Allen C. Robinson, and John N. Shadid. A Plasma Modeling Hierarchy and Verification Approach. Technical Report SAND-2020-3576, Sandia National Lab, Albuquerque, NM USA, March 2020.
- ⁸F. Taccogna, F. Cichocki, D. Eremin, G. Fubiani, and L. Garrigues. Plasma propulsion modeling with particle-based algorithms. *Journal of Applied Physics*, 134(15):150901, October 2023.
- ⁹Francesco Taccogna and Pierpaolo Minelli. Three-dimensional particle-in-cell model of Hall thruster: The discharge channel. *Physics of Plasmas*, 25(6):061208, June 2018.
- ¹⁰Willca Villafana, G. Fubiani, L. Garrigues, G. Vigot, B. Cuenot, and O. Vermorel. 3D Particle-In-Cell modeling of anomalous electron transport driven by the Electron Drift Instability in Hall thrusters. In *37th International Electric Propulsion Conference*, June 2022.
- ¹¹W. Villafana, B. Cuenot, and O. Vermorel. 3D particle-in-cell study of the electron drift instability in a Hall Thruster using unstructured grids. *Physics of Plasmas*, 30(3):033503, March 2023.
- ¹²Julio L. Nicolini, Dong-Yeop Na, and Fernando L. Teixeira. Model Order Reduction of Electromagnetic Particle-in-Cell Kinetic Plasma Simulations via Proper Orthogonal Decomposition. *IEEE Transactions on Plasma Science*, 47(12):5239–5250, December 2019.
- ¹³Indranil Nayak, Fernando L. Teixeira, Dong-Yeop Na, Mrinal Kumar, and Yuri A. Omelchenko. Accelerating Particle-in-Cell Kinetic Plasma Simulations via Reduced-Order Modeling of Space-Charge Dynamics using Dynamic Mode Decomposition. (arXiv:2303.16286), May 2024.
- ¹⁴Maryam Reza, Farbod Faraji, and Aaron Knoll. Resolving multi-dimensional plasma phenomena in Hall thrusters using the reduced-order particle-in-cell scheme. *Journal of Electric Propulsion*, 1(1):19, October 2022.
- ¹⁵F. Faraji, M. Reza, A. Knoll, and J. N. Kutz. Dynamic mode decomposition for data-driven analysis and reduced-order modeling of $E \times B$ plasmas: I. Extraction of spatiotemporally coherent patterns. *Journal of Physics D: Applied Physics*, 57(6):065201, November 2023.
- ¹⁶Farbod Faraji, Maryam Reza, Aaron Knoll, and J. Nathan Kutz. Dynamic Mode Decomposition for data-driven analysis and reduced-order modelling of ExB plasmas: II. dynamics forecasting. (arXiv:2308.13727), August 2023.
- ¹⁷Peter J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, August 2010.

¹⁸J. L. Vay, A. Almgren, J. Bell, L. Ge, D. P. Grote, M. Hogan, O. Kononenko, R. Lehe, A. Myers, C. Ng, J. Park, R. Ryne, O. Shapoval, M. Thévenet, and W. Zhang. Warp-X: A new exascale computing platform for beam-plasma simulations. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 909:476–479, November 2018.

¹⁹T. Charoy, J. P. Boeuf, A. Bourdon, J. A. Carlsson, P. Chabert, B. Cuenot, D. Eremin, L. Garrigues, K. Hara, I. D. Kaganovich, A. T. Powis, A. Smolyakov, D. Sydorenko, A. Tavant, O. Vermorel, and W. Villafana. 2D axial-azimuthal particle-in-cell benchmark for low-temperature partially magnetized plasmas. *Plasma Sources Science and Technology*, 28(10):105010, October 2019.

²⁰Iain D. Boyd and Rainer A. Dressler. Far field modeling of the plasma plume of a Hall thruster. *Journal of Applied Physics*, 92(4):1764–1774, August 2002.

²¹Samuel J Araki, Robert S Martin, David L Bilyeu, Lubos Brieda, Carrie Hill, and Maria Choi. Current Capabilities of AFRL’s Spacecraft Simulation Tool. In *37th International Electric Propulsion Conference*, Cambridge, MA, USA, June 2022.

²²Thomas A. Marks and Alex A. Gorodetsky. Hall thruster simulations in Warp-X. In *38th International Electric Propulsion Conference*, Toulouse, France, June 2024.

²³Doruk Aksoy, Sruti Vutukury, Thomas A. Marks, Joshua D. Eckels, and Alex A. Gorodetsky. Compressed analysis of electric propulsion simulations using low rank tensor networks. Toulouse, France, June 2024.

²⁴David Manzella, Robert Jankovsky, Frederick Elliott, Ioannis Mikellides, Gary Jongeward, and Doug Allen. Hall Thruster Plume Measurements On-Board the Russian Express Satellites. In *27th International Electric Propulsion Conference*, Pasadena, CA, USA, October 2001.

²⁵Jonathan H. Tu, Clarence W. Rowley, Dirk M. Luchtenburg, Steven L. Brunton, and J. Nathan Kutz. On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, 1(2):391–421, December 2014.

²⁶Doruk Aksoy, David J Gorsich, Shravan Veerapaneni, and Alex A Gorodetsky. An incremental tensor train decomposition algorithm. *SIAM Journal on Scientific Computing*, 46(2):A1047–A1075, 2024.