

RESEARCH

Open Access



# Hall thruster model improvement by multidisciplinary uncertainty quantification

Joshua D. Eckels<sup>1\*</sup>, Thomas A. Marks<sup>1</sup>, Madison G. Allen<sup>1</sup>, Benjamin A. Jorns<sup>1</sup> and Alex A. Gorodetsky<sup>1\*</sup>

\*Correspondence:  
eckelsjd@umich.edu;  
goroda@umich.edu

<sup>1</sup>Department of Aerospace  
Engineering, University  
of Michigan, Ann Arbor,  
Michigan, USA

## Abstract

We study the analysis and refinement of a predictive engineering model for enabling rapid prediction of Hall thruster system performance across a range of operating and environmental conditions and epistemic and aleatoric uncertainties. In particular, we describe an approach by which experimentally-observed facility effects are assimilated into the model, with a specific focus on facility background pressure. We propose a multifidelity, multidisciplinary approach for Bayesian calibration of an integrated system comprised of a set of component models. Furthermore, we perform uncertainty quantification over the calibrated model to assess the effects of epistemic and aleatoric uncertainty. This approach is realized on a coupled system of cathode, thruster, and plume models that predicts global quantities of interest (QoIs) such as thrust, efficiency, and discharge current as a function of operating conditions such as discharge voltage, mass flow rate, and background chamber pressure. As part of the calibration and prediction, we propose a number of metrics for assessing predictive model quality. Based on these metrics, we found that our proposed framework produces a calibrated model that is more accurate, sometimes by an order of magnitude, than engineering models using nominal parameters found in the literature. We also found for many QoIs that the remaining uncertainty was not sufficient to account for discrepancy with experimental data, and that existing models for facility effects do not sufficiently capture experimental trends. Finally, we confirmed through a global sensitivity analysis the prior intuition that anomalous transport dominates model uncertainty, and we conclude by suggesting several paths for future model improvement. We envision that the proposed metrics and procedures can guide the refinement of future model development activities.

**Keywords:** Uncertainty quantification, Hall thrusters, Facility effects, Anomalous electron transport, Multidisciplinary surrogate

## Introduction

A primary goal within the Hall thruster modeling community is the development of predictive models that are practical for simulation-guided design, online control, test, and optimization [1]. Possession of a predictive Hall thruster model would accelerate Hall thruster development and qualification and improve safety margins for mission planning. It would also allow one to calibrate system parameters from ground-based testing and to disentangle the errors associated with facility effects for in-space prediction.

We contrast predictive models with the existing utility of models in the community. State-of-the-art modeling uses a fluid or fluid/particle hybrid simulation of the discharge channel and near-field to predict key quantities of interest (QoIs) such as thrust or efficiency. Some of these models, for example those based on fluid approximations, may approach computational accessibility for rapid prediction (i.e.  $\mathcal{O}(\text{hrs to days})$  for a single simulation), but are not “predictive” to date in the sense that they cannot make standalone, reliable predictions without significant experimental validation. Instead, they are often tuned to reproduce results for a *specific* operating condition and used as a “soft sensor” to inform unobserved quantities of the plasma and thruster performance at that condition.

Alternatively, we define a “predictive” Hall thruster model as one that reliably predicts key QoIs over a broad range of operating conditions. Such a model would have a single set of globally-relevant parameters over the full range of operating conditions and would necessarily be able to make accurate predictions even for conditions where no experimental data is available. In this work, we are primarily interested in predicting time-averaged performance metrics, including thrust, efficiency, discharge current, and ion velocity, over practical ranges of three specific operating conditions, namely the discharge voltage, anode mass flow rate, and background chamber pressure. We attempt global calibration of a set of model parameters and test the model’s extension to unseen operating conditions to determine its predictive quality by comparison to a prior baseline.

Within this predictive modeling framework, two primary issues inhibit the predictive quality of fluid Hall thruster models: *anomalous electron transport* and *facility effects*. The former results from simplifying kinetic plasma phenomena into a fluid framework. In real thrusters, these kinetic effects lead to turbulence-enhanced transport of electrons from the cathode to the anode [2, 3]. While expensive kinetic plasma simulations may be able to resolve these dynamics (analogously to direct numerical simulations in the turbulence community [4]), their closure in a fluid framework introduces significant uncertainty to model predictions [5] and has so far seen limited success [6]. The latter issue of facility effects deals with the well-known but poorly-understood feedback between thruster and environment during ground testing in a vacuum chamber [7, 8]. Unavoidable consequences of ground testing like elevated background pressure and thruster-facility electrical coupling are not representative of the space environment and inhibit model validation against ground test data. It is an ongoing task to develop models for facility effects that can be calibrated on and compared to ground test data, and ultimately that can be extrapolated to space-like conditions [9].

Prior work on closure modeling is extensive and empirical attempts at facility effects modeling are ongoing; we do not propose a new Hall thruster model nor any new insight into anomalous transport or facility effects here. Instead, due to the large uncertainties introduced by closure models and the coupled, multidisciplinary nature of thruster-facility interactions, we posit that our proposed approach can aid in the refinement and improvement of Hall thruster models by providing rapid feedback on the major sources of error and uncertainty. Uncertainty quantification (UQ) is of central importance to understand confidence in model predictions and the risks involved in using these

predictions for Hall thruster design, test, and optimization. UQ methods such as sensitivity analysis, model calibration, and experimental design also provide several ways to improve the model itself. However, these analyses are expensive, multi-query operations. As a result, we use a multifidelity (MF) and multidisciplinary (MD) surrogate-based approach to accelerate analysis.

Overall, we propose a bottom-up approach for Hall thruster model improvement, where we start with a simple baseline coupling between cathode, thruster, and plume and perform UQ to improve the most critical parts of the model iteratively. Similar couplings have been proposed in the literature [10–14], but we stress our focus on iterative model improvement guided by UQ and the incorporation of facility effects. To the best of our knowledge, this is the first attempt at global calibration of a fluid Hall thruster model on system-level data with uncertainty quantification of both anomalous transport and facility effects.

While this work analyzes a particular set of models for the Hall thruster system, we also propose a number of metrics to benchmark model development progress. These metrics include the mean and standard deviation of relative model errors to experimental data for each QoI, as well as comparisons of model error to experimental noise. Ideally, model error is on the same order as experimental noise and the variance in error is small, indicating a good fit of the model to the data and small remaining uncertainty, respectively. By these metrics, the models developed in this work demonstrated improved performance compared to nominal settings found in the literature, sometimes by an order of magnitude.

We first outline our MD Hall thruster modeling approach, then we describe our UQ framework and the specific algorithms we use, and finally we apply the framework for model improvement. This includes model calibration for tuning parameters, forward propagation for studying effects of uncertainty on model outputs, and sensitivity analysis for finding important parameters. These techniques require many evaluations of expensive computational models, and so we improve upon an existing MD surrogate method [15] to make these tasks feasible. Our contributions are **1)** a baseline integrated MD Hall thruster model implementation, **2)** an expanded MD surrogate framework that includes dimension reduction of field quantities, and **3)** model improvement suggestions from performing UQ. Ultimately, we found that the proposed framework provided an order of magnitude improvement in model predictions on a validation set for some QoIs, while remaining discrepancies with experimental data highlighted that further refinement of the underlying physical models may be required. A key benefit of our modular approach is that remaining shortcomings in the model can be addressed iteratively by targeting the most critical components first, such as adding new parameters or updating the relevant physics. We hope this framework will guide future iterations of predictive Hall thruster model development.

## Methods

In this section, we describe the Hall thruster system model, the uncertainty quantification problem, and the algorithmic approach.

### Hall thruster system model

A Hall thruster is an axisymmetric, cylindrical device that ionizes and accelerates a noble gas (typically Xenon) through an applied potential between an anode and an external cathode. The propellant is released through the anode and gets ionized and accelerated in the cylindrical discharge channel. The ion beam expands in the plume of the thruster and is neutralized by electrons released from the cathode. We model the Hall thruster system as a set of interacting component models: a cathode model, the thruster discharge channel model, and the far-field plume model, as shown in Fig. 1. These models are combined into a system model that provides desired quantities of interest such as cathode coupling voltage, thrust, efficiency, ion velocity, etc. The system model is a mapping

$$y = f(x) : \Gamma \rightarrow \mathcal{Y}, \tag{1}$$

from  $D$  system inputs  $x \in \Gamma \subseteq \mathbb{R}^D$  to  $Q$  system outputs  $y \in \mathcal{Y} \subseteq \mathbb{R}^Q$ . The inputs consist of three categories of variables  $x = (d, \theta, \phi)$ , where  $d$  corresponds to experimental conditions under which the thruster is operating,  $\theta$  are tunable model calibration parameters (e.g. anomalous transport coefficients), and  $\phi$  are other required inputs not in  $d$  or  $\theta$ . Tables 1 and 2 list the inputs and outputs, respectively.

The system model  $f$  is formed through the integration of component models:

$$y_k = f_k(x_k, \xi_k) : \Gamma_k \times \Xi_k \rightarrow \mathcal{Y}_k, \quad k = 1 \dots K, \tag{2}$$

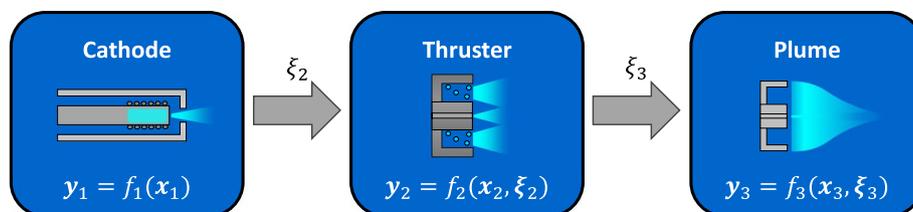
where  $x_k \in \Gamma_k \subseteq \mathbb{R}^{D_k}$  and  $y_k \in \mathcal{Y}_k \subseteq \mathbb{R}^{Q_k}$ . The coupling variables  $\xi_k \in \Xi_k \subseteq \mathbb{R}^{S_k}$  serve to map outputs from some components into inputs of others. The Hall thruster system we study here consists of a feed-forward integration of three components — a cathode, thruster, and plume model as shown in Fig. 1. These component models are described in the following subsections.

#### Cathode

The cathode is modeled by the semi-empirical approach of Ref. [16]. This model predicts the cathode coupling voltage  $V_{cc}$  as a function of background pressure  $P_B$ :

$$V_{cc} = V_{vac} + T_e \log \left[ 1 + \frac{P_B}{P_T} \right] - \left[ \frac{T_e}{P_T + P^*} \right] P_B, \tag{3}$$

where  $T_e$  is the temperature of the electrons at the cathode and  $(V_{vac}, P^*, P_T)$  are other tunable model parameters given in Table 1. This model accounts for  $V_{cc}$ -driven performance loss, which is observed to depend on the chamber background pressure



**Fig. 1** Hall thruster system model as a three-component feed-forward system

**Table 1** Summary of multidisciplinary system model inputs

System index	Variable description	Symbol	Units	Component inputs	Distribution	Domain
1	Background pressure	$P_B$	Torr	$\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3$	$\mathcal{U}(\pm 20\%)$	$[10^{-8}, 10^{-4}]$
2	Anode voltage	$V_a$	V	$\mathbf{d}_1, \mathbf{d}_2$	-	[200, 400]
3	Anode mass flow rate	$\dot{m}_a$	mg/s	$\mathbf{d}_2$	$\mathcal{U}(\pm 3\%)$	[2, 7]
4	Cathode electron temperature	$T_e$	eV	$\theta_1, \theta_2$	$\mathcal{U}(1, 5)$	[1, 5]
5	Vacuum coupling voltage	$V_{vac}$	V	$\theta_1$	$\mathcal{U}(0, 60)$	[0, 60]
6	Turning point pressure	$P^*$	$\mu\text{Torr}$	$\theta_1$	$\mathcal{U}(10, 100)$	[10, 100]
7	Thruster to facility plasma density ratio	$P_T$	$\mu\text{Torr}$	$\theta_1$	$\mathcal{U}(10, 100)$	[10, 100]
8	Anode injection neutral velocity	$u_n$	m/s	$\theta_2$	$\mathcal{U}(100, 500)$	[100, 500]
9	Inner to outer transition length	$l_t$	mm	$\theta_2$	$\mathcal{U}(1, 20)$	[1, 20]
10	Anomalous transport coefficient	$a_1 (10^6)$	-	$\theta_2$	$\mathcal{U}(-2.5, -1)$	$[-2.5, -1]$
11	Anomalous transport coefficient	$a_2$	-	$\theta_2$	$\mathcal{U}(10, 100)$	[10, 100]
12	Shift displacement	$\Delta z$	-	$\theta_2$	$\mathcal{N}(0.2, 0.07)$	[0, 0.4]
13	Upstream shift axial limit	$z_0$	-	$\theta_2$	$\mathcal{N}(-0.12, 0.04)$	$[-0.25, 0]$
14	Upstream shift pressure limit	$p_0$	$\mu\text{Torr}$	$\theta_2$	$\mathcal{N}(45, 7)$	[25, 65]
15	Scatter/main beam ratio	$c_0$	-	$\theta_3$	$\mathcal{U}(0, 1)$	[0, 1]
16	Divergence angle ratio	$c_1$	-	$\theta_3$	$\mathcal{U}(0.1, 0.9)$	[0.1, 0.9]
17	Divergence angle slope	$c_2$	rad/Pa	$\theta_3$	$\mathcal{U}(-15, 15)$	$[-15, 15]$
18	Divergence angle offset	$c_3$	rad	$\theta_3$	$\mathcal{U}(0, \pi/2)$	$[0, \pi/2]$
19	Neutral density slope	$c_4 (10^6)$	$m^{-3}/\text{Pa}$	$\theta_3$	$\mathcal{U}(18, 22)$	[18, 22]
20	Neutral density offset	$c_5 (10^6)$	$m^{-3}$	$\theta_3$	$\mathcal{U}(14, 18)$	[14, 18]
21	Charge exchange cross sectional area	$\sigma_{cex}$	$\text{\AA}^2$	$\phi_3$	$\mathcal{U}(51, 58)$	[51, 58]
22	Plume radius	$r_p$	m	$\phi_3$	-	[0.5, 1.5]

The system index denotes the variable's index into the aggregate set of system inputs  $\mathbf{x}$ . The elements in component inputs indicate the variable's classification as a **1) d** - environment operating condition, **2)  $\theta$**  - model calibration parameter, or **3)  $\phi$**  - other variable input; the subscripts indicate component model inputs 1, 2, 3 for the cathode, thruster, plume models, respectively. Variables with the  $(10^6)$  notation indicate a log-uniform distribution. The arguments of the normal distribution  $\mathcal{N}(\mu, \sigma)$  are the mean  $\mu$  and standard deviation  $\sigma$

(assuming  $P_B$  is measured at or near the thruster exit plane). The predicted cathode coupling voltage  $V_{cc}$  is then passed downstream to the thruster model.

### Thruster

The thruster discharge chamber and near-field plume is modeled as a one-dimensional fluid and simulated with the *Hallthruster.jl* code [17]. This thruster model couples to the cathode model by using  $V_{cc}$  as the potential boundary condition at the cathode end of the 1d domain. The thruster couples to the far-field plume model by predicting the total ion beam current  $I_b$  in the plasma discharge. All species are treated as fluids in the plasma flow axial direction with appropriate continuity and

**Table 2** Summary of multidisciplinary system model outputs

System index	Variable description	Symbol	Units	Component outputs	Coupling variables
1	Cathode coupling voltage	$V_{cc}$	V	$\mathbf{y}_1$	$\xi_2$
2	Total ion beam current	$I_b$	A	$\mathbf{y}_2$	$\xi_3$
3	Discharge current	$I_d$	A	$\mathbf{y}_2$	-
4	Thrust	$T$	mN	$\mathbf{y}_2$	-
5	Current utilization efficiency	$\eta_c$	-	$\mathbf{y}_2$	-
6	Average axial ion velocity	$u_{ion}(z)$	m/s	$\mathbf{y}_2$	-
7	Plume ion current density	$j_{ion}(\gamma)$	A/m <sup>2</sup>	$\mathbf{y}_3$	-

The system index denotes the variable's index into the aggregate set of system outputs  $\mathbf{y}$ . The component outputs column indicates the variable's associated component model: 1, 2, 3 for the cathode, thruster, plume models, respectively. Coupling variable connections are indicated in that column when applicable

energy/momentum conservation. The background pressure is treated as an additional “ingested” flux of neutrals  $\Gamma_{\text{ingest}}$ , equivalent to the thermal flux of background neutrals across the exit plane of the thruster:

$$\Gamma_{\text{ingest}} = \frac{n_{n,B}}{4} \sqrt{\frac{8k_B T_B}{\pi m}}, \quad (4)$$

for neutral mass  $m$ , background temperature  $T_B$ , and background neutral density  $n_{n,B} = P_B/(k_B T_B)$  with the Boltzmann constant  $k_B$ . Effectively, the code adds an “ingested propellant” to the actual propellant  $\Gamma_{\text{prop}}$ , resulting in a total flux of neutrals at the anode of  $\Gamma_{\text{anode, tot}} = \Gamma_{\text{prop}} + \Gamma_{\text{ingest}}$ .

We study two performance metrics predicted by the thruster model, namely the thrust  $T$  and the current utilization efficiency  $\eta_c$ . The thrust is computed as the net momentum flux of all ion charge states and the current utilization efficiency is computed as the ratio of ion beam current to total discharge current  $I_d$ , i.e.  $\eta_c = I_b/I_d$  from Ref. [18].

To model the anomalous electron transport, we employ a simple two-zone Bohm-like closure model of the anomalous collision frequency,  $\nu_{\text{AN}}$  [19]. In this model, the anomalous collision frequency has a stepped profile, with a comparatively low value at axial locations upstream of a pressure-dependent threshold location,  $z^*$ , and a higher value downstream of  $z^*$ . Formally, the anomalous collision frequency is given as a function of axial location  $z$  by

$$\nu_{\text{AN}}(z) = \begin{cases} a_1 \omega_{ce}, & z < z^*, \\ a_2 a_1 \omega_{ce}, & z > z^*, \end{cases} \quad (5)$$

where  $\omega_{ce}$  is the electron cyclotron frequency and  $(a_1, a_2)$  are tunable model parameters corresponding to the anomalous collision frequency inside and outside of the channel. Piecewise Bohm-like models like this are common in fluid simulations of Hall thrusters [19]; however, they are typically specific to a given thruster and operating condition and do not readily generalize across devices and operating regimes [20]. We believe the proposed model will be suitable for the single thruster under study and the relatively narrow range of operating conditions considered.

We account for the upstream shift in acceleration region due to facility pressure [21] by allowing the transition location between low and high anomalous collision frequency zones ( $z^*$ ) to vary with pressure. Based on the ansatz that this location should not shift arbitrarily far upstream or downstream, we adopt a simple logistic model of  $z^*$ , given by

$$z^*(P_B) = L_{ch} \left( 1 + z_0 + \frac{\Delta z}{1 + b^{2P_B/p_0 - 1}} \right), \quad (6)$$

where  $L_{ch}$  is the length of the thruster channel and we have introduced four new parameters:

- $\Delta z$  — the total displacement from  $P_B \rightarrow 0$  to  $P_B \rightarrow \infty$ , expressed as a fraction of  $L_{ch}$ ,
- $z_0$  — the upstream axial limit of the velocity profile shift as  $P_B \rightarrow \infty$ , expressed as a fraction of  $L_{ch}$ ,
- $p_0$  — the background pressure at which the upstream shift plateaus, and
- $b$  — the slope of the pressure response curve (which is treated as a constant  $b = 14$  due to relative insensitivity during initial tests).

We note the model in Eq. (6) is phenomenological rather than physical, and is included to better capture the experimentally-observed upstream shift in acceleration region at higher background pressures. We consider the ion velocity profile  $u_{ion}(z)$  predicted by the thruster model along channel centerline in the axial direction  $z$ , which allows us to study the impact of Eq. (6) on ion acceleration in the channel.

Lastly, we note that the discharge current  $I_d = I_e + I_b$  is composed of the electron current streaming backward from cathode to anode ( $I_e$ ) and the ion beam current  $I_b$ . For constant  $I_d$ , an increase in backstreaming electron current decreases the current efficiency since a smaller fraction of the total current through the device is usable as thrust in the beam. We are particularly interested in this efficiency mode as it is the primary efficiency loss mechanism in Hall thrusters ( $\eta_c < 80\%$ , typically [18]). Additionally, it directly captures the effects of anomalous electron transport on thruster performance. As the electron current across the thruster's magnetic field is directly related to the effective anomalous collision frequency [2], high values of collision frequency will tend to decrease the current utilization efficiency.

### Plume

The far-field plume is modeled via another semi-empirical form [22] to predict the expansion of ion current density  $j_{ion}$  in the Hall thruster far-field plume resulting from the total beam current  $I_b$  at the thruster exit plane. The plume model treats the ion current density in the plume as the superposition of the main beam plasma  $j_{beam}$ , inelastic ion-neutral scattering  $j_{scatter}$ , and charge-exchange collisions  $j_{cex}$ :

$$j_{ion} = j_{beam} + j_{scatter} + j_{cex}. \quad (7)$$

The model assumes empirically-guided shapes for each ion population and encodes a facility pressure  $P_B$  dependence of characteristic divergence angles within the plume. The model is detailed in Ref. [22], but we provide a condensed and modified form of the model in Appendix A.

### Model summary

As stated previously, our main goal is to capture the impact of *anomalous transport* and *facility effects* on Hall thruster performance. We account for anomalous transport in the thruster model via Eqs. (5) and (6). We account for facility effects by coupling cathode, thruster, and plume models in a multidisciplinary framework (Fig. 1). Facility effects specifically enter the framework as the background chamber pressure  $P_B$  in four distinct locations:

1. Equation (3) as an impact on the cathode coupling voltage,
2. Equation (4) as an additional “ingested” flux of neutrals from the vacuum chamber,
3. Equation (6) as a shift to the two-zone Bohm anomalous transport closure model, and
4. Equation (7) as an impact on characteristic divergence angles within the plume (see Appendix A).

### Uncertainty quantification

Many engineering models have inherent uncertainty due to concerns like manufacturing tolerances, incomplete or missing knowledge, and natural variability. Accounting for this uncertainty is critical to establish confidence in model predictions. In this section, we describe the uncertainty quantification problem for the Hall thruster system model provided in the previous section.

The uncertainties in our problem are captured in the model inputs, which are endowed with a probability density function (PDF)  $\rho(\boldsymbol{x})$ . The goal is to characterize the uncertainty in the model predictions — which can be fully characterized by a density  $\rho(\boldsymbol{y})$  induced by the system model — due to uncertainty in the inputs. To this end, we treat  $\rho(\boldsymbol{d}, \boldsymbol{\phi})$  as *aleatoric* sources of uncertainty, i.e. they are irreducible for a given set of processes (measurement uncertainty of pressure, for example, can't be reduced except by improving the measurement process). We treat  $\rho(\boldsymbol{\theta})$  as *epistemic* sources of uncertainty, i.e. the uncertainty is due to lack of knowledge and can be improved by gathering data/information.

The uncertainties for the inputs are provided in Table 1. For each input  $x_i$ ,  $i = 1 \dots D$ , we specify the system index  $i$ , the uncertainty  $\rho(x_i)$  as a probability distribution, and the domain  $\Gamma(x_i)$  as a set of bounds  $[l_i, u_i]$ . This approach treats each input as an independent random variable (i.e.  $\rho(\boldsymbol{x}) = \prod_{i=1}^D \rho(x_i)$ ) since we have no prior knowledge on any existing correlations, (provided enough data, the posterior distribution after model calibration is insensitive to the choice of prior). For operating conditions  $\boldsymbol{d}$ , the uncertainty is specified relative to some nominal value, e.g.  $\mathcal{U}(\pm 20\%)$  is a uniform distribution between values 20% above and below the current set point. Variables with no specified distribution are deterministic (i.e. negligible uncertainty). All uncertainties represent our prior knowledge on the expected range and distribution for each variable, based on available experimental or numerical evidence. We conservatively assign large bounds to parameters where prior knowledge is limited. Details on the formulation of each variable can be found in the appropriate references for each component model.

The ultimate goal of this UQ framework is to provide feedback for *model improvement* activities. We accomplish this by **1)** reducing epistemic uncertainty  $\rho(\theta)$  with model calibration, **2)** quantifying output uncertainty  $\rho(y)$  with forward propagation, **3)** narrowing the input space and identifying important parameters with sensitivity analysis, and **4)** making other suggestions from model validation results (such as increasing model fidelity, updating model assumptions, adding extra parameters, etc.). For all of these UQ tasks, the computation of output statistics is performed using sample-based estimators, such as the Monte Carlo estimate of the mean:

$$\mathbb{E}[f(\mathbf{x})] \approx \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}^{(i)}), \quad \mathbf{x}^{(i)} \sim \rho(\mathbf{x}), \quad (8)$$

which requires  $N$  evaluations of the system model  $f(\mathbf{x}^{(i)})$ . Often,  $N$  can approach  $\mathcal{O}(10^6)$  for sufficient accuracy. In our case, the cost of a single evaluation of the thruster model is  $\mathcal{O}(1 \text{ min})$ , so  $N$  evaluations could require weeks of runtime even on a high-performance computing system. We thus require a cheap-to-evaluate surrogate model that approximates the input-output behavior of the full model.

### Algorithms

We now summarize the specific methods used in this study to perform uncertainty quantification on the Hall thruster model in a computationally feasible way. We first describe a surrogate method that enables UQ, and then we describe the UQ algorithms for model calibration, forward uncertainty propagation, and sensitivity analysis. The results of these analyses are presented in the [Results](#) section.

### Multidisciplinary surrogate

We seek to build a surrogate of the MD system in Fig. 1 to reduce computational burden as efficiently as possible, and so we take advantage of multifidelity (MF) techniques that combine model evaluations from a hierarchy of modeling fidelities (for example, increasing cell resolution for PDE discretization). We choose the multi-index stochastic collocation (MISC) framework as it provides an effective way to build adaptive MF surrogates [23–25]. MISC approximates each component model  $f_k$  as:

$$f_k(\mathbf{x}_k, \boldsymbol{\xi}_k) \approx f_{k, \mathcal{I}_k}(\mathbf{x}_k, \boldsymbol{\xi}_k) = \sum_{[\alpha, \beta] \in \mathcal{I}_k} c_{k, [\alpha, \beta]} f_{k, [\alpha, \beta]}(\mathbf{x}_k, \boldsymbol{\xi}_k), \quad (9)$$

where  $\alpha = [\alpha_1, \dots, \alpha_{R_k}]$  is a set of  $R_k$  indices (i.e. a multi-index) that specifies model fidelity and  $\beta = [\beta_1, \dots, \beta_{N_k}]$  is a similar set of  $N_k$  indices that specifies surrogate fidelity. For a given set  $\mathcal{I}_k$  of concatenated multi-indices  $[\alpha, \beta]$ , the MISC approximation is a linear combination of surrogates of varying fidelity  $f_{k, [\alpha, \beta]}$ , whose coefficients  $c_{k, [\alpha, \beta]}$  have a closed-form expression given some assumptions on the structure of  $\mathcal{I}_k$  [15]. The individual surrogates  $f_{k, [\alpha, \beta]}$  can take on any form (i.e. Gaussian processes, neural networks, polynomials, etc.) so long as they have some way to tune the underlying physical model fidelity via  $\alpha$  and some way to tune the parametric surrogate fidelity via  $\beta$ . We follow prior work (and leave implementation details) in Ref. [15] to adaptively build the index sets  $\mathcal{I}_k$  in an MD/MF setting using multivariate Lagrange polynomial interpolants

for the individual  $f_{k,[\alpha,\beta]}$ . The Lagrange interpolants are a natural choice for the adaptive MISC approach assuming sufficient smoothness of the underlying function given their tensor-product grid structure and a nested sampling/collocation strategy [26]. We denote the full MD system surrogate as  $f_{\mathcal{J}}(\mathbf{x}) \approx f(\mathbf{x})$ ,  $\mathcal{J} = [\mathcal{I}_1, \dots, \mathcal{I}_K]$ .

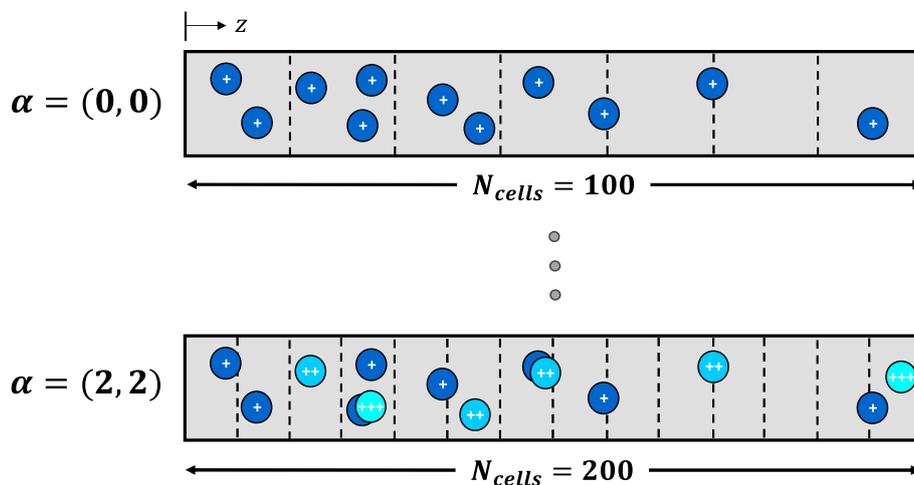
In our current setting, we use  $R_k = 2$  model fidelities for the thruster model. The first index  $\alpha_1$  controls the number of grid cells in the simulation via  $N_{\text{cells}} = 50(\alpha_1 + 2)$  for  $\alpha_1 = 0, 1, 2$ . The second index  $\alpha_2$  controls the number of ion charge states included in the simulation via  $N_{\text{charge}} = \alpha_2 + 1$  for  $\alpha_2 = 0, 1, 2$ , with increasing model fidelity in both cases for increasing  $\alpha$ . Figure 2 illustrates the difference between the lowest and highest fidelity versions of the thruster model.

The number of surrogate training points along each input dimension  $i$  is controlled by  $N_{\text{points},i} = 2\beta_i + 1$  for  $i = 1 \dots N_k$  (with  $N_k = 12$  inputs for the thruster model). The surrogate approximation in Eq. (9) increases in accuracy as more sets of  $[\alpha, \beta]$  are added to the index set  $\mathcal{I}_k$  using the adaptive procedure described in Ref. [15].

We apply the MISC methodology in the context of high-dimensional inputs, coupling variables, and outputs, which quickly encounters the curse of dimensionality. To deal with this challenge, we perform linear dimension reduction (i.e. principal orthogonal decomposition [27–29]) on high-dimensional spatial field quantities in the model. For example, we find a low-dimensional subspace (or “latent” space) of the numerical ion velocity profile, and the surrogate is constructed for the latent coefficients of the ion velocity in the reduced space; details are included in Appendix B.

**Model calibration**

One of the primary goals in our UQ framework is to calibrate the model coefficients  $\theta$  in Table 1 to fit a set of experimental data  $y_e$ . We use data from the 1.35 kW SPT-100 operating on Xenon for calibration due to the thruster’s widespread use and the availability of data in the literature (see datasets 1-4 described in Appendix C). Within the context of our UQ framework, we perform the calibration through Bayesian inference by sampling from the posterior distribution on  $\theta$  given the experimental data  $y_e$ , that is  $\rho(\theta | y_e) \propto \rho(y_e | \theta)\rho(\theta)$ . The prior  $\rho(\theta)$  represents our state of knowledge on the



**Fig. 2** 1d axial simulation domain of the thruster model from the lowest fidelity  $\alpha = (0, 0)$  to the highest fidelity  $\alpha = (2, 2)$

parameters *before* seeing the data and the likelihood  $\rho(\mathbf{y}_e | \boldsymbol{\theta})$  indicates how likely the data was generated by a given set of parameters. The posterior  $\rho(\boldsymbol{\theta} | \mathbf{y}_e)$  is ideally much narrower than the prior, indicating that we have *reduced* epistemic uncertainty by learning from the observed data. To formulate the likelihood, we model the experimental data as normally-distributed about the forward model predictions:

$$\mathbf{y}_e = f(\boldsymbol{\theta}) + \boldsymbol{\xi}, \quad \boldsymbol{\xi} \sim \mathcal{N}(0, \sigma^2), \quad (10)$$

where  $\boldsymbol{\xi}$  is additive Gaussian noise with characteristic variance  $\sigma^2$ . We have chosen here to neglect the aleatoric uncertainties  $\rho(\mathbf{d}, \boldsymbol{\phi})$  to simplify the analysis and since we expect their impact to be small. As a result, we are able to express the system model  $f(\mathbf{x})$  from Eq. (1) as simply  $f(\boldsymbol{\theta})$ , with  $(\mathbf{d}, \boldsymbol{\phi})$  fixed to nominal conditions. The likelihood  $\rho(\mathbf{y}_e | \boldsymbol{\theta})$  results from the Gaussian form of Eq. (10) and, assuming independent data  $\mathbf{y}_e = \{y_e^{(i)}\}_{i=1}^{N_e}$ , is given by:

$$\rho(\mathbf{y}_e | \boldsymbol{\theta}) = \prod_{i=1}^{N_e} \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{(y_e^{(i)} - f(\boldsymbol{\theta}))^2}{2\sigma_i^2}\right), \quad (11)$$

where we have expressed the experimental noise  $\sigma_i^2$  on a case-by-case basis for each  $y_e^{(i)}$ . It is understood that we extract from  $f(\boldsymbol{\theta})$  only the QoIs that correspond to the current data point  $y_e^{(i)}$ , e.g. if  $y_e^{(i)}$  is a cathode coupling voltage measurement, then we extract  $V_{cc}$  from index 1 of  $f(\boldsymbol{\theta})$ , as detailed in Table 2. We also substitute in place the surrogate  $f_{\mathcal{J}}(\boldsymbol{\theta})$  for  $f(\boldsymbol{\theta})$  during the actual computations, as using the full model for even a modest number of samples is simply infeasible.

Having defined the prior  $\rho(\boldsymbol{\theta})$  in Table 1 and the likelihood  $\rho(\mathbf{y}_e | \boldsymbol{\theta})$  in Eq. (11), the unnormalized posterior  $\rho(\boldsymbol{\theta} | \mathbf{y}_e)$  can be computed through Bayes' rule. We use a delayed-rejection adaptive Metropolis Markov chain Monte Carlo (MCMC) routine to sample from the 17d posterior [30].

### **Forward propagation and sensitivity analysis**

After calibrating the model, we study the impacts of remaining uncertainty on model predictions by propagating input uncertainties  $\rho(\mathbf{x})$  through the surrogate  $f_{\mathcal{J}}(\mathbf{x})$  to obtain the output distribution  $\rho(\mathbf{y})$ . We do this by propagating  $N$  Monte Carlo samples of the uncertain inputs  $\{\mathbf{x}^{(i)}\}_{i=1}^N \sim \rho(\mathbf{x})$  through the surrogate to obtain  $N$  predictions of the uncertain outputs  $\mathbf{y}^{(i)} = f_{\mathcal{J}}(\mathbf{x}^{(i)})$ . We then use the 5th and 95th percentiles of the outputs as credible intervals for the model predictions.

Lastly, we take a preliminary step in improving future versions of the model by way of sensitivity analysis. By studying the impacts of model input uncertainty on resulting output variance, we can find the parameters that have the greatest impact on model predictions and consequently remove unimportant parameters from future consideration. Since we are primarily interested in studying the impact of facility effects on the model, we perform the sensitivity analysis over varying background pressures. We choose the Sobol' method of global variance analysis to study all uncertain inputs in Table 1 over their full prior domains [31, 32]. In the Sobol' analysis, we compute the first-order ( $S_{1i}$ ) and total-order ( $S_{Ti}$ ) sensitivity indices for all inputs  $i = 1 \dots D_k$  for each component model  $k$ . The first-order indices measure the independent contributions of each input on total output

variance and have the property  $\sum_i S_{1i} < 1$ , where the remaining fraction  $1 - \sum_i S_{1i}$  is due to “higher-order” interactions between coupled inputs. The total-order indices additionally include the higher-order interactions in which input  $i$  is involved, with  $S_{Ti} > S_{1i}$  indicating the presence of higher-order interactions. The property  $\sum_i S_{Ti} \geq 1$  is due to double counting of interactions between inputs  $i$  and  $j$  in  $S_{Ti}$  and  $S_{Tj}$ ,  $i \neq j$ . We leave full descriptions of the Sobol’ indices in Ref. [32], but note that we must use sample-based estimators to obtain their values. As such, we include Monte Carlo 95% confidence intervals on all estimates of Sobol’ indices with  $N = 5000$  samples.

## Results

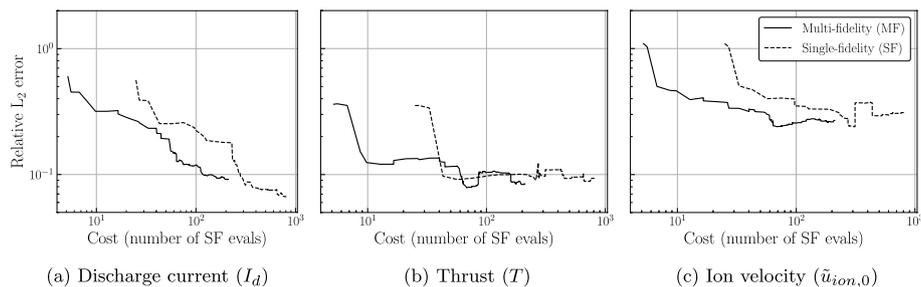
In this section, we present the results of the UQ methods applied to the Hall thruster model and assess goodness of fit with experimental data. We first quantify surrogate training accuracy, then we summarize overall reduction in uncertainty from model calibration, next we compare model trends to experiment, and finally we present results from sensitivity analysis.

### Hall thruster model surrogate

We first report on the construction of a surrogate for the MD Hall thruster system described in Fig. 1. Since the cathode and plume components are analytical, only the thruster model requires a surrogate. We evaluate the accuracy of the surrogate by comparing to an independent test set with the relative  $L_2$  error metric. We define the relative  $L_2$  percent error between two vector quantities  $\mathbf{v}$  and  $\mathbf{w}$  as

$$\text{Relative } L_2 \text{ error} = E(\mathbf{v}, \mathbf{w}) = \frac{\|\mathbf{v} - \mathbf{w}\|_2}{\|\mathbf{v}\|_2} \quad (12)$$

and report the error  $E(\mathbf{y}_m, \mathbf{y}_s)$  between the full model  $\mathbf{y}_m = f(\mathbf{x})$  and surrogate  $\mathbf{y}_s = f_{\mathcal{J}}(\mathbf{x})$  on  $N = 500$  test set samples  $\mathbf{x} = \{\mathbf{x}^{(i)}\}_{i=1}^N$ . Figure 3 shows the  $L_2$  test set error during training as a function of the total incurred computational cost (measured in units of total number of *equivalent* high-fidelity model evaluations). The  $L_2$  error is compared when using the multi-fidelity strategy vs. a single high-fidelity strategy (i.e.  $\alpha = (2, 2)$ ). Only three QoIs are shown, namely the discharge current ( $I_d$ ), the thrust ( $T$ ), and the first ion velocity latent coefficient ( $\tilde{u}_{ion,0}$ ).



**Fig. 3** Comparison of multi-fidelity vs. single-fidelity ( $\alpha = (2, 2)$ ) relative  $L_2$  test set error during surrogate training against computational cost (in units of number of *equivalent* high-fidelity model evaluations). The multi-fidelity strategy exhibits better accuracy for a given cost, and surrogate error reaches 10% during training for most QoIs

In most cases, employing the multi-fidelity strategy results in a *more accurate* surrogate than using the single highest fidelity alone for a given computational budget. This result highlights the importance of the MISC approach we adopted in this study. We expect the relative improvement to be even greater as higher-fidelity (and more costly) models are incorporated. We also found that the surrogate prediction of many QoIs reaches 10% accuracy or better for the equivalent cost of around 200 full high-fidelity model evaluations over a 12d input space. Further details on surrogate performance are provided in Appendix B.

### Model calibration

Next, we use the surrogate to calibrate all system model coefficients  $\theta$  simultaneously to fit experimental data using the Bayesian inference procedure described in the [Model calibration](#) section. MCMC is performed for 100,000 iterations with a burn-in fraction of 10% and a resulting acceptance ratio of 22%. The posterior marginals are summarized in Table 3 for each calibrated variable. We note that one variable was essentially reduced to a constant with standard deviation less than 0.1%, and so we only indicate a point value. We also include the original priors for each variable from Table 1 for comparison. Appendix D additionally provides plots of the 1d and 2d marginals for each component model.

We make four remarks regarding the posterior distribution:

**Remark 1** The cathode model parameters  $\theta_1 = (T_e, V_{vac}, P^*, P_T)$  agree with previously reported values in Ref. [16].

**Table 3** Statistics of the 1d marginal posteriors of the Hall thruster model fit coefficients  $\theta$

Variable	Prior	Posterior					
		Min	5 <sup>th</sup> pctile	50 <sup>th</sup> pctile	95 <sup>th</sup> pctile	Max	Std dev
$T_e$	$\mathcal{U}(1, 5)$	1.00	1.03	1.33	1.81	2.41	0.25
$V_{vac}$	$\mathcal{U}(0, 60)$	31.2	31.4	31.6	31.8	31.9	0.1
$P^*$	$\mathcal{U}(10, 100)$	25.8	30.2	34.6	41.5	50.3	3.5
$P_T$	$\mathcal{U}(10, 100)$	10.0	10.0	10.2	10.9	12.5	0.3
$u_n$	$\mathcal{U}(100, 500)$	133.0	137.1	141.2	145.3	149.7	2.5
$l_t$	$\mathcal{U}(1, 20)$	1.57	1.73	1.88	2.03	2.27	0.09
$a_1 (10^4)$	$\mathcal{U}(-2.5, -1)$	-2.183	-2.174	-2.167	-2.161	-2.153	0.004
$a_2$	$\mathcal{U}(10, 100)$	14.0	14.3	14.6	14.9	15.3	0.2
$\Delta z$	$\mathcal{N}(0.2, 0.07)$	-	-	0.4	-	-	< 0.1%
$z_0$	$\mathcal{N}(-0.12, 0.04)$	-0.035	-0.033	-0.031	-0.029	-0.027	0.001
$p_0$	$\mathcal{N}(45, 7)$	40.5	49.9	56.9	63.1	65.0	4.0
$c_0$	$\mathcal{U}(0, 1)$	0.01	0.46	0.92	0.99	1.00	0.18
$c_1$	$\mathcal{U}(0.1, 0.9)$	0.53	0.70	0.81	0.89	0.90	0.06
$c_2$	$\mathcal{U}(-15, 15)$	7.9	11.4	14.0	14.9	15.0	1.1
$c_3$	$\mathcal{U}(0, \pi/2)$	0.28	0.38	0.45	0.50	0.54	0.04
$c_4 (10^4)$	$\mathcal{U}(18, 22)$	18.0	18.1	18.8	19.3	19.6	0.4
$c_5 (10^4)$	$\mathcal{U}(14, 18)$	14.0	14.1	15.1	16.3	16.9	0.7

Variables with the  $(10^4)$  notation indicate a log-uniform distribution. The arguments of the normal distribution  $\mathcal{N}(\mu, \sigma)$  are the mean  $\mu$  and standard deviation  $\sigma$

**Remark 2** Most of the model parameters have significantly reduced uncertainty under the posterior as compared to the prior, indicating that the data was very informative in learning the parameters.

**Remark 3** The uncertainty was reduced so significantly for the  $\Delta z$  parameter that it can practically be treated as a point value. This likely indicates the model is over-parameterized and can be simplified by removing extraneous parameters. It also implies that further data collection will likely not reduce the epistemic uncertainty further.

**Remark 4** Lastly, many parameters ( $P_T, \Delta z, c_0, c_2$ ) were pushed to the boundaries of the prior, indicating a better fit may lie outside the original expected bounds. In the case of  $c_0$ , where values greater than one are nonphysical, this may indicate shortcomings intrinsic in the model itself.

### Quantifying uncertainty

In the last section, we obtained global, best-fit values for the model parameters and their uncertainties through Bayesian inference; this was enabled by the use of a surrogate approximation to the true model. Ideally, the best-fit parameters learned by the surrogate would improve the quality and reduce the uncertainty of predictions made by the true model. The overarching idea with this approach is that the purely data-driven surrogate is only valid within the narrow bounds of training data, while the true model is grounded in physics or empirical evidence and should have predictive power outside of the training regime. Therefore, the primary goal is to improve the true model's fit to experimental data by proxy of improving the surrogate.

In this section, we measure how well this goal has been met by showing several plots of model predictions with uncertainty against the training data for each component model, both before and after parameter calibration. "Prior predictive" plots give uncertainty bounds obtained with  $N = 1000$  Monte Carlo samples of the uncertain inputs under the prior distribution  $\{\mathbf{x}^{(i)}\}_{i=1}^N \sim \rho(\mathbf{d}, \boldsymbol{\phi}, \boldsymbol{\theta})$  described in Table 1, and they represent the 5th and 95th percentiles of resulting surrogate predictions  $\mathbf{y}^{(i)} = f_{\mathcal{J}}(\mathbf{x}^{(i)})$  via shaded regions. "Posterior predictive" plots give uncertainty bounds in the same way, except the model parameters are instead sampled from the posterior distribution  $\{\boldsymbol{\theta}^{(i)}\}_{i=1}^N \sim \rho(\boldsymbol{\theta} | \mathbf{y}_e)$ . The posterior predictive contains the same aleatoric uncertainty as the prior predictive from  $\rho(\mathbf{d}, \boldsymbol{\phi})$  and any remaining epistemic uncertainty from  $\rho(\boldsymbol{\theta} | \mathbf{y}_e)$ . Except where otherwise indicated, we plot the median surrogate prediction with a dashed line and the true system model prediction with a solid line for comparison. The true model predictions use the median parameter values in Table 3.

We also note that when the uncertainty from model predictions is on the same order as the experimental noise (as was typical under the posterior predictive), it is sometimes useful to further corrupt model predictions by additive Gaussian noise according to Eq. (10) to more fully represent the posterior predictive (this is equivalent to emulating new experimental data  $\mathbf{y}^*$  from the posterior via  $\mathbf{y}^*, \boldsymbol{\theta} \sim \rho(\mathbf{y}^*, \boldsymbol{\theta} | \mathbf{y}_e)$ ). The additional "measurement" uncertainty emulated by this process is indicated by the shaded blue regions where applicable.

### Cathode model uncertainty

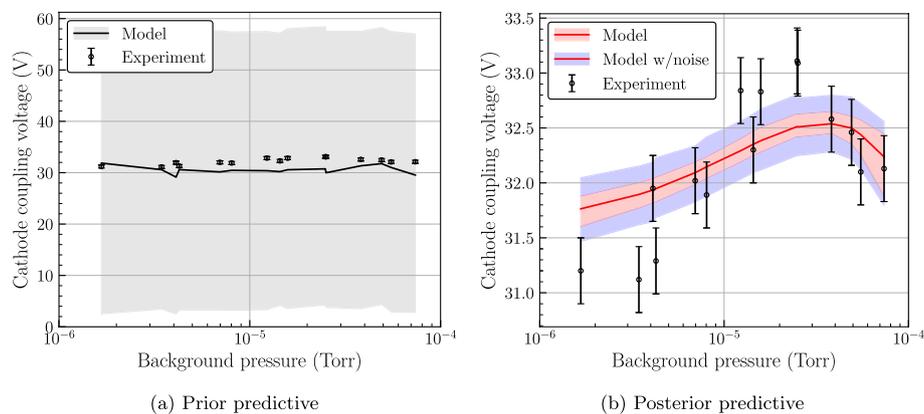
The system-level output of the cathode model is the cathode coupling voltage. Experiments show that this performance-loss mechanism is influenced by elevated vacuum chamber pressures. Higher coupling voltages decrease the total potential drop seen by the accelerating ions, so it is important that our models capture changes in coupling voltage due to changing background pressure to accurately predict thruster performance. Figure 4 compares the model predictions to experimental data from Ref. [16]. We observe a significant reduction in uncertainty under the posterior, mostly attributable to the tight posterior bounds on the  $V_{vac}$  parameter (see Appendix D). We also observe higher uncertainty for predictions at higher pressures due to the relative uncertainty in the pressure itself. As the background pressure approaches 0, the model asymptotes at  $V_{vac} = 31.6$  from Eq. (3) while the data appears to continue decreasing. We hope that such a physics-based extrapolation to zero pressure is well-founded, but ultimately further data collection at lower pressures may be required to validate this result.

### Thruster model uncertainty

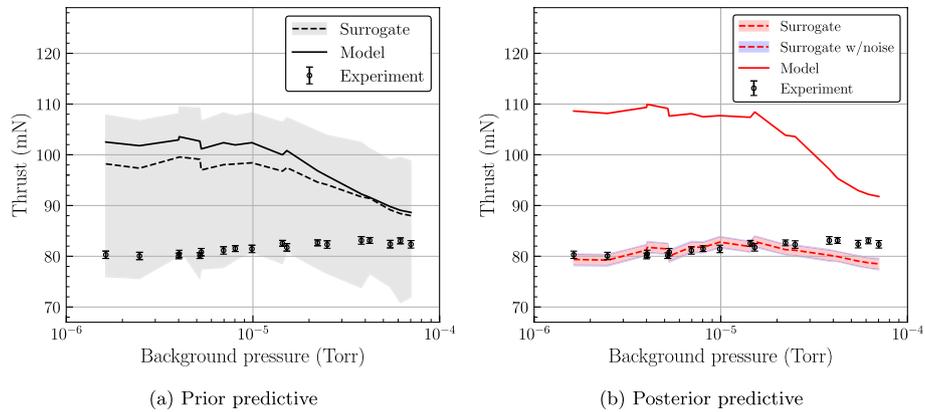
We next evaluate the thruster model predictions on the system-level outputs of thrust, discharge current, and axial ion velocity. Figures 5 and 6 respectively show thrust and discharge current against background pressure for the facility effect characterization test data from Ref. [33]. The model uses all operating conditions from the experiment, including a decreasing anode mass flow rate for increasing back-pressure that was effected in the experiment to maintain a constant discharge current of 4.5 A. We assume a 10% relative uncertainty on the discharge current in the absence of any reported values.

We make four observations regarding the thrust and discharge current predictions in Figs. 5 and 6:

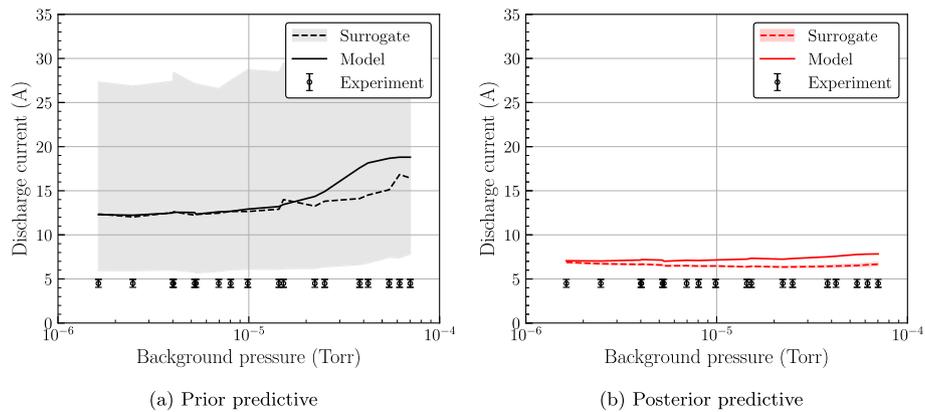
**Observation 1:** *The model predicts a decreasing trend of thrust with pressure, which disagrees with the experimental trend.* — The experimental data was collected over various pressure such that the discharge power remained approximately constant [33], and so the observed “artificial” increase in thrust at elevated pressures is likely due to ingestion and ionization of background neutrals in the acceleration region. However,



**Fig. 4** Comparison of model predictions to experimental cathode coupling voltage from Ref. [16] over varying background pressures. The uncertainty is significantly reduced under the posterior predictive and the model fits the experimental trend well



**Fig. 5** Comparison of model predictions to experimental thrust from Ref. [33] over varying background pressures. The uncertainty is significantly reduced under the posterior predictive. The surrogate fits the data well while the true model over-predicts the data and exhibits the wrong trend compared with experiment



**Fig. 6** Comparison of model predictions to experimental discharge current from Ref. [33] over varying background pressures. The uncertainty is significantly reduced under the posterior predictive. The surrogate and true model predictions fit the experiment better under the posterior predictive but still exhibit a large discrepancy

the thermal ingestion of neutrals predicted by the model from Eq. (4) is not sufficient to make up for the decreasing anode flow rates used in the experiment, resulting in a decreasing trend of thrust with pressure. We also observed a decrease in current efficiency  $\eta_c$  from 66% at the lowest pressure case to 51% at the highest pressure case, primarily driven by an increase in discharge current  $I_d$  as shown in Fig. 6. This result can be understood in context of the shifted-Bohm anomalous transport in Eq. (5), where larger regions of the discharge exhibit higher electron conductivity at higher pressures, and so a larger discharge current is required to maintain the same beam current. Worse electron confinement at higher pressures also corresponds to lower ionization rates and lower thrust.

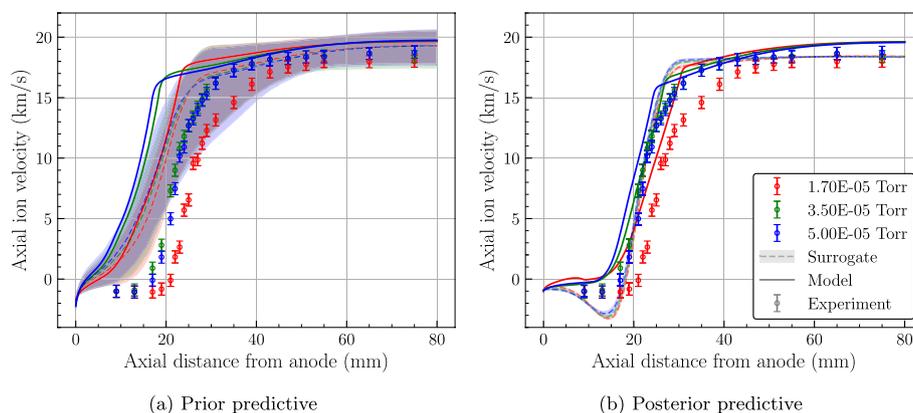
**Observation 2:** *Uncertainty under the posterior is significantly reduced from the prior for both thrust and discharge current, and the remaining uncertainty is mostly due to aleatoric sources.* — This result agrees with the tight bounds on the posterior and implies that uncertainty cannot be reduced further by collecting more data; any remaining

discrepancy between the model and data must be due to error in the model form itself. This result also illustrates that surrogate predictions are overconfident since the uncertainty bounds do not account for model form error nor the error in the surrogate itself. It would be critical to account for these other sources of uncertainty outside of our current validation context.

**Observation 3:** *We have learned the best possible fit for the surrogate, but not necessarily for the true model, i.e. we sacrifice accuracy in the optimization of the parameters for computational feasibility.* — This limitation is evident in Fig. 5, where the surrogate matches the data well while the true model predictions have actually moved in the *opposite* direction. We show some tests in Appendix B where surrogate predictions have high  $\mathcal{O}(25\%)$  error and can even move in opposite directions as the true model for some parameter regimes. Specifically, we found that anomalous coefficients of  $a_1 < -2$  cause a discontinuous, unstable prediction of thrust. Since the posterior lies in this regime, the best-fit parameters learned by the surrogate *do not* necessarily correspond to improvements in the quality of the true model. For discharge current however, the surrogate error is lower and so both surrogate and model predictions improve under the posterior.

**Observation 4:** *The posterior discharge current predictions appear to be as close to the data as was possible under the prior, yet they are still >50% higher than the experiment.* — This result may indicate that our prior was over-constrained such that no set of parameters was capable of reproducing the experimental 4.5 A. However, since the prior predictive of thrust does allow capturing the experimental data, it seems possible that over-prediction of discharge current is a limitation intrinsic to the model itself. We suspect that a three-zone anomalous transport model would allow slower electron transport in the peak magnetic field region and thereby reduce discharge current to more realistic values.

Next, we show a comparison to ion velocity data for the thruster model. Figure 7 compares surrogate and full model predictions to laser-induced fluorescence (LIF) measurements of the axial ion velocity profile from Ref. [21]. Much work in closure modeling tries to match such ion velocity profiles to validate anomalous transport models [20]. We also



**Fig. 7** Comparison of model predictions to experimental ion velocity profiles from Ref. [21] over varying background pressures. The uncertainty is significantly reduced under the posterior predictive. The surrogate and true model predictions fit the experiment better under the posterior predictive but do not fully capture the upstream shift in acceleration region with increasing pressure

hope to determine here if our shifted-Bohm transport model is sufficient for capturing trends with pressure.

One of the main goals of Ref. [21] was to characterize the upstream shift in the acceleration region due to increasing background pressure, as apparent in the shift of the steepest section of the LIF data in Fig. 7. Our shifted-Bohm transport model in Eq. (5) sought to replicate this behavior by applying a logistic upstream shift in the anomalous transport that increased up to a limit with increasing pressure. While this shift in the model is apparent in Fig. 7, it is not sufficient to match the experimentally-observed shift. In addition, the predicted velocity profiles are not steep enough to match the experimental acceleration region. Uncertainty is significantly reduced under the posterior similar to thrust and discharge current shown previously. However, the remaining mismatch between the model and data cannot be accounted for by input uncertainty and so must be due to the model form itself.

It is interesting to note the physical correspondence of the learned shifted-Bohm parameters in Table 3 to the results in Fig. 7. Higher values of  $\Delta z$ , for example, correspond to a greater shift in ion velocity profile by proxy of a greater shift in anomalous collision frequency. Since the experimental shift is greater than the predicted shift, the calibration routine wanted  $\Delta z$  as large as possible, up to the maximum prior bound of  $\Delta z = 0.4$ . The location of the acceleration zone, however, is further downstream than the prior uncertainty could account for. It is possible that allowing  $z_0 > 0$  may have better captured this experimental result, as the calibration routine pushed  $z_0 \rightarrow 0$ . The upstream shift in acceleration zone also seems to taper at  $P_B = 50 \mu\text{Torr}$ , which agrees with the calibrated parameter  $p_0 \sim 57 \mu\text{Torr}$ . Under the current parameterization, it seems unlikely that the 1d fluid model can simultaneously capture both the location and steepness of the acceleration region, as well as the slower asymptote of ion velocity outside the discharge channel.

Figure 7 also shows the first test of our dimension reduction surrogate methods. The surrogate ion velocity profiles are reconstructed from predictions of the latent coefficients  $\tilde{u}_{ion,0} \dots \tilde{u}_{ion,3}$ . Formally, the predicted profiles  $\hat{u}_{ion}(z)$  labeled as ‘‘Surrogate’’ in Fig. 7 are given by:

$$\hat{u}_{ion}(z) = \tilde{V}f_{\mathcal{J}}(\mathbf{x}) \quad (13)$$

for a given system input  $\mathbf{x}$  and the projection matrix  $\tilde{V}$  (see Appendix B for details). If we let  $u_{ion}(z) = f(\mathbf{x})$  be the true model prediction, then the total error in the surrogate prediction is given by:

$$\epsilon_{tot}(\mathbf{x}) = \hat{u}_{ion} - u_{ion} \quad (14)$$

$$= \tilde{V}f_{\mathcal{J}}(\mathbf{x}) - u_{ion} \quad (15)$$

$$= \tilde{V}(\epsilon_{int} + \tilde{V}^T u_{ion}) - u_{ion}, \quad (16)$$

where  $\epsilon_{int} = f_{\mathcal{J}}(\mathbf{x}) - \tilde{V}^T u_{ion} \rightarrow$  interpolation error,

$$\epsilon_{tot}(\mathbf{x}) = \tilde{V}\epsilon_{int} + \epsilon_{rec}, \quad (17)$$

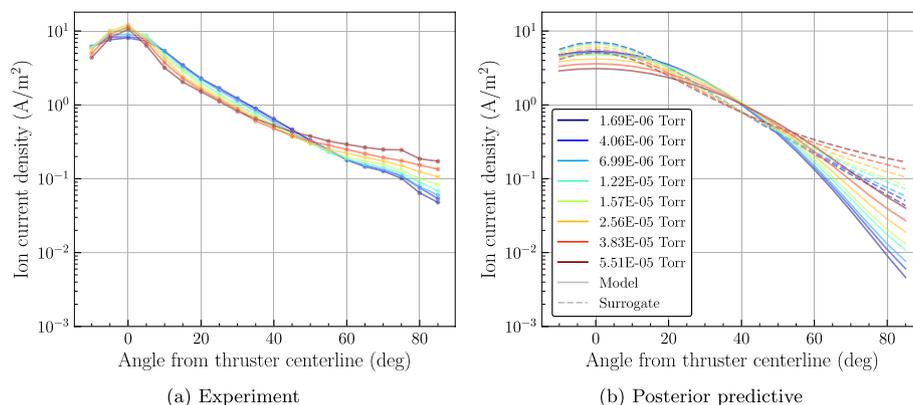
where  $\epsilon_{rec} = \tilde{V}\tilde{V}^T u_{ion} - u_{ion} \rightarrow$  reconstruction error.

The result of Eq. (17) is that the total error of the reconstructed fields can be decomposed into **1)** reconstruction error that results from truncating the SVD projection  $\tilde{V}$ , and **2)** interpolation error that results from inaccurate surrogate estimations of latent coefficients. In our experience, we found the reconstruction error to be high in steep regions of the ion velocity profile, often in the form of high frequency SVD modes that overshoot sharp transitions. This error is amplified when the surrogate approximation is poor, as was typically the case for the ion velocity latent coefficients. These errors are clearly evident in the nonphysical overshoots of the acceleration region in Fig. 7. Even so, we were able to get a cheap estimate of the full ion velocity profile using the surrogate in the latent space, enabling the calibration and UQ results we have shown. We leave it as future work to further investigate trade-offs in accuracy when employing dimension reduction within the surrogate framework.

### Plume model uncertainty

We now turn to the last system-level output of the plume model: ion current density. We encode in the plume model the experimentally-observed effect at elevated facility pressures of decreasing divergence angle [33]. This manifests as a more peaked current density along thruster centerline and wider “wings” of current density at larger angles due to scattering and charge-exchange collisions. Figure 8 shows this effect in the current density profiles in both the model and experiment for increasing facility pressure. We note that the plume model itself is analytical and does not require a surrogate; however, it is dependent on the total beam current  $I_b$  predicted by the thruster and so it relies on the thruster surrogate for UQ predictions. To prevent clutter, we do not plot error bars or uncertainty bounds in Fig. 8.

While we do observe agreement in the wings of the plume (angle > 50 deg), we see a decreasing trend in peak current density on centerline in the model instead of the expected increase in experiments. This is likely due to an over-allocation of current by the model in the middle (15-50 deg) range, which by current conservation forces the peak density to be lower. This actually has the effect of *increasing* beam divergence for



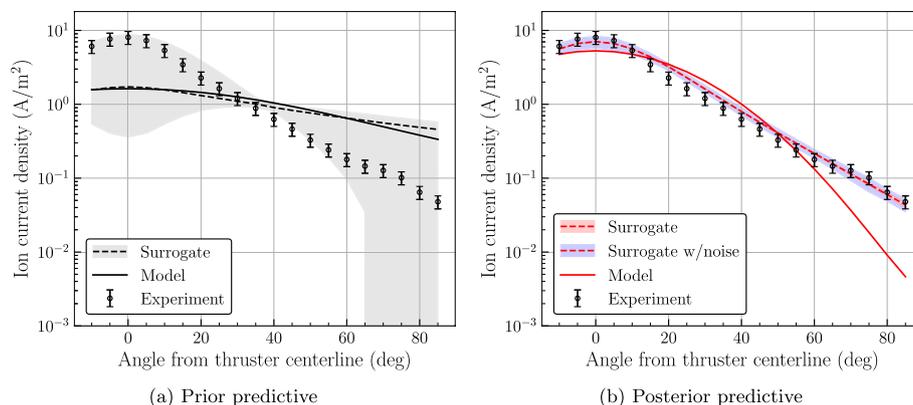
**Fig. 8** Comparison of posterior model predictions to experimental ion current density profiles from Ref. [33] over varying background pressures. The surrogate matches the experimental trend in the “wings” (> 50 deg) of the plume, but neither the surrogate nor true model capture the increasing centerline current density with increasing pressure

increasing pressure, which is opposite the expected experimental trend. We leave the details of the plume model in Appendix A, but we note the surprising calibration results of  $c_0 > 0.5$  and  $c_2 > 0$ , which respectively indicate more allocation of ion current to scattering effects and increasing divergence angle for increasing pressure. These unfavorable results are likely due to numerical artifacts in the calibration procedure, either due to high surrogate errors or too much weight placed on matching the wings of the plume rather than the centerline. It is also possible that this empirical plume model cannot faithfully account for the experimental trends under its current parameterization, since the best-fit parameters are opposite those expected from experiment.

To gain a better understanding of how uncertainty affects the plume model predictions, we show a closer comparison between the model and experiment at  $P_B = 1.69\text{E-}6$  Torr in Fig. 9.

We first observe, similar to the thruster model results, that uncertainty is significantly reduced under the posterior and that the best-fit parameters learned by the surrogate do not necessarily mean a best-fit for the true model. Although the true model does improve from prior to posterior, it does not converge as well as the surrogate to the experimental data, especially at large angles and near thruster centerline. We suspect that these regions, similar to the acceleration region in the ion velocity profiles, exhibit high reconstruction error due to sharp changes in ion current density (note the log scale of Fig. 9). As surrogate error is higher in these regions, the best-fit parameters learned by the surrogate correspond less to improvement in the true model.

It is also interesting to note the apparent presence of an “inflection” angle in Fig. 8 in both the experiment (around 50 deg) and in the model (around 40 deg) where the ion current density is roughly constant for varying pressures (i.e.  $dj_{ion}/dP_B \approx 0$ ). To the left of the inflection point, greater facility pressure compresses the ion current along centerline (in experiments), while it spreads ion current in the wings of the plume to the right of the inflection point. Other than the current conservation that must occur as ion current is shifted between regions of the plume, it is not immediately clear from experiments why such an inflection point develops at a single angular location in the plume. It is also not



**Fig. 9** Comparison of model predictions to experimental ion current density from Ref. [33] for a background pressure of  $P_B = 1.69\text{E-}6$  Torr. The uncertainty is significantly reduced under the posterior predictive. The surrogate matches the experiment well under the posterior but the true model exhibits large discrepancy near thruster centerline and at large angles ( $> 50$  deg)

intuitive from the plume model equations why this same phenomenon occurs, albeit at a different location than experiment. This effect is further emphasized by the narrow uncertainty band at the inflection point even due to all other model parameters, as shown in the prior predictive in Fig. 9. Despite the complex coupling of pressure with all the other model uncertainties, the model predicts a roughly constant current density at the inflection point. Along with the discrepancies previously mentioned, this may motivate a better parameterization or model form for the plume component.

**Training performance summary**

We conclude this section by summarizing the model/surrogate performance before and after calibration on the training data (Figs. 4, 5, 6, 7, 8 and 9). Table 4 summarizes the distribution of surrogate relative  $L_2$  error for each QoI with the mean  $\mu$  and standard deviation  $\sigma$  over  $N = 1000$  samples of the uncertain inputs  $\{\mathbf{x}^{(i)}\}_{i=1}^N$ , given by:

$$\mu = \frac{1}{N} \sum_{i=1}^N E(\mathbf{y}_e, f_{\mathcal{J}}(\mathbf{x}^{(i)})), \tag{18}$$

$$\sigma = \sqrt{\frac{\sum_{i=1}^N (E(\mathbf{y}_e, f_{\mathcal{J}}(\mathbf{x}^{(i)})) - \mu)^2}{N}}, \tag{19}$$

where  $E(\cdot, \cdot)$  is computed from Eq. (12) between the experimental data  $\mathbf{y}_e$  and the surrogate predictions  $f_{\mathcal{J}}(\cdot)$ , respectively. We compare the average  $L_2$  error to the characteristic experimental noise  $\xi$  to assess goodness of fit, with values  $\mu/\xi \sim \mathcal{O}(1)$  indicating good agreement (i.e. model error is on the same order as measurement error). In general, a reduction in  $\mu$  indicates improved accuracy and a reduction in  $\sigma$  indicates reduced uncertainty. We also include the point error estimate  $\mu_{50} = E(\mathbf{y}_e, f_{\mathcal{J}}(\mathbf{x}_{50}))$  for the surrogate at median

**Table 4** Relative  $L_2$  error of the model and surrogate with experimental SPT-100 training data under the prior and posterior distributions for several quantities of interest (QoIs)

QoI	Distribution	$\xi$ (%)	$\epsilon_{50}$	Surrogate $L_2$ (%)				Model $L_2$ (%)	
				$\mu_{50}$	$\mu$	$\sigma$	$\mu/\xi$	$\mu_{50}$	$\mu_{50}/\xi$
$V_{cc}$ (V)	Prior	1	-	6	54	6	54	-	-
	Posterior	1	-	2	1.3	0.1	1.3	-	-
$T$ (mN)	Prior	1	1	21	19	2	19	22	22
	Posterior	1	23	2.5	2.6	0.2	2.6	29	29
$I_d$ (A)	Prior	10	10	207	290	40	29	230	23
	Posterior	10	11	45	45	0.3	4.5	63	6.3
$u_{ion}$ (ms <sup>-1</sup> )	Prior	5	7	44	35	5	7	50	10
	Posterior	5	13	21	21	0.2	4.2	17	3.4
$j_{ion}$ (Am <sup>-2</sup> )	Prior	20	17	77	79	6	4	80	4
	Posterior	20	29	33	33	0.3	1.7	49	2.5

Surrogate  $L_2$  errors are summarized by the mean  $\mu$  and standard deviation  $\sigma$  over  $N = 1000$  samples. The point estimate  $\mu_{50}$  is computed at median parameter values only, and  $\epsilon_{50}$  gives the corresponding approximation error between surrogate and model. The ratio of mean  $L_2$  error ( $\mu$ ) to relative experimental noise ( $\xi$ ) indicates the degree of fit with experiment, with  $\mu/\xi \sim \mathcal{O}(1)$  indicating good agreement and  $\mu/\xi > 1$  indicating larger mismatch

parameter values  $\mathbf{x}_{50}$  (see Table 1 for the prior and Table 3 for the posterior), with the corresponding approximation error  $\epsilon_{50} = E(f(\mathbf{x}_{50}), f_{\mathcal{J}}(\mathbf{x}_{50}))$  between model and surrogate at the median predictions. Note that the prediction of cathode coupling voltage ( $V_{cc}$ ) is exact and so  $\epsilon_{50}$  is not included for this QoI. Also note that we only compute the true model at the median parameter values and provide the point error estimate  $\mu_{50} = E(\mathbf{y}_e, f(\mathbf{x}_{50}))$ .

We make three observations regarding the training performance summary in Table 4:

**Observation 1:** *Uncertainty under the posterior is significantly reduced from the prior for all QoIs.* — This fact was already observed in Figs. 4, 5, 6, 7, 8 and 9 and is confirmed by the posterior standard deviation in error of  $\sigma < 0.3\%$  for all QoIs in Table 4. The low remaining uncertainty along with  $\mu/\xi > 1$  for most QoIs again indicates that model form error is responsible for the remaining discrepancy with experiment.

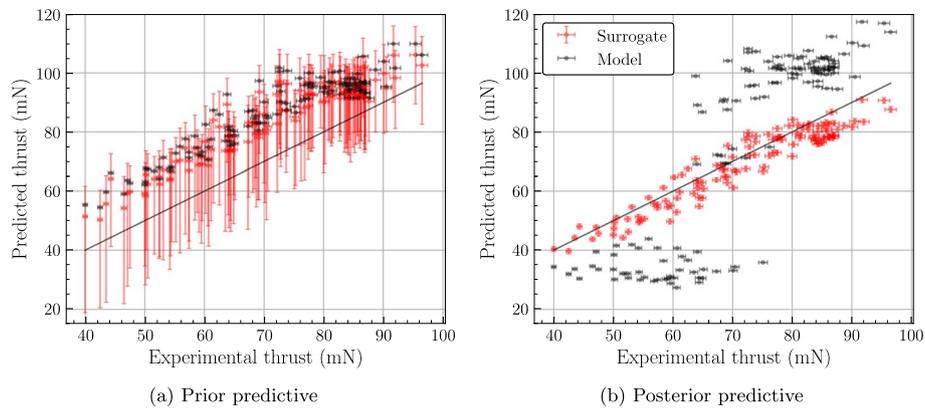
**Observation 2:** *Surrogate predictions under the posterior are significantly improved from the prior baseline.* — For example, the average surrogate  $L_2$  error improves for discharge current by an order of magnitude, from  $\mu = 290\%$  to 45%. Recall that the prior baseline for this model is based on best available numerical or experimental evidence; this is a significant improvement in accuracy despite the relatively high remaining error. The accuracy in cathode coupling voltage ( $V_{cc}$ ) even improves enough that the model error is on the order of experimental error (i.e.  $\mu/\xi \sim \mathcal{O}(1)$ ), indicating that further calibration for this QoI will likely not improve accuracy further. Except for thrust, improvement in surrogate predictions was accompanied by similar improvement in true model predictions.

**Observation 3:** *Surrogate approximation error of the true model limits observed improvement in the true model during calibration.* — This fact was previously observed in Fig. 5 where the true model prediction of thrust actually degrades under the posterior due to high surrogate approximation error ( $\epsilon_{50}$ ). Since the calibration is performed with the surrogate, the surrogate-experiment error naturally improves for all QoIs, ideally inducing the model-experiment error to improve as well. However, when surrogate-model approximation error is high, then true model improvement may be limited as was observed with thrust, where approximation error increased from  $\epsilon_{50} = 1\%$  to 23% under the posterior causing true model performance to degrade.

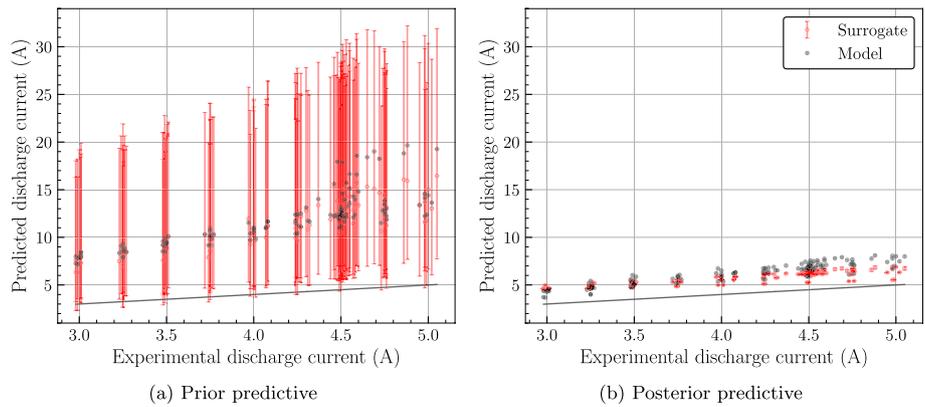
### Model extrapolation

In the last section, we extensively quantified the impacts of uncertainty on model predictions and compared to the training data. We now test the model on a large, independent dataset that was not used during calibration. The dataset provides SPT-100 thrust and discharge current measurements at many new operating conditions (such as off-nominal voltages) and is described in Dataset 5 of Appendix C. Figures 10 and 11 summarize the model performance on this test set by plotting the predicted thrust and discharge current, respectively, against the measured experimental values. Points closer to the  $y = x$  line indicate better agreement between the model and data. Horizontal error bars indicate experimental noise and vertical error bars indicate the 5th and 95th percentiles of  $N = 1000$  surrogate predictions. True model predictions only use the median parameter values where applicable.

The performance of both the model and surrogate on the test set is summarized in Table 5 using the relative  $L_2$  error metric. We again quantify the distribution of error with the mean  $\mu$  and standard deviation  $\sigma$  over  $N = 1000$  samples. As seen in Fig. 10 and as previously observed in the [Quantifying uncertainty](#) section, the high error between



**Fig. 10** Validation test set comparison of model predictions to experimental thrust from Ref. [34] over a variety of operating conditions, with points closer to the  $y = x$  line indicating a closer agreement between model and experiment. The uncertainty is significantly reduced under the posterior predictive. The surrogate fits the experiment better under the posterior predictive but the true model still exhibits a large discrepancy



**Fig. 11** Validation test set comparison of model predictions to experimental discharge current from Ref. [34] over a variety of operating conditions, with points closer to the  $y = x$  line indicating a closer agreement between model and experiment. The uncertainty is significantly reduced under the posterior predictive. The surrogate and model fit the experiment better under the posterior predictive but still exhibit a large discrepancy

surrogate and model ( $\epsilon_{50}$ ) in the posterior regime causes the true model thrust predictions to have higher error under the posterior vs. the prior. However, the calibrated surrogate improves significantly from the prior baseline and matches the experimental thrust data well with an average relative  $L_2$  of  $\mu < 10\%$  on the test set.

For the discharge current, we observe a significant improvement for both the model and surrogate predictions in Fig. 11. The model improves in-line with the surrogate as intended due to relatively low surrogate approximation error in both the prior and posterior regime. However, as observed previously, there is still a large discrepancy between the model and discharge current data, which points to limitations intrinsic in the model itself rather than the surrogate or calibration procedure.

Overall, the results from this section show good agreement of our globally-calibrated model over a wide range of operating conditions for the SPT-100 that lie outside the

**Table 5** Relative  $L_2$  error of the model and surrogate with experimental SPT-100 validation data from Ref. [34] under the prior and posterior distributions for thrust ( $T$ ) and discharge current ( $I_d$ )

QoI	Distribution	$\xi$ (%)	$\epsilon_{50}$	Surrogate $L_2$ (%)				Model $L_2$ (%)	
				$\mu_{50}$	$\mu$	$\sigma$	$\mu/\xi$	$\mu_{50}$	$\mu_{50}/\xi$
$T$ (mN)	Prior	1	1	21	20	1	20	22	22
	Posterior	1	27	7	7	0.1	7	30	30
$I_d$ (A)	Prior	10	7	183	270	15	27	193	19
	Posterior	10	11	40	40	0.1	4	53	5.3

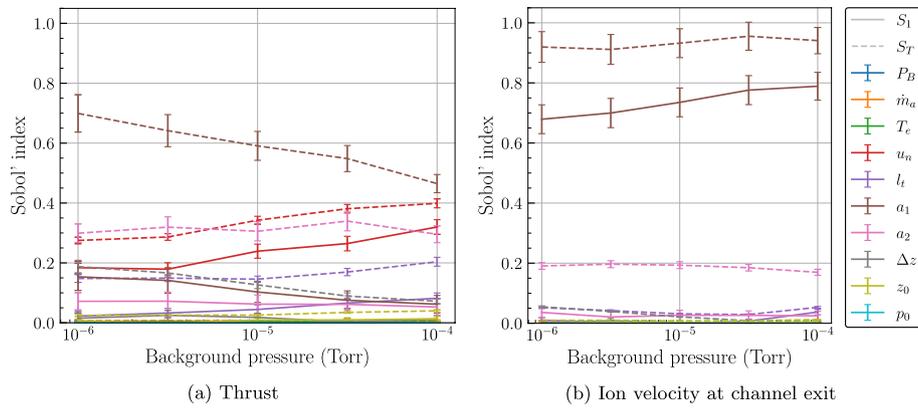
Surrogate  $L_2$  errors are summarized by the mean  $\mu$  and standard deviation  $\sigma$  over  $N = 1000$  samples. The point estimate  $\mu_{50}$  is computed at median parameter values only, and  $\epsilon_{50}$  gives the corresponding approximation error between surrogate and model. The ratio of mean  $L_2$  error ( $\mu$ ) to relative experimental noise ( $\xi$ ) indicates the degree of fit with experiment, with  $\mu/\xi \sim \mathcal{O}(1)$  indicating good agreement and  $\mu/\xi > 1$  indicating larger mismatch

training data. The results also show substantial improvement of the calibrated surrogate from the prior baseline, with an order of magnitude improvement for the QoIs shown and a large reduction in uncertainty. However, due to high posterior surrogate approximation error ( $\epsilon_{50}$ ) for thrust and model inadequacy for discharge current, the true model still has significant error compared to experiment ( $> 30\%$  for thrust and  $> 50\%$  for discharge current) and the remaining uncertainty is not sufficient to account for the discrepancy.

### Sensitivity analysis

Having reduced model parameter uncertainty significantly and compared results to data, we have a better understanding of the strengths and weaknesses of the current model parameterization. We now take a preliminary step in improving future versions of the model using global Sobol' sensitivity analysis. As described in the [Forward propagation and sensitivity analysis](#) section, the goal of this sensitivity analysis is to identify the relative importance of model inputs and to remove those with negligible impact on model outputs from future consideration. We consider all inputs over their full prior domains in Table 1. Figure 12 shows the first- and total-order Sobol' indices for all uncertain inputs of the thruster model for varying background pressures. We show the Sobol' indices only for thrust predictions and the ion velocity at channel exit. The nomenclature in the legend is given in Table 1. An expected result in Fig. 12 is the high contribution of anomalous transport coefficients  $a_1$  and  $a_2$  in the operation of the thruster, with  $S_T > S_1$  for these inputs likely indicating high coupling between them. We observe that uncertainty in  $a_1$  is especially dominant in predicting ion velocity, with increasing effect at higher pressures ( $S_T > 0.8$ ). The fact that the shifted-Bohm parameters ( $\Delta z, z_0, p_0$ ) have minimal impact supports the conclusion that the model is likely over-parameterized. The low contribution of uncertainty in operating conditions  $P_B$  and  $\dot{m}_a$  also suggests these uncertainties can be safely ignored.

A surprising result is the high contribution of anode injection neutral velocity  $u_n$  in the thrust predictions, with increasing impact at higher pressures. This effect can be understood in the context of how ingested neutrals enter the simulation via:



**Fig. 12** Sobol' indices over varying background pressures for thrust and ion velocity at channel exit. Dashed lines indicate total-order ( $S_T$ ) indices and solid lines indicate first-order ( $S_1$ ) indices. Error bars give 95% Monte Carlo confidence intervals. The anomalous transport coefficients ( $a_1, a_2$ ) are shown to have the largest impact on variance in these model outputs

$$\Gamma_{\text{anode, tot}} = \Gamma_{\text{prop}} + \Gamma_{\text{ingest}} \tag{20}$$

$$n_{n, \text{tot}} u_n = \frac{\dot{m}_a}{mA_e u_n} u_n + \frac{n_{n, B}}{4} \sqrt{\frac{8k_B T_B}{\pi m}} \tag{21}$$

$$n_{n, \text{tot}} = \frac{\dot{m}_a}{mA_e u_n} + \frac{P_B}{4u_n} \sqrt{\frac{8}{\pi m k_B T_B}}, \tag{22}$$

where  $n_{n, \text{tot}}$  is the total neutral density at the anode. Equation (22) suggests an increase in total available propellant for higher background pressure  $P_B$  and slower injection velocity  $u_n$ . At higher pressures, the model is more sensitive to changes in  $u_n$ , as this changes the total amount of available propellant and thereby the total thrust predicted by the model.

We additionally performed a Sobol' analysis for the cathode and plume parameters and included these results in Appendix E.

### Discussion

In this work, we have used a multifidelity framework to build a surrogate of a multi-disciplinary Hall thruster model, we have used Bayesian inference to calibrate model parameters on experimental data, we have used forward uncertainty propagation to validate the model against experimental data, and we have used sensitivity analysis to identify the most critical components of the model. We now discuss the results of these analyses in the context of the overarching goal of this work: Hall thruster model improvement. We first summarize the core findings of this work, then we make suggestions for future model improvement, and finally we discuss several ways to improve the framework itself.

## Core findings

1. **A multifidelity, multidisciplinary (MF/MD) treatment of the Hall thruster system model provides significant computational gains compared to a monolithic, black-box treatment.** The UQ analyses applied in this work are infeasible for expensive models and so we used surrogate methods in an MF/MD setting to maximize computational efficiency. We found that a more accurate surrogate can be obtained for the Hall thruster system model for a given computational cost by applying the multifidelity MISC method [15] rather than using a single high-fidelity, black-box version of the model. We also expect this relative improvement to increase as more expensive models are included.
2. **The UQ and validation framework has shown some success in improving the Hall thruster model over a prior baseline.** A key benefit of the proposed framework is thorough accounting for sources of uncertainty and clear validation against experimental data, which allows us to determine *how*, *where*, and, to some extent, *why* the model is underperforming. For example, even though the prediction of thrust and discharge current on an independent validation set still exhibits large discrepancy with experiment, the accuracy was significantly improved and the minimal remaining uncertainty suggested issues with the model form itself. In addition, assuming sufficiently large training bounds, the surrogate can make confidence-bounded predictions at unseen operating conditions fairly well, such as the 7% error observed on the thrust validation set. This accuracy can potentially be improved further by collecting more surrogate training data.
3. **The modular nature of the analysis allows iteration on the most critical components of the model.** For example, we found that the cathode model sufficiently predicts the dependence of cathode coupling voltage on pressure, while the thruster model does not predict correct trends with pressure. The thruster model can therefore be “swapped” out, re-parameterized, or otherwise improved independently without requiring any change to the cathode model. The benefit of this approach becomes more salient as surrogates are constructed for more component models, i.e. changing one component model will not require reconstructing surrogates for any other component. The framework readily generalizes to more complicated thruster-facility interactions and can iteratively incorporate these improvements without losing any prior progress. In addition, the framework itself suggests new paths for model improvement as we summarize next.

## Suggestions for model improvement

1. **Increase the physical fidelity of the thruster and plume models.** Our attempt at an integrated, multidisciplinary model in this study uses a 1d fluid model for the thruster discharge and a semi-empirical model of the ion current density in the plume. This choice was made for simplicity to develop and test the UQ framework (even though some of the relevant physics in Hall thrusters is known to be 2d or even 3d in nature [3]). However, after calibrating the model to fit experimental data, we

found that epistemic uncertainty in the model parameters was reduced to essentially zero, and the remaining uncertainty was not sufficient to explain the discrepancy between the model and data. Specifically, predictions of thrust and discharge current on an independent validation set were  $> 30\%$  off from experiment. Therefore, we suggest improving the model form itself to incorporate more of the relevant physics and to better fit the data. This likely means upgrading to a 2d thruster model [35] and to a more physics-based plume model (the modularity of our MD approach makes this task feasible). The problem of anomalous electron transport will still be present in a 2d fluid code, and so a thorough UQ analysis will still be required to assess confidence in predictions and goodness of fit. Indeed, this is a primary motivation of the approach taken in this study. The incorporation of a 2d code also opens some interesting opportunities for the multifidelity methods used in this study.

- 2. Refine the dependence on facility effects.** Background pressure is currently the only facility effect considered by the model. The background pressure enters the model in four empirically-guided ways: by modulating the cathode coupling voltage, by increasing available propellant through thermal ingestion, by shifting the anomalous transport upstream, and by altering the plume divergence. We found that the dependence of cathode coupling voltage on pressure is sufficiently captured by the model, but the other three concerns are lacking. Namely, the dependence of thrust on pressure is not captured by the thermal ingestion model, the upstream shift in the acceleration region is not captured by the shifted anomalous transport, and the decreasing plume divergence is not captured by the empirical plume model. Further experimental and numerical investigation into these effects can refine the model's dependence on pressure. For example, existing studies also find the thermal ingestion model to be insufficient and instead use a background flow model [36]. There is also the open question of how facility effects *other* than pressure affect performance at higher powers, such as back-sputtering of carbon from the chamber walls or electrical coupling of the beam to the walls [9]. We leave it as a task for future iterations to bring these interactions into the model as well.
- 3. Focus parameterization on anomalous transport.** It is already well-known that anomalous electron transport dominates the behavior of Hall thrusters. Our sensitivity analysis confirms that the choice of anomalous transport coefficients outweighs most other uncertainties present in the system. Future refinements should focus effort on this aspect of the model. For example, recent advancements in experimentally measuring anomalous transport can help guide future modeling decisions [37]. One change that may have a large impact would be transitioning to a three-region transport model, which would significantly improve the ability of the model to fit the data [19] at the cost of introducing more calibration parameters.

### Suggestions for framework improvement

- 1. Refine the dimension reduction technique.** We were enabled in this work by dimension reduction to include high-dimensional field quantities in the surrogate analyses (i.e. the Bayesian inference, uncertainty propagation, and sensitivity analysis). We chose lin-

ear projection for simplicity as the compression technique, but we encountered several limitations, including high frequency oscillatory modes and high reconstruction error that was amplified by surrogate approximation error. Additionally, this method required a relatively high offline cost to construct the original data matrix. Future iterations can explore nonlinear methods such as autoencoders and iterative construction of the latent-space projection [38]. The adaptive surrogate training procedure can also be augmented to optimally balance reconstruction errors with interpolation errors during training.

2. **Incorporate time-resolved analysis.** In the present work, we are limited to studying the *time-averaged* input-output behavior of the multidisciplinary model, even though the dynamics of Hall thrusters are highly oscillatory and time-dependent. Efficiently coupling distinct component models in a time-resolved manner can be incorporated in future iterations of the framework presented here.
3. **Account for model form uncertainty and reduce surrogate error.** While we thoroughly assessed uncertainty from model inputs in this study, the uncertainty bounds on our model predictions are overly-conservative as they do not account for uncertainty in the form of the model itself. This is especially important as the 1d thruster model is largely simplified from the real physical system, which is known to exhibit 2d and 3d phenomena. Quantifying this uncertainty can be considered in future iterations. For now, any remaining discrepancy between our models and the data is likely attributed to model form uncertainty and so we intend to take steps in increasing physical model fidelity in future work. Similarly, the discrepancy between the models and their *surrogate* approximations is not accounted for in the uncertainty bounds. This error can likely be mitigated by gathering more training data and using locally adaptive surrogate methods to fit non-smooth regions of the model response. Surrogate methods that have built-in probabilistic predictions (like Gaussian processes) can also be used to account for this uncertainty.

## Conclusion

In this work, we have implemented an iterative framework for improving Hall thruster models by applying uncertainty quantification methods on a multidisciplinary system. We have shown with a simple feedforward coupling between a cathode model, a 1d fluid thruster code, and a plume model the impacts of uncertainty on model predictions both before and after Bayesian calibration to experimental data. We have extensively compared the model to experimental SPT-100 data, especially in the context of anomalous electron transport and facility effects which limit the practical use of these models for Hall thruster design, test, and optimization. By exercising the proposed framework, we demonstrated significant improvement over a prior baseline for some quantities of interest. Ultimately however, we found that remaining uncertainties in the model were not sufficient to account for discrepancy with experimental data. This highlights that model form error is the dominant factor limiting our model's accuracy, and that future work should focus on refining the involved physics.

We were also enabled in this work by an existing multifidelity surrogate method to reduce the computational expense of the uncertainty analysis. We improved upon this method to include approximation of high-dimensional spatial quantities using data compression techniques. Future work can additionally include time-resolved analysis and incorporate reconstruction accuracy into the adaptive training procedure.

### Appendix A: Plume model details

We detail the semi-empirical plume model first described in Ref. [22]. All model inputs described here correspond to those outlined in Table 1 except where otherwise noted. The model predicts the ion current density  $j_{\text{ion}}$  (A/m<sup>2</sup>) in the thruster plume as a superposition of the main beam, an ion-neutral scattering beam, and a charge-exchange beam (Eq. (7) copied here for convenience):

$$j_{\text{ion}} = j_{\text{beam}} + j_{\text{scatter}} + j_{\text{cex}}. \quad (23)$$

The main and scattered beams are assumed to be normally-distributed about the thruster centerline from angle  $\gamma = -90$  to  $90$  deg and expand radially with an inverse square dependence on distance  $r_p$  from the thruster exit plane:

$$j_{\text{beam}} = \frac{I_B}{r_p^2} \left( A_1 \exp \left[ - \left( \frac{\gamma}{\gamma_1} \right)^2 \right] \right), \quad (24)$$

$$j_{\text{scatter}} = \frac{I_B}{r_p^2} \left( A_2 \exp \left[ - \left( \frac{\gamma}{\gamma_2} \right)^2 \right] \right), \quad (25)$$

where  $A_1, A_2$  and  $\gamma_1, \gamma_2$  are characteristic magnitudes and divergence angles of each beam respectively. The charge-exchange beam is assumed to expand spherically with no angular dependence:

$$j_{\text{cex}} = \frac{I_{\text{cex}}}{2\pi r_p^2}, \quad (26)$$

where  $I_B$  and  $I_{\text{cex}}$  are the total currents carried by the main/scattered and charge-exchange beams respectively, given by:

$$I_B = I_b \exp(-r_p n_n \sigma_{\text{cex}}), \quad (27)$$

$$I_{\text{cex}} = I_b (1 - \exp(-r_p n_n \sigma_{\text{cex}})). \quad (28)$$

$I_b$  is the total ion beam current exiting the thruster, as provided by *Hallthruster.jl*. The neutral density (in m<sup>-3</sup>) is given a simple linear dependence on background chamber pressure:  $n_n = c_4 P_B + c_5$ , with model fit coefficients  $c_4 \approx \frac{1}{k_B T} \sim \mathcal{O}(10^{20})$  and  $c_5 \sim \mathcal{O}(10^{18})$ . The charge-exchange collision cross-sectional area for Xenon is estimated from Ref. [39] as  $\sigma_{\text{cex}} \sim \mathcal{O}(55 \text{ \AA}^2)$ . We also assign a linear pressure dependence of the characteristic divergence angles and introduce the appropriate fit coefficients:

$$\gamma_1 = c_2 P_B + c_3, \quad (29)$$

$$\gamma_2 = \frac{1}{c_1} \gamma_1, \quad (30)$$

where  $c_1 \sim \mathcal{U}(0.1, 0.9)$  reflects the intuition that the divergence of the scattered beam should be larger than the main beam. We expect the divergence of the main beam should decrease as background pressure increases [40], i.e. that  $c_2 < 0$ , but we allow  $c_2 > 0$  generally. Finally, we include from Ref. [22] the definitions for the main and scattered beam magnitudes:

$$A_1 = \frac{1 - c_0}{\frac{\pi^{3/2}}{2} \gamma_1 \exp\left(-\left(\frac{\gamma_1}{2}\right)^2\right) \left(2\operatorname{erfi}\left(\frac{\gamma_1}{2}\right) + \operatorname{erfi}\left(\frac{i\pi - \gamma_1^2}{2\gamma_1}\right) - \operatorname{erfi}\left(\frac{i\pi + \gamma_1^2}{2\gamma_1}\right)\right)}, \quad (31)$$

$$A_2 = \frac{c_0}{\frac{\pi^{3/2}}{2} \gamma_2 \exp\left(-\left(\frac{\gamma_2}{2}\right)^2\right) \left(2\operatorname{erfi}\left(\frac{\gamma_2}{2}\right) + \operatorname{erfi}\left(\frac{i\pi - \gamma_2^2}{2\gamma_2}\right) - \operatorname{erfi}\left(\frac{i\pi + \gamma_2^2}{2\gamma_2}\right)\right)}, \quad (32)$$

where the model fit coefficient  $c_0 \sim \mathcal{U}(0, 1)$  determines the ratio between the main and scattered beam magnitudes.

## Appendix B: Surrogate analysis

We provide additional details in this section for the surrogate methodology used in this study.

### Dimension reduction

We expanded upon the original framework in Ref. [15] to consider approximation of high-dimensional spatial field quantities. The ion velocity  $u_{ion}(z)$  and ion current density  $j_{ion}(\gamma)$  outputs from Table 2 are 1d spatially-varying QoIs, where  $z$  is the axial location in the thruster channel (measured from the anode) and  $\gamma$  is the angular location in the plume (measured from thruster centerline). Numerical simulations of these field quantities return the solution “ $u$ ” on a set of  $M$  discretized points of the form  $u \in \mathbb{R}^M$ . Typically,  $M$  is large to ensure PDE convergence on the mesh (especially in  $> 1$  spatial dimensions), so  $u$  is high-dimensional. Training surrogates in high-dimensional space is infeasible due to ballooning computational requirements (i.e. the curse of dimensionality) and so we require a low-dimensional mapping  $\tilde{u} = g(u)$ ,  $\tilde{u} \in \mathbb{R}^r$ ,  $r \ll M$  and train the surrogate in the low-dimensional (or “latent”) space  $\tilde{u}$  instead. We apply ideas from similar studies [27–29] to the current task and obtain the map  $g(u)$  using the singular value decomposition (SVD). First, we form the data matrix  $A \in \mathbb{R}^{N \times M}$  by running  $N$  simulations  $\{\mathbf{y}^{(i)}\}_{i=1}^N = \{f(\mathbf{x}^{(i)})\}_{i=1}^N$ ,  $\mathbf{x}^{(i)} \sim \rho(\mathbf{x})$  and filling the rows of  $A$  with the field quantity extracted from each simulation output:

$$A = N \left\{ \underbrace{\begin{bmatrix} \text{---} & u^{(1)} & \text{---} \\ & \vdots & \\ \text{---} & u^{(i)} & \text{---} \\ & \vdots & \\ \text{---} & u^{(N)} & \text{---} \end{bmatrix}}_M, \quad u^{(i)} \subseteq \mathbf{y}^{(i)}. \right. \quad (33)$$

The data matrix forms a basis for all the variation we expect to see in  $u$  over the input space. We then decompose the matrix according to  $A = U_{N \times N} S_{N \times M} V_{M \times M}^T$  and truncate to the first  $r$  singular values  $\lambda_i$  of the diagonal matrix  $S$  such that the remaining fraction of total variance is above a target threshold  $\epsilon \in [0, 1]$ :

$$\frac{\sum_{i=1}^r \lambda_i^2}{\sum_{i=1}^{\text{rank}(S)} \lambda_i^2} > \epsilon. \quad (34)$$

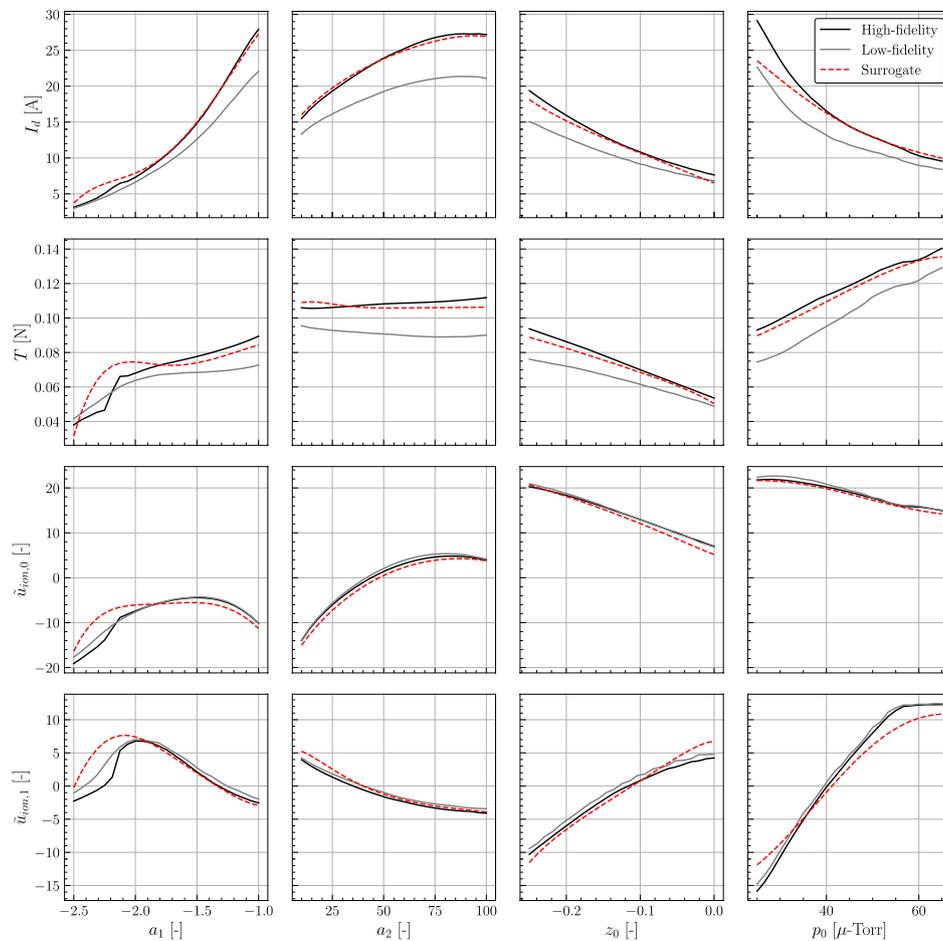
The data matrix is approximated by  $A \approx \tilde{U}_{N \times r} \tilde{S}_{r \times r} \tilde{V}_{r \times M}^T$ , where we have kept only the first  $r$  columns of  $U$  and  $V$ . The matrix  $\tilde{V}^T$  is a linear projection  $\mathbb{R}^M \rightarrow \mathbb{R}^r$  that we use as our low-dimensional mapping  $g(u)$ , that is  $\tilde{u} = \tilde{V}^T u$ . We refer to  $\tilde{u}_i, i = 1 \dots r$  as the latent coefficients of  $u$  and use the inverse mapping  $\tilde{V}$  to reconstruct the full field  $\hat{u} = \tilde{V} \tilde{V}^T u$  (the result  $\hat{u} \neq u$  is due to reconstruction error from the truncation). For example, the quantity  $\tilde{u}_{ion,0}$  is the first latent coefficient obtained by compressing the ion velocity profile via  $\tilde{V}^T u_{ion}(z)$ . The surrogate is trained on the latent coefficients and the full field is reconstructed when making predictions.

In practice, the SVD is better conditioned when the data matrix is first standardized via  $(A - \bar{A})/\sigma(A)$ , where the mean  $\bar{A}$  and standard deviation  $\sigma(A)$  are applied over the samples  $i = 1 \dots N$ . The principal orthogonal decomposition (or principal component analysis) is a congruent method for obtaining the linear projection  $\tilde{V}^T$  to the latent space.

In the current study, we reduced the ion velocity profile  $u_{ion}(z)$  from  $M = 200$  to  $r = 4$  with a target threshold of  $\epsilon = 95\%$  in Eq. (34). With the same threshold, we reduced the current density profile  $j_{ion}(\gamma)$  from  $M = 100$  to  $r = 2$ . This required an offline cost of  $N = 500$  simulations to form the data matrix  $A$  in Eq. (33).

### Additional surrogate performance

In the [Hall thruster model surrogate](#) section, we reported an *average* 10% surrogate accuracy over a test set. To gain some practical/visual intuition for this result, we plot the surrogate prediction against both the lowest-fidelity (0, 0) and highest-fidelity (2, 2) thruster models for several 1d sweeps across the 12d input space in Fig. 13. For practical reasons, we only show the four outputs ( $I_d, T, \tilde{u}_{ion,0}, \tilde{u}_{ion,1}$ ) against four select inputs ( $a_1, a_2, z_0, p_0$ ) arranged in a  $4 \times 4$  grid (see Tables 1 and 2 for details on these variables). The 1d sweeps were performed by fixing the end points of the input variable in each column, randomly selecting values for all other “non-column” inputs, and smoothly stepping from left to right. This implies that four independent 1d sweeps were performed in Fig. 13, one for each column.



**Fig. 13**  $4 \times 4$  grid showing surrogate predictions against low- and high-fidelity model predictions for four select outputs (rows) over four select inputs (columns)

The key takeaway from Fig. 13, while only showing a small portion of the full 12d space, is that the surrogate can be used as a reliable substitute for the high-fidelity model. The surrogate fits the high-fidelity model reasonably well given the sparse training data and high-dimensionality. While the surrogate may have high error at specific locations (such as  $a_1 < -2$ ), *on average* it provides good accuracy at negligible cost as a cheap replacement for the full model.

We note that many engineering simulation QoIs exhibit a smooth response to changing inputs, such as shown in Fig. 13. In these cases, interpolating polynomials or similar methods provide a very efficient approximation to the response surface. However, noisy, mode-hopping, or otherwise discontinuous responses can significantly erode polynomial accuracy; we have seen some of this behavior in our own experiments and leave it as future work to investigate alternative surrogate methods (i.e. neural networks, Gaussian processes, etc.) to handle these cases.

### Appendix C: SPT-100 thruster data

This section is provided for convenience to organize all SPT-100 data used in this study.

#### Thruster details

In the notation of the [Hall thruster system model](#) section, the operating conditions  $\mathbf{d}$  can be divided into test article details like geometry or material properties ( $\mathbf{d}_T$ ) and test environment details like voltage settings, flow rates, etc. ( $\mathbf{d}_E$ ). We work only with the SPT-100 Hall thruster in this study, that is  $\mathbf{d}_T = \text{constant}$ . Here, we group  $\mathbf{d}_T$  into parameters describing the thruster geometry, wall material, propellant, cathode, and magnetic field design, as summarized in Table 6.

**Table 6** SPT-100 thruster details

Group	Parameter	Symbol	Units	Value	Reference
Geometry	Channel length	$L_{ch}$	m	0.025	[41]
Geometry	Inner radius	$r_i$	m	0.035	[41]
Geometry	Outer radius	$r_o$	m	0.05	[41]
Wall	Wall material	-	-	Borosil (BNSiO <sub>2</sub> )	[42]
Propellant	Propellant material	-	-	Xenon	
Cathode	Mount location	-	-	External	
Magnetic field	Peak strength	$B_{max}$	T	0.016	[43]
Magnetic field	Centerline profile	$B_r(z)$	T	Eq. (24) in Ref	[44]

For a given thruster, all parameters in  $\mathbf{d}_T$  are assumed known and constant. The geometry of most axisymmetric, cylindrical Hall thrusters can be simplified to three parameters: the channel length  $L_{ch}$  and inner and outer radii ( $r_i$  and  $r_o$ ). Alternatively, the channel width  $b = r_o - r_i$  and mean channel diameter  $d_{ch} = r_o + r_i$  are frequently used. The wall material and propellant can be parameterized by key material properties such as molar mass, specific heats, or properties related to secondary electron emission. The cathode can be parameterized by its mount location (external or center), geometry, or properties related to thermionic emission. The magnetic field can be parameterized as a smooth function of axial/radial location  $\mathbf{B} = f(z, r)$  or by pointwise data. We give all the  $\mathbf{d}_T$  parameter values used in this study for the SPT-100 in Table 6. We assume a boron nitride—silica mixture for the wall material with material properties included in *Hall-thruster.jl* [17]. The magnetic field profile is given as the radial field strength  $B_r$  as a function of axial location  $z$ , obtained from curve fits to prior electromagnetic simulations of the SPT-100 circuit (see references therein).

We note that Table 6 may be an oversimplification of true thruster designs; modern designs that use central-mounted cathodes and magnetic shielding might require a more detailed parameterization, especially when used in a 2d finite-element simulation. Such mesh-based simulations would ideally handle arbitrary geometry files, materials, magnetic fields, etc. as do modern FEA codes. We provide Table 6 simply to indicate the design parameters used in this study. Critically, this implies that all of our results are likely limited to this specific set of  $\mathbf{d}_T$  for the SPT-100. All calibration results were performed assuming  $\mathbf{d}_T = \text{constant}$  and exclusively with SPT-100

experimental data; model fit coefficients very likely don't extrapolate to other thrusters. We leave it as future work to extend the analysis to other thrusters and operating regimes.

### Experimental data

Table 7 summarizes the experimental data used in this study that was extracted from the SPT-100 literature (most of this data was collected as part of facility effect characterization tests). The measurement quantities for each dataset are expressed in the nomenclature of Table 2.

**Table 7 Summary of experimental datasets used in this study**

Number	Description	Measurement quantity	Reference
1	Cathode coupling voltage measurements over varying background pressures.	$V_{cc}$	[16]
2	Thrust and discharge current measurements at two facilities over varying background pressures.	$T, I_d$	[33]
3	Ion current density measurements over varying background pressures.	$j_{ion}$	[33]
4	Laser-induced fluorescence measurements of average ion velocity along channel centerline.	$u_{ion}$	[21]
5	Thrust and discharge current measurements for various operating conditions.	$T, I_d$	[34]

We additionally provide tabulated forms of Datasets 1 and 2 in Tables 8 and 9. The tables provide the background pressure, anode voltage, and anode mass flow rate for each data point after some processing from the original sources. Except where otherwise noted, the experimental uncertainty in the data is expressed as error ( $\pm$ ) bars reported or inferred from the original sources. In this study, we treat these errors as additive Gaussian noise centered at the reported value with a characteristic standard deviation equal to 1/4 the error range. The discharge current for Datasets 1-3 were reported as a constant  $I_d = 4.5$  A for all operating conditions, for which we assume a  $\pm 10\%$  uncertainty.

**Table 8 Cathode coupling voltage data from Ref. [16]**

Pressure (Torr)	Anode voltage (V)	Anode flow rate (mg/s)	Cathode coupling voltage (V)	Error ( $\pm$ V)
1.67E-6	300	5.194	31.20	0.3
3.45E-6	300	5.222	31.12	0.3
4.11E-6	300	5.252	31.95	0.3
4.27E-6	300	5.215	31.29	0.3
6.97E-6	300	5.209	32.02	0.3
8.04E-6	300	5.195	31.89	0.3
1.23E-5	300	5.229	32.84	0.3
1.44E-5	300	5.151	32.30	0.3
1.58E-5	300	5.199	32.83	0.3
2.49E-5	300	5.077	33.11	0.3
2.51E-5	300	5.055	33.09	0.3

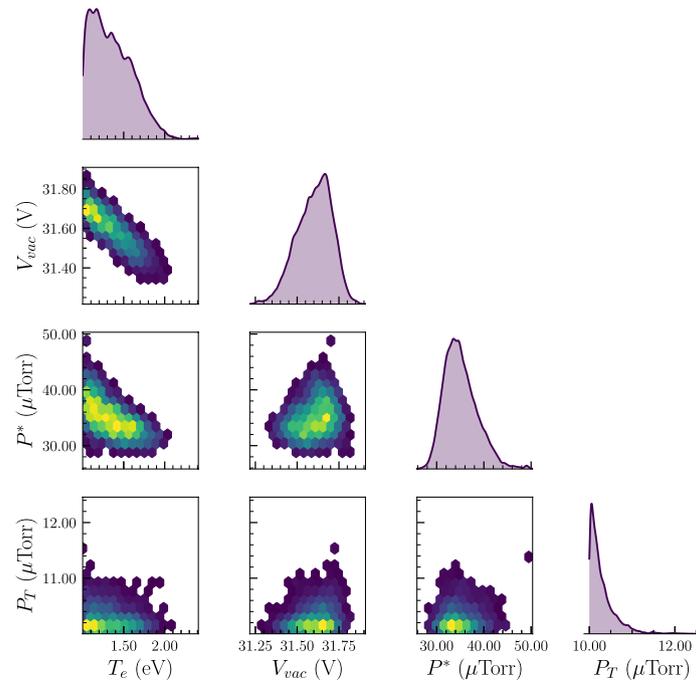
Pressure (Torr)	Anode voltage (V)	Anode flow rate (mg/s)	Cathode coupling voltage (V)	Error ( $\pm$ V)
3.82E-5	300	4.951	32.58	0.3
4.91E-5	300	4.926	32.46	0.3
5.51E-5	300	4.839	32.10	0.3
7.37E-5	300	4.726	32.13	0.3

**Table 9 Thrust data from Ref. [33]**

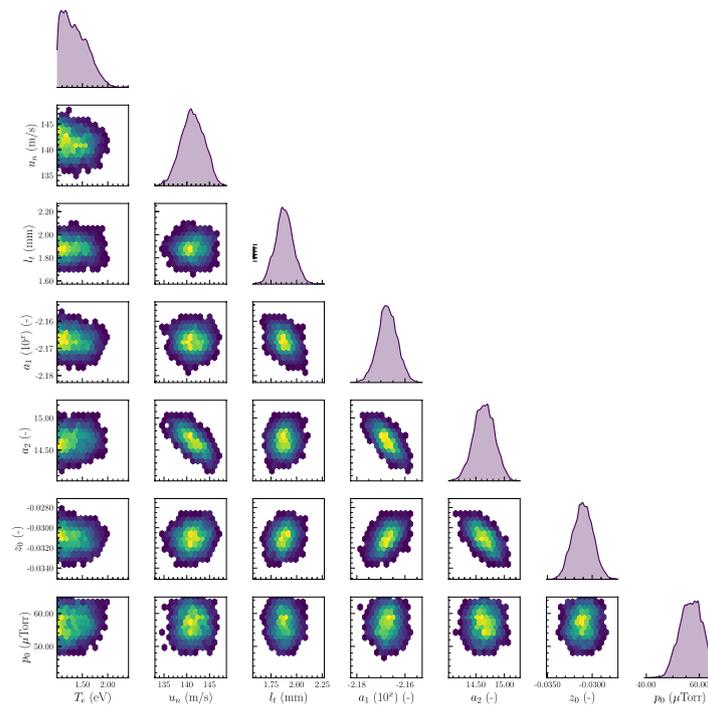
Pressure (Torr)	Anode voltage (V)	Anode flow rate (mg/s)	Thrust (mN)	Error ( $\pm$ %)
1.63E-6	300	5.194	80.29	0.9%
2.47E-6	300	5.161	80.06	0.9%
4.00E-6	300	5.222	80.05	0.7%
4.03E-6	300	5.252	80.46	0.9%
5.20E-6	300	5.215	80.13	0.7%
5.29E-6	300	5.142	80.81	0.9%
6.96E-6	300	5.209	81.15	0.9%
8.04E-6	300	5.195	81.51	0.7%
9.88E-6	300	5.229	81.44	0.9%
1.44E-5	300	5.151	82.55	0.7%
1.52E-5	300	5.199	81.76	0.9%
2.22E-5	300	5.077	82.64	0.7%
2.50E-5	300	5.055	82.31	0.9%
3.81E-5	300	4.951	83.11	0.9%
4.22E-5	300	4.926	83.11	0.7%
5.44E-5	300	4.839	82.39	0.9%
6.17E-5	300	4.797	83.04	0.7%
7.02E-5	300	4.764	82.35	0.9%

**Appendix D: MCMC marginals**

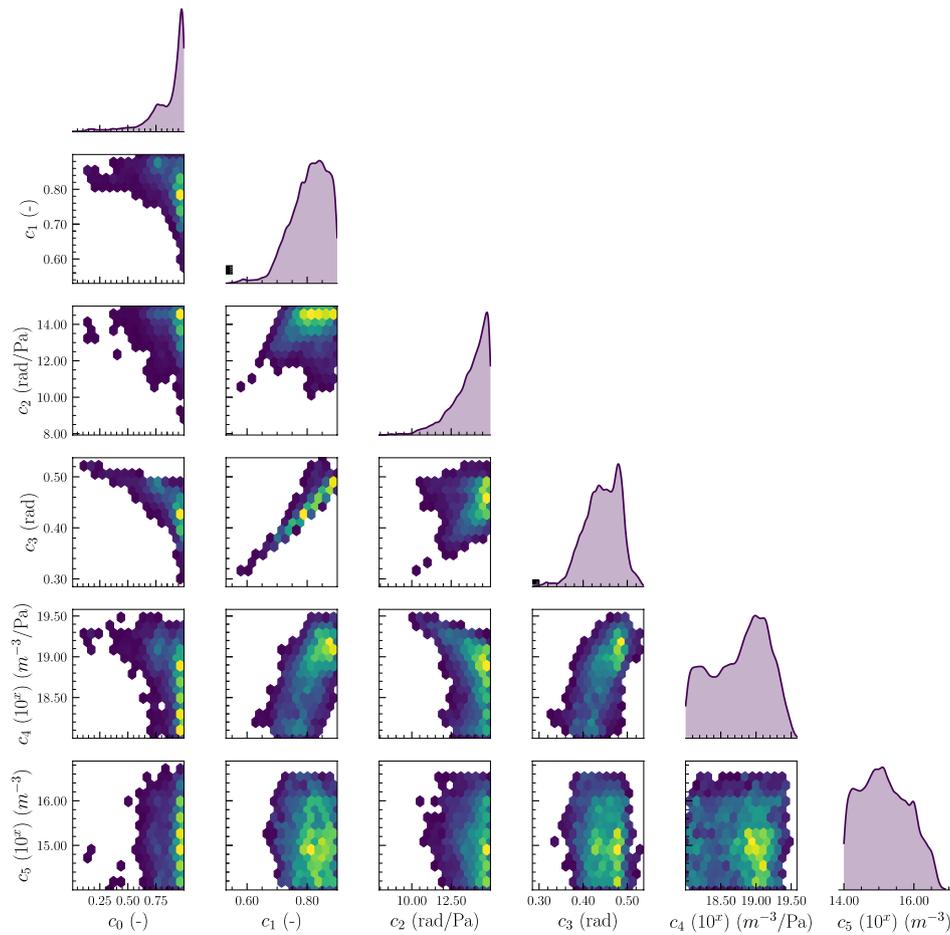
In this section, we provide plots of the 1d and 2d marginal posteriors to supplement the data in Table 3. These plots give a better understanding of the actual *shape* of the posterior distribution. Since there are 17 parameters total, we break the posterior into groups of parameters corresponding to each component model. Figure 14 shows the marginals of the cathode model parameters. Likewise, Figs. 15 and 16 show the marginals for the thruster and plume model parameters, respectively.



**Fig. 14** 1d and 2d marginal posteriors for the cathode model parameters



**Fig. 15** 1d and 2d marginal posteriors for the thruster model parameters



**Fig. 16** 1d and 2d marginal posteriors for the plume model parameters

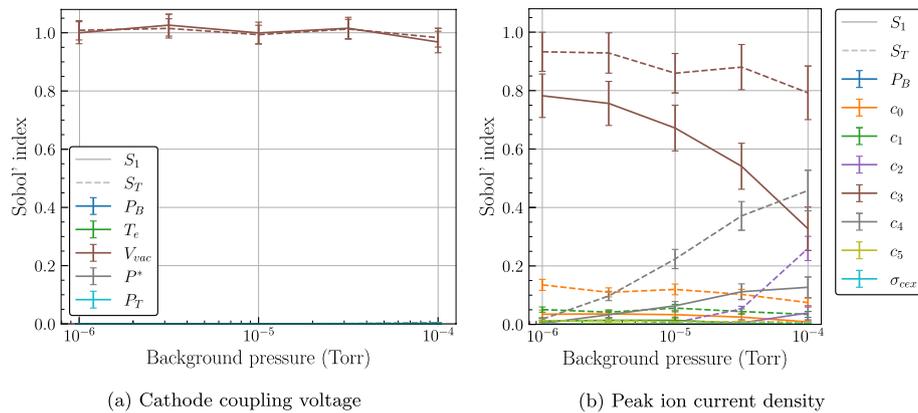
All the 1d marginal plots use a Gaussian kernel density estimate to smooth out the distributions. The 2d marginals show histograms of the 2d distribution, with lighter colors indicating higher probability density. It is evident from these plots that most parameters have a Gaussian shape with very tight bounds. A strong positive correlation exists between  $c_1$  and  $c_3$  which can be understood by combining Eqs. (29) and (30) for the scattered beam characteristic divergence angle:

$$\gamma_2 = \frac{c_2}{c_1} P_B + \frac{c_3}{c_1}. \tag{35}$$

This suggests that a better parameterization may exist for  $\gamma_2$  that is independent of the  $c_1$  proportionality to  $\gamma_1$ . Another interesting result in Fig. 16 is the wide posterior of  $(c_4, c_5)$  that suggests the data was not informative for learning these parameters. Since these parameters are tied to the plume radius  $r_p$  in the exponential of Eqs. (27) and (28) and since data was only collected at a single location  $r_p = 1$  m, it may be that data collection at different  $r_p$  is necessary to reduce the uncertainty in these parameters further.

## Appendix E: Additional Sobol' analysis

We provide the results of the Sobol' sensitivity analysis for both the cathode and plume model parameters here. Figure 17 shows the first- and total-order Sobol' indices for the cathode and plume parameters over varying background pressures.



**Fig. 17** Sobol' indices over varying background pressures for cathode coupling voltage and peak ion current density. Dashed lines indicate total-order ( $S_T$ ) indices and solid lines indicate first-order ( $S_1$ ) indices. Error bars give 95% Monte Carlo confidence intervals. The vacuum coupling voltage  $V_{vac}$  is shown to have the dominant impact in the variance of cathode coupling voltage. The plume parameters ( $c_3, c_4$ ) are shown to have the largest impact on variance of peak ion current density

The result in Fig. 17a trivially demonstrates that the cathode coupling voltage is nearly fully determined by the  $V_{vac}$  parameter as is evident in Eq. (3), where  $V_{vac}$  sets the asymptotic value for  $V_{cc}$  as  $P_B \rightarrow 0$ . The other parameters only serve to slightly modulate  $V_{vac}$  at higher pressures. If the prior on  $V_{vac}$  is known to a smaller domain than  $[0, 60]$  V, then the other parameters might be shown to have more relative importance.

Figure 17b shows the Sobol' indices for the plume parameters over varying background pressures. The two dominating parameters in predicting peak ion current density in the plume are the divergence angle offset ( $c_3$ ) and the neutral density slope ( $c_4$ ), with  $c_4$  naturally only becoming important at higher pressures. Recall from Eq. (29) that  $c_3$  is critical in determining the characteristic divergence angle of the plume; a higher divergence spreads the current away from centerline and vice versa. This impact is less at higher pressures due to the  $c_2 P_B$  term in Eq. (29). Uncertainty in parameters like  $\sigma_{cex}$  and the pressure  $P_B$  can safely be ignored based on this analysis.

### Acknowledgements

This research was supported in part through computational resources provided by Advanced Research Computing at the University of Michigan.

### Authors' contributions

Code, data, analysis, primary writing - J.E.; Primary revisions, conceptualization - A.G.; Thruster model, analysis - T.M.; Plume model - M.A.; Cathode model, conceptualization, review - B.J. All authors have read and agreed to the published version of the manuscript.

### Funding

Funding for this work was provided by NASA in part through the Joint Advanced Propulsion Institute (JANUS), a NASA Space Technology Research Institute, under grant number 80NSSC21K1118, as well as in part through a NASA Space Technology Graduate Research Opportunity grant 80NSSC23K1181.

**Data availability**

The datasets generated and analyzed during this study are available from the corresponding author on reasonable request.

**Code availability**

The code generated during this study is available at <https://github.com/JANUS-Institute/HallThrusterPEM.git>. The surrogate code originally described in Ref. [15] has been implemented and made openly available at <https://github.com/eckelsjd/amisc.git>.

**Declarations****Competing interests**

The authors declare no competing interests.

Received: 20 May 2024 Accepted: 30 August 2024

Published online: 27 September 2024

**References**

1. Taccogna F, Garrigues L (2019) Latest progress in Hall thrusters plasma modelling. *Rev Mod Plasma Phys* 3(1):12. <https://doi.org/10.1007/s41614-019-0033-1>
2. Boeuf JP (2017) Tutorial: Physics and modeling of Hall thrusters. *J Appl Phys* 121(1):011101. <https://doi.org/10.1063/1.4972269>
3. Kaganovich ID, Smolyakov A, Raitses Y, Ahedo E, Mikellides IG, Jorns B, Taccogna F, Gueroult R, Tsikata S, Bourdon A, Boeuf JP, Keidar M, Powis AT, Merino M, Cappelli M, Hara K, Carlsson JA, Fisch NJ, Chabert P, Schweigert I, Lafleur T, Matyash K, Khrabrov AV, Boswell RW, Fruchtman A (2020) Physics of  $E \times B$  discharges relevant to plasma propulsion and similar technologies. *Phys Plasmas* 27(12):120601. <https://doi.org/10.1063/5.0010135>
4. Spalart PR (2015) Philosophies and fallacies in turbulence modeling. *Prog Aerosp Sci* 74:1–15. <https://doi.org/10.1016/j.paerosci.2014.12.004>
5. Jorns B, Marks TA, Dale ET (2020) A Predictive Hall Thruster Model Enabled by Data-Driven Closure. In: *AIAA Propulsion and Energy 2020 Forum*. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2020-3622>
6. Marks TA, Jorns BA (2023) Evaluation of algebraic models of anomalous transport in a multi-fluid Hall thruster code. *J Appl Physics* 134(15):153301. <https://doi.org/10.1063/5.0171824>
7. Randolph T, Kim V, Kaufman H, Kozubsky K, Zhurin V, Day M (1993) Facility effects on stationary plasma thruster testing. In: *Proc. of the 23rd international electric propulsion conference*. Electric Rocket Propulsion Society, Seattle
8. Dankanich JW, Walker M, Swiatek MW, Yim JT (2017) Recommended Practice for Pressure Measurement and Calculation of Effective Pumping Speed in Electric Propulsion Testing. *J Propuls Power* 33(3):668–680. <https://doi.org/10.2514/6.2017-1135>
9. Walker MLR, Lev D, Saeedifard M, Jorns BA, Foster J, Gallimore AD, Gorodetsky AA, Rovey JL, Chew HB, Levin D, Williams JD, Yalin A, Wirz RE, Marian J, Boyd I, Hara K, Lemmer K (2022) Overview of the Joint Advanced Propulsion Institute (JANUS). In: *Proc. of the 37th International electric propulsion conference*. Electric Rocket Propulsion Society, Boston
10. Brieda L, Keidar M (2013) Multiscale Modeling of Hall Thrusters. In: *Proc. of the 33rd International Electric Propulsion Conference*. Electric Rocket Propulsion Society, Washington DC
11. Lopez Ortega A, Katz I, Mikellides IG, Goebel DM (2015) Self-Consistent Model of a High-Power Hall Thruster Plume. *IEEE Trans Plasma Sci* 43(9):2875–2886. <https://doi.org/10.1109/TPS.2015.2446411>
12. Brieda L, Keidar M (2017) Towards Full Chamber HET Simulations with a Single 2D Hall Thruster Model. In: *Proc. of the 35th International Electric Propulsion Conference*. Electric Rocket Propulsion Society, Atlanta
13. Boyd I, Garrigues L, Koo J, Keidar M (2000) Progress in development of a combined device/plume model for Hall thrusters. In: *36th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, Joint Propulsion Conferences*. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2000-3520>
14. Qarnain S, Martinez-Sanchez M (1998) Issues regarding the generation of an end-to-end Hall thruster computational model. In: *34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, Joint Propulsion Conferences*. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.1998-3796>
15. Jakeman JD, Friedman S, Eldred MS, Tamellini L, Gorodetsky AA, Allaire D (2022) Adaptive experimental design for multi-fidelity surrogate modeling of multi-disciplinary systems. *Int J Numer Methods Eng* 123(12):2760–2790. <https://doi.org/10.1002/nme.6958>
16. Jorns BA, Byrne MP (2021) Model for the dependence of cathode voltage in a Hall thruster on facility pressure. *Plasma Sources Sci Technol* 30(1):015012. <https://doi.org/10.1088/1361-6595/abd3b6>
17. Marks T, Schedler P, Jorns B (2023) HallThruster.jl: A Julia package for 1D Hall thruster discharge simulation. *J Open Source Softw* 8(86):4672. <https://doi.org/10.21105/joss.04672>
18. Hofer RR (2004) Development and Characterization of High-Efficiency, High-Specific Impulse Xenon Hall Thrusters. PhD thesis, University of Michigan
19. Hofer R, Katz I, Goebel D, Jameson K, Sullivan R, Johnson L, Mikellides I (2008) Efficacy of Electron Mobility Models in Hybrid-PIC Hall Thruster Simulations. In: *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2008-4924>

20. Marks TA, Jorns BA (2023) Challenges with the self-consistent implementation of closure models for anomalous electron transport in fluid simulations of Hall thrusters. *Plasma Sources Sci Technol* 32(4):045016. <https://doi.org/10.1088/1361-6595/accd18>
21. Macdonald N, Pratt Q, Nakles M, Pilgram N, Holmes M, Hargus W (2019) Background Pressure Effects on Ion Velocity Distributions in an SPT-100 Hall Thruster. *J Propuls Power* 35:1–10. <https://doi.org/10.2514/1.B37133>
22. Allen MG, Eckels JD, Byrne MP, Gorodetsky AA, Jorns BA (2022) Application of Optimal Experimental Design to Characterize Pressure Related Facility Effects in a Hall Thruster. In: Proc. of the 37th International Electric Propulsion Conference. Electric Rocket Propulsion Society, Boston
23. Haji-Ali AL, Nobile F, Tamellini L, Tempone R (2016) Multi-Index Stochastic Collocation for random PDEs. *Comput Methods Appl Mech Eng* 306:95–122. <https://doi.org/10.1016/j.cma.2016.03.029>
24. Jakeman JD, Eldred M, Geraci G, Gorodetsky A (2020) Adaptive Multi-index Collocation for Uncertainty Quantification and Sensitivity Analysis. *Int J Numer Methods Eng* 121(6):1314–1343. <https://doi.org/10.1002/nme.6268>
25. Piazzola C, Tamellini L, Pellegrini R, Broglia R, Serani A, Diez M (2022) Comparing multi-index stochastic collocation and multi-fidelity stochastic radial basis functions for forward uncertainty quantification of ship resistance. *Eng Comput*. <https://doi.org/10.1007/s00366-021-01588-0>
26. Narayan A, Jakeman JD (2014) Adaptive Leja Sparse Grid Constructions for Stochastic Collocation and High-Dimensional Approximation. *SIAM J Sci Comput* 36(6):A2952–A2983. <https://doi.org/10.1137/140966368>
27. Higdon D, Gattiker J, Williams B, Rightley M (2008) Computer Model Calibration Using High-Dimensional Output. *J Am Stat Assoc* 103(482):570–583. <https://doi.org/10.1198/01621450700000888>
28. Kapsuzoglu B, Mahadevan S, Matsumoto S, Miyagi Y, Watanabe D (2022) Adaptive surrogate modeling for high-dimensional spatio-temporal output. *Struct Multidiscip Optim* 65(10):300. <https://doi.org/10.1007/s00158-022-03402-x>
29. Crevillén-García D (2018) Surrogate modelling for the prediction of spatial fields based on simultaneous dimensionality reduction of high-dimensional input/output spaces. *R Soc Open Sci* 5(4):171933. <https://doi.org/10.1098/rsos.171933>
30. Haario H, Laine M, Mira A, Saksman E (2006) DRAM: Efficient adaptive MCMC. *Stat Comput* 16(4):339–354. <https://doi.org/10.1007/s11222-006-9438-0>
31. Sobol'IM (2001) Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Math Comput Simul* 55(1):271–280. [https://doi.org/10.1016/S0378-4754\(00\)00270-6](https://doi.org/10.1016/S0378-4754(00)00270-6)
32. Saltelli A, Annoni P, Azzini I, Campolongo F, Ratto M (2010) Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index. *Comput Phys Commun* 181(2):259–270. <https://doi.org/10.1016/j.cpc.2009.09.018>
33. Diamant KD, Liang R, Corey RL (2014) The Effect of Background Pressure on SPT-100 Hall Thruster Performance. In: 50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, AIAA Propulsion and Energy Forum. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2014-3710>
34. Sankovic J, Hamley J, Haag T (1993) Performance evaluation of the Russian SPT-100 thruster at NASA LeRC. In: 23rd International Electric Propulsion Conference. Electric Rocket Propulsion Society, Seattle
35. Mikellides IG, Katz I (2012) Numerical simulations of Hall-effect plasma accelerators on a magnetic-field-aligned mesh. *Phys Rev E* 86(4):046703. <https://doi.org/10.1103/PhysRevE.86.046703>
36. Frieman JD, Liu TM, Walker MLR (2017) Background Flow Model of Hall Thruster Neutral Ingestion. *J Propuls Power* 33(5):1087–1101. <https://doi.org/10.2514/1.B36269>
37. Roberts PJ, Jorns B (2024) Inferring Electron Heat Flux in a High-Power Hall Thruster with Incoherent Thomson Scattering. In: AIAA SCITECH 2024 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2024-1957>
38. Aksoy D, Gorsich DJ, Veerapaneni S, Gorodetsky AA (2022) An Incremental Tensor Train Decomposition Algorithm. <https://doi.org/10.48550/arXiv.2211.12487>
39. Miller JS, Pullins SH, Levandier DJ, Chiu Yh, Dressler RA (2002) Xenon charge exchange cross sections for electrostatic thruster models. *J Appl Phys* 91(3):984–991. <https://doi.org/10.1063/1.1426246>
40. Huang W, Kamhawi H, Haag T, Ortega AL, Mikellides IG (2016) Facility Effect Characterization Test of NASA's HERMeS Hall Thruster. In: 52nd AIAA/SAE/ASEE Joint Propulsion Conference. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2016-4828>
41. Mikellides I, Katz I, Mandell M, Snyder J (2001) A 1-D model of the Hall-effect thruster with an exhaust region. In: 37th Joint Propulsion Conference and Exhibit, Joint Propulsion Conferences. American Institute of Aeronautics and Astronautics. <https://doi.org/10.2514/6.2001-3505>
42. Gascon N, Dudeck M, Barral S (2023) Wall material effects in stationary plasma thrusters. I. Parametric studies of an SPT-100. *Phys Plasmas* 10(10):4123–4136. <https://doi.org/10.1063/1.1611880>
43. Panelli M, Morfei D, Milo B, D'Aniello FA, Battista F (2021) Axisymmetric Hybrid Plasma Model for Hall Effect Thrusters. *Particles* 4(2):296–324. <https://doi.org/10.3390/particles4020026>
44. Sahu R, Mansour AR, Hara K (2020) Full fluid moment model for low temperature magnetized plasmas. *Phys Plasmas* 27(11):113505. <https://doi.org/10.1063/5.0021474>

## Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.