Bayesian calibration of a multi-component model of the SPT-140 Hall thruster

Collin B. Whittaker*, Thomas A. Marks†, and Alex A. Gorodetsky‡
University of Michigan, Ann Arbor, MI, 48109

We deploy the Joint Advanced Propulsion Institute's (JANUS's) Hall thruster Predictive Engineering Model (PEM) model on the SPT-140 system qualified and flown for NASA's Psyche spacecraft. The purpose of the PEM is to enable rapid predictions and uncertainty quantification of system performance while accounting for facility effects associated with ground-based testing. The most recent instantiation of the PEM framework consists of an analytical model predicting cathode coupling voltage, a 1-D axial fluid Hall thruster code predicting the main discharge, and an analytical plume expansion model predicting ion current density traces. The primary facility effect they capture is background pressure. To deploy the PEM for the SPT-140, we extended the models to enable scaling the Hall thruster's magnetic field at different operating conditions. We then calibrated the PEM using data from ground qualification tests performed by NASA's Jet Propulsion Laboratory (JPL). These datasets include thrust, discharge current, and ion acceleration measurements. A Bayesian calibration of the model parameters yields a posterior consistent with previous calibrations on the SPT-100 and H9 thrusters. Unlike standard practice, where anomalous collisions are calibrated to individual operating conditions, we use a single anomalous transport model parameterized by pressure. Thus we provide initial evidence for predictive capabilities of target quantities of interest. Following calibration, we extrapolate the model predictions to vacuum. Thrust predictions at higher powers are shown to be comparable to modeling by JPL using Hall2De done during qualification, while discharge current data are not reliably predicted. The mismatch of discharge current is attributed to the fact that the PEM assumes a fixed anomalous collision frequency profile whose position shifts with pressure.

I. Introduction

Hall thrusters are space-based plasma devices that must be tested in ground vacuum facilities that incompletely simulate the space environment. One critical difference between ground and space that we focus on here is the the effect of background gas in the facility. More specifically, ground testing necessitates extrapolating performance measured with finite background pressure to the vacuum of space. Extrapolating the performance and internal states of Hall thrusters at varying background pressures is challenging due the lack of *ab initio* models of the thruster plasma—many of the physical mechanisms leading to changes in operation with pressure are not completely understood. In principle, many of these effects could be captured in a kinetic model of the discharge plasma, but facility-scale kinetic simulations do not yet exist at a feasible computational cost. To make up for this gap in understanding, engineering models of Hall thrusters incorporate parameters that stand in for the missing physics. These coarse-scale models, however, must be calibrated to match data. Traditionally, the model's parameters are optimized for each unique set of operating conditions (discharge voltage and current, background pressure, etc.), and the models obtained thus lack the power to extrapolate [1–3]. A more extensible approach is to calibrate a single model across a wide range of conditions, and we follow this approach here. To this end, we require a model that is flexible enough to reproduce and extrapolate the observed variation of thruster performance with facility conditions (in this work, primarily pressure), while simultaneously constrained by critical physical principles. At the same time, to enable automatic calibration and uncertainty quantification, both of which require potentially thousands of sequential model evaluations, we require an inexpensive model.

To respond to these needs, we developed the Hall thruster Predictive Engineering Model (PEM) [4] as part of NASA's Joint Advanced Propulsion Institute (JANUS) [5]. This model is built from three components—a model for the cathode coupling voltage, a 1-D fluid Hall thruster discharge model, and an analytical model for the expansion of the

---

*Assistant Research Scientist, Aerospace Engineering, cbwhitt@umich.edu, Member AIAA
†Assistant Research Scientist, Aerospace Engineering, Member AIAA
‡Associate Professor, Aerospace Engineering, Associate Fellow AIAA

plasma plume. This model is designed to predict several quantities of interest (including thrust, discharge current, and ion velocity) as they vary with background pressure. Crucially, each component model builds in an explicit pressure dependence as well as several tunable parameters. This infrastructure allows the full system to be calibrated on data that range across multiple background pressures and operating conditions and then extrapolate to unseen environments.

In our initial work [4], we demonstrated the PEM on the SPT-100 Hall thruster and obtained good agreement with the calibration dataset and satisfactory extrapolation to a validation dataset. In a subsequent paper [6], we improved the flexibility of the model and both re-tested it on the SPT-100 and extended it the H9 magnetically-shielded Hall thruster. For both thrusters, the model exhibited predictive errors of less than 10% on the validation set. In this work, we apply the Hall thruster PEM to a new thruster—the SPT-140 recently flown on NASA's Psyche spacecraft [7].

The SPT-140 thruster and the available datasets necessitate certain changes compared with our previous efforts. In particular, in this paper we consider the the following changes and differences: we 1) incorporate multiple power conditions that have different magnetic field settings; 2) train without access to cathode coupling voltage and ion current density data; and 3) study a more consistent and well-defined, which we track from ground-based qualification testing to flight.

After the incorporation of these new capabilities and dataset modifications, our model produces thrust and discharge currents within 7% and 15% of the experimental data used for calibration, respectively. Additionally, we demonstrate 8% and 17% errors in the same quantities of interest when extrapolating ground test data to orbit. We compare our calibrated simulation results to those obtained using Hall2De in prior works [2], a higher-fidelity code. We find that the error in our discharge current predictions is the result of using the same anomalous electron collision frequency profile across multiple power levels. This contrasts with the Hall2De results, which were calibrated for each operating condition individually. Finally, we analyze the amount of data required to calibrate the PEM on this dataset. We find that the model predictions for thrust on-orbit change by less than 3% when calibrating the model on only data taken above 20 µTorr. However, we find that uncertainty in the position of the ion acceleration region increases by a factor 2.5. These findings suggest that ground testing at ultra-low pressures (< 5 µTorr) may not be needed to extrapolate system-level parameters like thrust to space, but may still be required to accurately capture the internal discharge dynamics.

This paper is organized as follows. In Sec. II we describe the PEM and its components (Sec. II.A), the data sources we employ (Sec. II.B), and the Bayesian inference procedure we apply for model calibration (Sec. II.C). Then, in Sec. III, we present the results of this calibration, including the values of the calibrated model parameters (Sec. III.A, a comparison of the model to its training dataset (Sec. III.B), and an extrapolation to a space validation dataset (Sec. III.C). Finally, we discuss these results in Sec. IV, highlighting the role changes in anomalous mobility play in our predictions (Sec. IV.A), examining what limiting calibration to only the highest pressure data modifies in our model (Sec. IV.B), and investigating how the dependence of mass flow rate and discharge current might be inverted in the model to more accurately reflect how data are acquired experimentally (Sec. IV.C).

II. Methodology

In this section, we describe our methodologies. We first provide an overview of the engineering model used to predict thruster behavior (Sec. II.A). Then, we describe the sources of training and validation data for that model (Sec. II.B). Finally, we detail the inference problem by which we calibrate the model to the training data (Sec. II.C).

A. Model

We leverage the Hall thruster Predictive Engineering Model (PEM) [4] of the Joint Advanced Propulsion Institute (JANUS) [5] to model thruster behavior. This model's purpose is to capture facility effects endemic to testing Hall thrusters in ground based vacuum chambers, which incompletely simulate the space environment where the thrusters operate—in particular the effect of finite background pressure in the test facility. The PEM is constituted by three components: 1) an analytical model which predicts the pressure-dependent coupling between the cathode and the main discharge [8]; 2) a 1-D axial fluid code that simulates the main discharge, including the ionization and acceleration of the plasma [9]; and 3) a semi-empirical formulae which analytically propagate the plume of the discharge through the background gas of the facility [10]. These components and their connection to each other are more expansively detailed in our prior work [6] and summarized in Fig. 1. We briefly review and highlight key features in the following subsections.

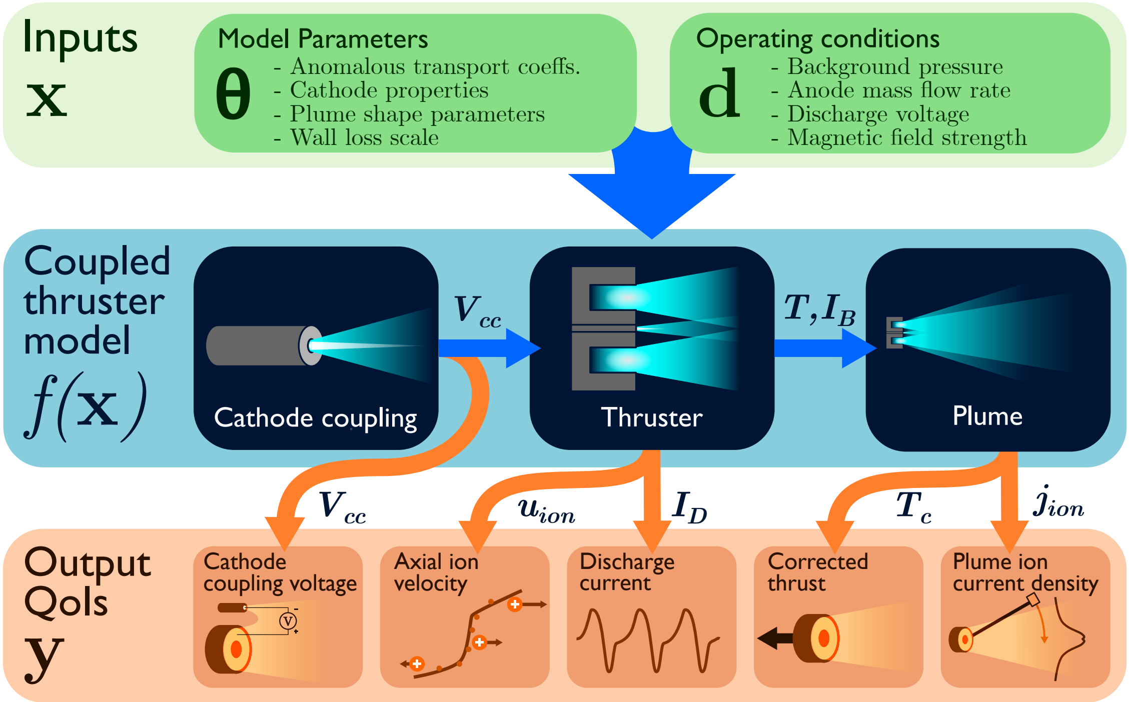

Fig. 1 Schematic overview of the JANUS Hall thruster PEM; figured adapted from [6].

1. Cathode Coupling

The cathode model (originally from [8]) is given compactly as:

$$V_{cc} = V_{vac} + T_{ec} \ln\left(1 + \frac{P_B}{P_T}\right) - \left(\frac{T_{ec}}{P_T + P^*}\right) P_B. \tag{1}$$

Here, V_{cc} denotes the cathode coupling voltage—that is, the potential drop from the cathode to the main discharge necessary to extract electrons—which is expressed as a function of background pressure in the facility P_B and the effective temperature of the cathode plasma T_{ec}. Correspondingly, V_{vac} is the coupling voltage of the system in an ideal vacuum, because $V_{cc} = V_{vac}$ when $P_B = 0$. The quantities P_T and P^* are reference pressures which serve as scaling coefficients. Equation (1) arises from balancing the effects of pressure gradients (logarithmic term) and enhanced conductivity (linear term) in the cathode plasma as a result of excess neutrals from the facility, and is derived in detail in

2. Thruster

To model the main discharge, we use `HallThruster.jl`, a 1-D axial fluid Hall thruster code [9]. `HallThruster.jl` models the quasineutral plasma as composed of separate neutral, ion, and electron fluids. Continuity is solved for the neutrals assuming they are subject to ionization but otherwise traveling at a fixed drift speed u_n, and both continuity and momentum are solved for the ions. For electrons, `HallThruster.jl` solves an energy equation but otherwise treats them as inertialess, invoking the generalized Ohm's law in combination with current conservation to solve for electron current density and the electrostatic potential. The upstream boundary of the simulation is set to the anode potential, while the downstream boundary of the simulation is set to the cathode coupling voltage from Eq. (1).

Part of the domain (the channel region) is bounded by walls, and the plasma is partially magnetized by a fixed transverse magnetic field profile; these are configured to match the design of the thruster. The radial magnetic field B_r is a function of position $B_r = B_r(z)$, with z the axial coordinate. For magnetic circuits composed of electromagnets, the magnetic field can be adjusted by modifying the currents passing through the coils. `HallThruster.jl` exposes a magnetic field scale parameter $\hat{B}_r$ to account for these changes, calculating

$$B_r = B_{r,ref} \hat{B}_r, \tag{2}$$

where $B_{r,ref}$ is some reference field profile. While this simple linear scaling cannot model distortions in the field profile caused by saturation of the magnetic circuit or a change in balance between multiple coils, magnetic circuits are often designed with a fixed topology restricted to operate within a linear range and scaled therein to meet stability margin [11], so this scaling often suffices. In particular, if I_{ref} is a reference electromagnetic coil current and I_c an actual coil current, then we have (see further discussion in Sec. II.B)

$$\hat{B}_r = I_c / I_{ref}. \tag{3}$$

`HallThruster.jl`—like other fluid Hall thruster codes [12–15]—includes an "anomalous" electron mobility term to capture the effect of instability-induced electron transport across the radial magnetic field. These kinetic effects are not yet self-consistently captured in the model—doing so more generally is an active area of research [16–18] . In `HallThruster.jl`, the heightened mobility is implemented as a position-dependent anomalous collision frequency, in particular following the phenomenological form,

$$\nu_{anom} = \omega_{ce}\alpha_{anom}\left(1 - \beta_{anom}\exp\left[-\left(\frac{\hat{z} - (\hat{z}_{anom} - \Delta\hat{z})}{\hat{L}_{anom}}\right)^2\right]\right), \tag{4}$$

which we illustrate in Fig. 2. Here, ν_{anom} is the anomalous electron collision frequency and $\omega_{ce} \propto B_r$ is the electron cyclotron frequency. Noting the right hand side of Eq. 4 scales as ω_{ce} times some quantity $c \ll 1$ and that the anomalous collision frequency tends to dominate the classical collision frequency, this induces a Bohm-like scaling for the electron mobility across magnetic field lines:

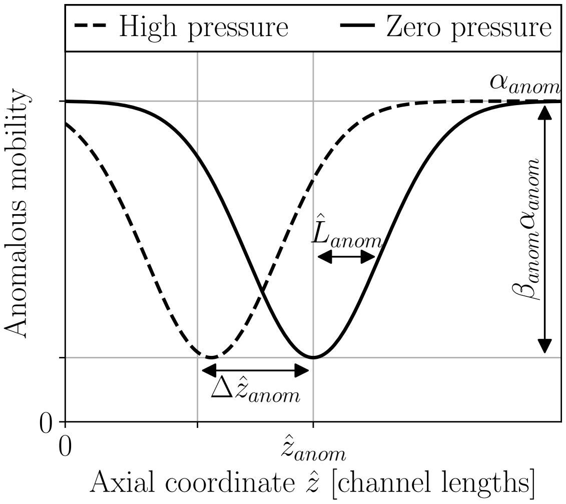

Fig. 2 Diagram of the anomalous collision frequency model of Eq. (4), adapted from [6].

$$\mu_\perp \propto \frac{\frac{1}{\nu_{anom}}}{1 + \left(\frac{\omega_{ce}}{\nu_{anom}}\right)^2} \propto \frac{\frac{1}{c\,\omega_{ce}}}{1 + \frac{1}{c^2}} \propto \frac{c}{\omega_{ce}} \propto \frac{c}{B_r}. \tag{5}$$

Additionally, there exists a localized transport barrier which serves to concentrate potential drops in the plasma to a narrow "acceleration region" [3, 19]—here modeled by the terms in the largest parentheses of Eq. (4) as a Gaussian dip in the collision frequency. The quantity $1 - \beta_{anom}$ is the multiplicative decrease in frequency at the nadir of the barrier, $\hat{L}_{anom}$ is a characteristic width of the barrier, $\hat{z}$ is the axial coordinate, and $\hat{z}_{anom}$ is the center of the barrier. The circumflexes denote that these terms have been normalized by the channel length of the thruster—i.e., $\hat{z} = z/L_{ch}$.

The final term $\Delta\hat{z} = \Delta\hat{z}(P_B)$ is a pressure-dependent shift in the location of the transport barrier which captures observations that the acceleration region migrates with changes in facility pressure [20]. In particular, we adopt a logistic form,

$$\Delta\hat{z}(P_B) = \Delta\hat{z}_{anom}\left(\frac{1}{1+\exp\left[2\left(1-\frac{P_B}{P_0}\right)\right]} - \frac{1}{1+\exp\left[2\right]}\right). \tag{6}$$

The term $\Delta\hat{z}_{anom}$ is the magnitude of the shift, and P_0 is its center. Consistent with our previous work [6], we fix the center of the logistic shift to $P_0 = 25\,\mu\text{Torr}$.

Two final scale parameters are exposed by the model. By default `HallThruster.jl` estimates the flux of neutrals ingested by the thruster from the background gas in the facility assuming the population is Maxwellian [9]. f_n is a multiplicative factor which scales the magnitude of this flux. Further, since the code is only 1-D, it does not directly compute radial gradients in the plasma density and so assumes that the density at the walls of the channel is a fixed fraction of the cell density (which is implicitly averaged over the cross section). We take this ratio to have a nominal value 0.5, modified by a scale factor c_w.

Independent of these scale factors and model parameters, `HallThruster.jl` has four key inputs: the discharge voltage V_d, the total voltage applied between the anode and cathode; the cathode coupling voltage V_{cc}, which we take as predicted by Eq. (1); the background pressure in the facility P_B; and the mass flow rate of propellant to the anode $\dot{m}_a$. These inputs specified, the model produces four key outputs: the discharge current I_d from the anode to cathode; the beam current I_B, the amount of the discharge current carried by ions and leaving the thruster; the “uncorrected” thrust T, an integrated momentum flux out these ions out of the domain assuming a perfectly collimated beam; and the velocity of the ion fluid as a function of position, $u_{ion} = u_{ion}(z)$.

3. Plume

`HallThruster.jl` is a 1-D code and *ipso facto* does not internally model the divergence and expansion of the ion beam into the test facility. We leverage the plume model of [10] to capture these behaviors. In particular, this analytical model represents the thruster as a point source of ion current; the corresponding ion current density j_{ion} is a function of radial distance away from the thruster r and polar angle off the thruster axis θ and is considered composed of three distinct populations, a “main” beam j_m, a “scattered” beam j_s, and a charge exchange beam j_{cx}:

$$j_{ion} = j_m + j_s + j_{cx}, \tag{7a}$$

$$j_m = \frac{(1-c_0)I_B}{r^2}\exp\left(-rn_n\sigma_{cx}\right)A_1\exp\left[-\left(\frac{\theta}{\theta_1}\right)^2\right], \tag{7b}$$

$$j_s = \frac{c_0 I_B}{r^2}\exp\left(-rn_n\sigma_{cx}\right)A_2\exp\left[-\left(\frac{\theta}{\theta_2}\right)^2\right], \tag{7c}$$

$$j_{cx} = \frac{I_B}{2\pi r^2}\left[1-\exp\left(-rn_n\sigma_{cx}\right)\right]. \tag{7d}$$

Equations (7a)-(7d) indicate the current density of each beam attenuates by default as r^{-2}. The main scattered beams expand in Gaussian profiles with characteristic widths θ_1 and θ_2, respectively, while the charge exchange beam is a uniform hemisphere. A_1 and A_2 are corresponding normalizing constants computed in [10]. The quantity c_0 is the proportion of the beam that is scattered into j_s, and the $\exp\left(-rn_n\sigma_{cx}\right)$ terms represent that as the beam expands radially and propagates through the background gas, it undergoes charge exchange collisions which attenuate the main and scattered beam and transfer this charge flux into the charge exchange beam. n_n is the effective neutral density from the background pressure and σ_{cx} is the cross section of the charge exchange process.

In recognition that the beam tends to diverge differently for varying background pressures as a result of interactions with the background gas and shifts in the acceleration region within the thruster [20], the beam widths and effective neutral density are functionalized by the background pressure as

$$\theta_1 = c_1\theta_2, \tag{8a}$$

$$\theta_2 = c_2 P_B + c_3, \tag{8b}$$

$$n_n = c_4 P_B + c_5, \tag{8c}$$

where c_1, c_2, c_3, c_4, and c_5 are calibration coefficients.

Aside from the calibration coefficients, the only inputs to the plume model are the background pressure P_B and the ion beam current I_B (the latter returned from the thruster model). The primary output is the position-dependent ion current density j_{ion}; however, at a higher level this predicted profile can be used to correct the thrust estimate produced by the thruster model, namely:

$$\cos\theta_d = \frac{\int_0^{\pi/2}(j_m + j_s)\cos\theta\sin\theta\,\mathrm{d}\theta}{\int_0^{\pi/2}(j_m + j_s)\sin\theta\,\mathrm{d}\theta}, \tag{9}$$

$$T_c = T\cos\theta_d. \tag{10}$$

Equation (9) is the classical computation of the divergence angle θ_d of the beam assuming the momentum flux is proportional to the current density [21]. We perform the computation excluding the current content of the charge exchange beam, since these collisions transfer momentum away from charged species, violating this assumption. Consequently, θ_d is solely a function of c_0, c_1, c_2, c_3, and P_B—i.e., is a calibrated function of background pressure. The thrust corrected by this pressure-dependent divergence angle, Eq. (10), is the principle output of the plume component and the overall model's best estimate for thrust to compare to data.

4. Composition

We summarize the inputs, outputs, and parameters of each component model in Tab. 1. The asterisks in the table indicate the quantities which couple between the various components—that is, where the output of one model is required as an input to another. Fortunately, the components do not feed back into each other, so these linkages are trivial. For our purposes, the aggregated model has four key inputs—P_B, V_d, $\dot{m}_a$, and $\hat{B}_r$—and three key outputs—I_d, u_{ion}, and T_c. However, the presence of the calibration parameters in between necessitates using experimental data to train the mapping between inputs and outputs.

Table 1 Summary of component model inputs and outputs; asterisks indicate variables that couple between component models, daggers denote notation revised from [6], and new variables are in bold.

Model	Inputs		Parameters		Outputs	
Cathode	background pressure	P_B	vacuum coupling voltage	V_{vac}	cathode coupling voltage*	V_{cc}
			cathode electron temperature	T_{ec}		
			"base" pressure	P_T		
			turning pressure	P^*		
Thruster	background pressure	P_B	anomalous mobility scale	α_{anom}	uncorrected thrust*	T
	discharge voltage	V_d	transport barrier scale	β_{anom}	ion beam current*	I_B
	cathode coupling voltage*	V_{cc}	transport barrier location†	$\hat{z}_{anom}$	discharge current	I_d
	anode mass flow rate	$\dot{m}_a$	transport barrier width†	$\hat{L}_{anom}$	ion velocity profile	u_{ion}
	magnetic field scale	$\hat{B}_r$	transport barrier shift†	$\Delta\hat{z}_{anom}$		
			neutral drift speed	u_n		
			wall density scale	c_w		
			neutral ingestion scale	f_n		
Plume	background pressure	P_B	scattered beam ratio	c_0	ion current density	j_{ion}
	ion beam current*	I_B	divergence angle ratio	c_1	corrected thrust	T_c
	uncorrected thrust*	T	divergence angle slope	c_2		
			divergence angle offset	c_3		
			neutral density slope	c_4		
			neutral density offset	c_5		

B. Data Sources

We briefly outline the experimental data we use to train and test the model in this section, which we also summarize this in Tab. 2. In contrast with our prior work [6], we do not have direct measurements of the cathode coupling voltage V_{cc} or ion current density profiles j_{ion} on which to calibrate the model. "Training" in the table indicates the data on which we calibrate the model, the ground testing qualification data of [22] and [2]. "Validation" indicates data that we do not train the model on and instead reserve to test its ability to extrapolate, the on-orbit measurements of [11].

Table 2 Summary of data sources used to train and validate the model.

Source	Reference	Category	Data Types	Power	Magnet Current I_c	Pressures P_B
Snyder et al. 2020	[22]	Training	I_d, T_c	4.5 kW	6.00 A	3.3-34 µTorr
				3.5 kW	5.50 A	2.8 µTorr
				3.0 kW	4.50 A	2.5 µTorr
				2.5 kW	4.00 A	2.2 µTorr
				2.0 kW	3.25 A	1.8 µTorr
				1.5 kW	3.25 A	1.6-31 µTorr
				0.9 kW	2.75 A	1.2 µTorr
Lopez Ortega et al. 2020	[2]	Training	I_d, u_{ion}	4.5 kW	6.00 A	9-30 µTorr
				4.5 kW	5.25 A	9-30 µTorr
				2.5 kW	4.00 A	6-30 µTorr
				1.0 kW	2.75 A	3-30 µTorr
Snyder et al. 2024	[11]	Validation	I_d, T_c	4.5 kW	6.00 A	0 µTorr
				3.0 kW	4.50 A	0 µTorr
				2.0 kW	3.25 A	0 µTorr
				1.0 kW	2.75 A	0 µTorr

1. Thruster

This work is focused on extending our existing modeling activities to include the SPT-140 Hall thruster. The SPT-140 is a latter part of the Experimental Design Bureau (EDB) Fakel's SPT portfolio, development of the series beginning in 1964 and continuing in the decades since then [23]. The SPT-140 began to undergo qualification activities in the west around the turn of the millennium [24–26]. As a result of this long heritage, the SPT-140 is not a singular thruster, but rather a series of similar but distinct configurations that have emerged over the years. We focus our modeling activities to the version of the SPT-140 manufactured for, qualified for, and eventually flown on the Psyche spacecraft. This configuration underwent extensive qualification testing on the ground [22] and since being launched on Psyche has been able to return telemetry and performance data from space [11]. JANUS's overriding goal is the capability to confidently extrapolate results from testing thrusters in ground-based vacuum facilities to space. Consequently, the activities for Psyche constitute a key test case for the institute.

2. Training Dataset

As a training dataset, we leverage experiments conducted by NASA's Jet Propulsion Laboratory to qualify the SPT-140 for Psyche [2, 22]. These activities were an extension of previous qualification activities for the SPT-140 [27]; they extended the operational envelope from the nominal 3 to 4.5 kW to a broader 1 to 4.5 kW. All tests used xenon as a propellant and had a discharge voltage of 300 V. The magnetic field profile of the SPT-140 is varied for each power level by adjusting the current energizing its electromagnetic circuit. The coil current ranges from 2.75 A coil current at 1 kW discharge power to 6 A at 4.5 kW. We represent this variation with Eq. (3). The full magnetic field profile and detailed geometry of the SPT-140 used are proprietary to NASA's Jet Propulsion Laboratory and its commercial partners, so are excluded from this report. Snyder et al. conducted a pressure study of the system, operating the thruster at several different power levels and varying the background pressure by flowing additional gas into the facility independent of that supplied to the thruster to achieve the targeted discharge currents [22]. At these operating points, thrust and discharge current (including oscillations) were measured; unfortunately, mass flow rate data were not published for some operating conditions, preventing us from using them as training data. Further data published by Lopez Ortega et al. provide ion velocity profiles for some of these operating conditions [2].

3. Validation Dataset

Psyche launched with four SPT-140's [11]. The flight model of the thrusters was slightly modified relative to the qualification model unit, but these changes were extrinsic to the performance of the thruster [22]. Consequently, we leverage the performance data measured in flight as a validation data set for our model. Radiometric velocimetry conducted during thruster burns provided extremely precise measurements of the thrust produced by the thrusters, and telemetry from the power processing unit provided information about the discharge current. Unfortunately, the mass flow rate supplied to the thruster is not reliably metered on orbit over short timescales, so when we make predictions to the space environment we assume a mass flow rate consistent with the minimum-pressure testing conditions achieved in ground testing [22]. This is consistent with the choice made by Lopez Ortega et al. in modeling the SPT-140 in the 2-D axial-radial code Hall2De [2]. The flow system controls to a target discharge current, not a fixed mass flow rate, however, a point we return to in the discussion (Sec. IV.C). Finally, we model the background pressure of space as an arbitrarily small value ($P_B \to 0$).

C. Calibration

The Hall thruster PEM architecture [4, 6] includes utilities for training the model, which we leverage here. We express the model training as a Bayesian parameter estimation problem, wherein we infer the values of the model parameters given some set of experimental data. Denoting generically the parameters as Θ and the data as D, we compute the posterior distribution

$$\mathrm{P}(\Theta \mid D) = \frac{\mathrm{P}(D \mid \Theta)\mathrm{P}(\Theta)}{\mathrm{P}(D)} \propto \mathrm{P}(D \mid \Theta)\mathrm{P}(\Theta). \tag{11}$$

The right hand side of Eq. (11) results from Bayes' rule, with $\mathrm{P}(D \mid \Theta)$ the likelihood and $\mathrm{P}(\Theta)$ the prior. The $\mathrm{P}(D)$ factor is the evidence and is merely a normalizing constant for our purposes, absorbed into the proportionality. By modeling the parameters Θ probabilistically, we account for uncertainty in their value arising from the fact that the data D are acquired by methods with finite error or representing that our model may incompletely capture the physics. In our case, Θ is the set of model parameters indicated in Tab. 1, namely

$$\Theta = \left\{V_{vac}, T_{ec}, P_T, P^*, \alpha_{anom}, \beta_{anom}, \hat{z}_{anom}, \hat{L}_{anom}, \Delta\hat{z}_{anom}, u_n, c_w, f_n, c_0, c_1, c_2, c_3, c_4, c_5\right\}. \tag{12}$$

The likelihood $\mathrm{P}(D \mid \Theta)$ is a model for how the data were generated (i.e., the underlying stochastic process representing error in the measurement). In our case, we have three streams of data, thrust measurements and discharge current measurements from [22] and ion velocity profile measurements from [2]. However, not all data are available at all such operating conditions. We define a distinct operating condition at which the thruster was operated by the set of model inputs corresponding to that operating condition, $\mathbf{d} = \left\{V_d, P_b, \dot{m}_a, \hat{B}_r\right\}$. Let n_{I_d}, n_{T_c}, and $n_{u_{ion}}$ represent the number of operating conditions where each kind of data was collected, and let j index those operating conditions, such that $\mathbf{d}_{I_d,j}$ is the j^{th} operating condition at which discharge current data were collected and so on. While discharge current and thrust are represented as singular measurements taken at an operating condition, ion velocity measurements are a collection of multiple data taken at different z locations. We denote an individual measurement location for operating condition j as $z_{j,l}$, and the number of such measurement locations m_j. Consequently, we define $y_{I_d,j}$ and $y_{T_c,j}$ to be the measured values of discharge current and at operating condition j and $y_{u_{ion},j,l}$ to be the measurement of ion velocity at operating condition j and axial location l. The corresponding model predictions for these quantities we label as $f_{I_d}(\mathbf{d}_{I_d,j};\Theta)$, $f_{T_c}(\mathbf{d}_{T_c,j};\Theta)$, and $f_{u_{ion}}(\mathbf{d}_{u_{ion},j}, z_l;\Theta)$, where we have explicitly indicated the functional dependence on the inputs and parameters of the model.

Subject to these definitions, the PEM models the likelihood as

$$\mathrm{P}(D \mid \Theta) \propto \exp\left[-\frac{1}{2}\frac{\sum_{j=1}^{n_{I_d}} \left(f_{I_d}(\mathbf{d}_{I_d,j};\Theta) - y_{I_d,j}\right)^2}{\gamma_{I_d}^2 \sum_{j=1}^{n_{I_d}} y_{I_d,j}^2}\right] \times \exp\left[-\frac{1}{2}\frac{\sum_{j=1}^{n_{T_c}} \left(f_{T_c}(\mathbf{d}_{T_c,j};\Theta) - y_{T_c,j}\right)^2}{\gamma_{T_c}^2 \sum_{j=1}^{n_{T_c}} y_{T_c,j}^2}\right] \times \exp\left[-\frac{1}{2}\frac{\sum_{j=1}^{n_{u_{ion}}} \sum_{l=1}^{m_j} \left(f_{u_{ion}}(\mathbf{d}_{u_{ion},j},\, z_l;\Theta) - y_{u_{ion},j,l}\right)^2}{\gamma_{u_{ion}}^2 \sum_{j=1}^{n_{u_{ion}}} \sum_{l=1}^{m_j} y_{u_{ion},j,l}^2}\right]. \tag{13}$$

We note that this closely follows the notation of [6], except where we have explicitly denoted the different data streams symbolically rather than enumerating them. Rather than model error in individual measurements, the PEM models error over entire streams of data, which are assumed independent—hence the three exponential terms in Eq. (13). In particular, the L_2 norm of the error between the model prediction and measurements (numerator) is modeled as a half-Gaussian distribution. The widths of these distributions is expressed as some fraction γ of the L_2 norm over the measurements, such that the width represents a relative error. Modeling the likelihood this way more equally weights different types of data (discharge current vs. thrust vs. ion velocity) in estimating the model parameters, rather than equally weighting individual measurements—which tends to be dominated by the multiplicity of ion velocity measurements [6]. For this study, we model $\gamma_{I_d} = \gamma_{T_c} = \gamma_{u_{ion}} = 0.05$, that is, a 5% relative error.

The prior distribution $\mathrm{P}(\Theta)$ represents any cogent information about the parameters that may result from previous studies, enforce physical constraints, or embody intuition. Within the PEM framework, individual thrusters are treated as unique systems that are not necessarily informative of each other, and we have not previously trained to model the SPT-140. We therefore elect the "uninformative" distributions of Tab. 3, which are the same as used in [6] for the SPT-100 thruster. The symbol $\mathcal{U}$ is used to denote the uniform distribution, and parameters preceded by a $\log_{10}$ indicate that the distribution has been prescribed over the logarithm of that parameter (because they may vary over several orders of magnitude). Units are indicated for dimensioned quantities.

Table 3 Prior distribution over the model parameters $\mathrm{P}(\Theta)$

Parameter	Distribution	Parameter	Distribution	Parameter	Distribution
V_{vac}	$\mathcal{U}(0, 60)$ V	α_{anom}	$\mathcal{U}(0, 1)$	c_0	$\mathcal{U}(0, 1)$
T_{ec}	$\mathcal{U}(1, 6)$ eV	β_{anom}	$\mathcal{U}(0, 1)$	c_1	$\mathcal{U}(0.1, 0.9)$
P_T	$\mathcal{U}(10, 100)$ µTorr	$\hat{z}_{anom}$	$\mathcal{U}(0.75, 1.5)$	c_2	$\mathcal{U}(-15, 15)$ mrad mPa^{-1}
P^*	$\mathcal{U}(10, 200)$ µTorr	$\hat{L}_{anom}$	$\mathcal{U}(0, 0.5)$	c_3	$\mathcal{U}(0.2, \pi/2)$ rad
		$\Delta\hat{z}_{anom}$	$\mathcal{U}(0, 0.5)$	$\log_{10} c_4$	$\mathcal{U}(18, 22)$ $\log_{10}$ m^{-3} Pa^{-1}
		u_n	$\mathcal{U}(100, 500)$ m s^{-1}	$\log_{10} c_5$	$\mathcal{U}(14, 18)$ $\log_{10}$ m^{-3}
		c_w	$\mathcal{U}(0.5, 1.5)$		
		f_n	$\mathcal{U}(1, 10)$		

Having given the likelihood in Eq. (13) and the prior in Tab. 3, our inference problem is fully specified. The irregular and multidimensional nature of the resulting posterior, however, renders it analytically intractable (i.e., to derive statistics). Rather, then, we sample from the posterior using Markov Chain Monte Carlo techniques, generating 15000 samples from the chain. The Markov chains are initialized at a nominal set of parameters and take some time to converge to the posterior distribution, so we discard the first half of the generated samples. We can then generate model predictions accounting for uncertainty in the parameters by evaluating the model at these samples. Note that in previous work we also treated the effect of uncertainty in the operating conditions themselves [6]—i.e., that the discharge voltages, background pressures, and other variables we treat as independent are not known precisely either. We found this to be a small contribution to overall uncertainty in our predictions and so neglect it here.

III. Results

In this section, we present the results of our model calibration, including the posterior parameter values inferred from the data (Sec. III.A), the corresponding model predictions when compared to the training data set (Sec. III.B), and an extrapolation to the space environment using the calibrated model (Sec. III.C).

A. Posterior Distribution

We first examine statistics over the samples we drew in training the model of Sec. II.A on the data indicated in Sec. II.B. We summarize these in Tab. 4, where we provide the *maximum a posteriori* (MAP) parameter values over the whole distribution as well as the marginal mean, standard deviation, 5%-ile, median (50%-ile), and 95%-ile parameter values over the samples. We also reiterate the prior distributions of Tab. 3 and the median parameters calibrated for the SPT-100 and H9 thrusters in [6] as comparisons. This is a necessary contrast because it indicates which parameters are informed by the training data—have a posterior distribution dissimilar from their prior. In our case, all the model parameters have posterior distributions narrower than their priors, indicating the data is at least somewhat informative for all parameters.

Table 4 Marginal statistics for different model parameters computed from posterior samples, contrasted with their prior distribution and previous results from [6].

Parameter	Units	Prior	MAP	Mean	Std. Dev.	5%	Median	95%	SPT-100	H9
V_{vac}	V	$\mathcal{U}(0,60)$	15.94	16.87	1.28	15.21	16.66	19.32	31.75	21.94
T_{ec}	eV	$\mathcal{U}(1,6)$	4.39	4.70	0.51	3.88	4.66	5.58	2.92	5.40
P_T	µTorr	$\mathcal{U}(10,100)$	95.8	94.5	3.7	87.1	95.1	99.3	48.7	3.2
P^*	µTorr	$\mathcal{U}(10,200)$	88.7	87.9	3.7	80.6	88.5	93.0	64.9	45.3
α_{anom}		$\mathcal{U}(0,1)$	0.089	0.083	0.035	0.040	0.075	0.151	0.06	0.13
β_{anom}		$\mathcal{U}(0,1)$	0.987	0.973	0.023	0.923	0.979	0.996	0.99	0.98
$\hat{z}_{anom}$		$\mathcal{U}(0.75,1.5)$	0.922	0.928	0.043	0.859	0.926	1.001	1.14	1.07
$\hat{L}_{anom}$		$\mathcal{U}(0,0.5)$	0.387	0.341	0.094	0.193	0.343	0.483	0.43	0.43
$\Delta\hat{z}_{anom}$		$\mathcal{U}(0,0.5)$	0.057	0.152	0.109	0.011	0.134	0.360	0.33	0.18
u_n	$\mathrm{m\,s^{-1}}$	$\mathcal{U}(100,500)$	249.6	250.9	3.3	246.8	250.2	257.9	278.1	268.9
c_w		$\mathcal{U}(0.5,1.5)$	0.68	0.74	0.20	0.52	0.68	1.16	0.67	1.19
f_n		$\mathcal{U}(1,10)$	4.81	4.84	1.67	2.01	4.92	7.45	5.23	3.85
c_0		$\mathcal{U}(0,1)$	0.136	0.47	0.28	0.05	0.46	0.94	0.76	0.32
c_1		$\mathcal{U}(0.1,0.9)$	0.35	0.62	0.20	0.21	0.66	0.88	0.32	0.39
c_2	$\mathrm{mrad\,mPa^{-1}}$	$\mathcal{U}(-15,15)$	-12.75	-12.2	1.4	-14.6	-12.2	-9.6	-12.36	2.71
c_3	rad	$\mathcal{U}(0.2,\pi/2)$	0.248	0.308	0.068	0.207	0.304	0.422	0.21	0.332
$\log_{10} c_4$	$\log_{10}\mathrm{m^{-3}\,Pa^{-1}}$	$\mathcal{U}(18,22)$	18.14	18.41	0.32	18.03	18.35	19.09	20.33	20.15
$\log_{10} c_5$	$\log_{10}\mathrm{m^{-3}}$	$\mathcal{U}(14,18)$	14.07	14.28	0.23	14.02	14.22	14.71	14.33	14.26

For the cathode coupling model, we find that V_{vac} has converged to a value of approximately 16.5 V, and P_T and P^* have settled at values approximately 100 µTorr. These large pressures all but eliminate any dependence of the cathode coupling voltage on background pressure per Eq. (1), meaning the model has effectively learned that there exists a fixed cathode coupling voltage of 16.5 V which is insensitive to pressure. While we lacked direct measurements of cathode coupling voltage to train the model against, Snyder et al. did report limited cathode to ground measurements, approximately 20 V for most power conditions [22]. The cathode to ground voltage tends to track the cathode coupling voltage because the beam plasma more readily couples to the facility, such that whatever potential drop exists from the beam to the cathode drives it farther below ground. This would suggest that while the magnitude of the cathode coupling voltage we obtained is modestly off, we have correctly captured an insensitivity to pressure. However, the simulations performed by Lopez Ortega et al. as part of the qualification activities predict that the beam potential near the thruster rises with respect to the cathode potential as background pressure decreases [2], which would instead suggest that the cathode coupling voltage should have a background pressure dependence. However, these potentials were evaluated less than one channel length downstream of the thruster exit, and so may not be fully representative of the far-field beam potential.

We find that the parameter values inferred for the thruster model (second column of Tab. 3) are broadly consistent with our previous results on the SPT-100 and H9 Hall thrusters [6]. In particular, we observe that the anomalous mobility scale α_{anom} has a magnitude about 0.1, the transport barrier scale β_{anom} reduces the mobility by approximately 98-99%, the transport barrier is manifest approximately coincident with the exit plane of the thruster ($\hat{z}_{anom} \approx 1$), and the width of the barrier is just under half a channel length ($\hat{L}_{anom} = 0.34$). While the shift in the transport barrier found here ($\Delta\hat{z}_{anom} = 0.13$) is comparable to prior study as well, the breadth of the distribution—a 90% credible interval between 0.011 and 0.360—indicate that there remains substantial uncertainty in how sensitive shifts in the acceleration region are to pressure.

The similarity of the thruster model parameters between different thrusters indicates that some fundamental features are common enough that one can obtain a decent first approximation of the discharge from a set of common parameter values. This similarity partially derives from underlying Hall thruster design principles, however. That is, they may not reflect fundamental physics as much as they reflect that different systems tend to converge in design. For example, the peak magnetic field strength tends to be approximately coincident with the exit plane of the thruster. This choice is driven by the ion optics, a balance of divergence in the beam and erosion of the channel [13, 28–30]. Consequently, resolving high-order phenomena is likely to remain an individualized process.

For the plume model parameters, we find that the model only learns anything conclusive about c_2. This scaling coefficient controls how much the beam is expected to collimate as a result of increases in background pressure, with the median value of$-12.2\,\mathrm{mrad\,mPa^{-1}}$ equivalent to $-1.62\,\mathrm{mrad\,\mu Torr^{-1}}$. This is a noticeable trend—from vacuum to $30\,\mathrm{\mu Torr}$ the total predicted change in divergence is 2.8 deg. This shift is close in magnitude to that of the SPT-100 ($-12.36\,\mathrm{mrad\,mPa^{-1}}$, but differs from that of the H9 (2.71). The SPT-100 is directly in family to the SPT-140, and both have unshielded magnetic typologies, as opposed to the magnetically-shielded H9 [30].

The values for c_0, c_1, and c_2, by contrast, vary more widely (standard deviations 60%, 32%, and 22% of their respective means). At face value, this suggests that the model is quite unsure about the divergence in the beam. However, we note that these parameters are highly correlated; that is, because we lack current density data to inform the detailed shape of the beam, there is great uncertainty about whether there is a large main beam and little scattered beam at a higher base beam width or a small main beam and a large scattered beam at a lower base beam width. The divergence angle itself, computed via Eq. (9), is more precisely determined than the marginal uncertainties in Tab. 4 imply.

Finally, the parameters c_4 and c_5 appear to have been inferred precisely but have been limited by the prescription of the prior (taking values close to the minimum allowed). Since per Eq. (9) the charge exchange current content is not used to compute thrust, the $\exp\left(-r n_n \sigma_{cx}\right)$ factor in Eqs. (7b) and (7c) cancels out, removing any dependence of the thrust calculation on c_4 and c_5 (which are used to compute n_n). Similarly, because the plume computation does not feed back into the cathode coupling or thruster models (see Fig. 1 and Tab. 1), parameters c_4 and c_5 have no effect whatsoever on predicting a quantity of interest we can compare to experimental data as the model was not trained on any ion current density data. Thus, the values of c_4 and c_5 are meaningless and we believe they have been limited to their minimum value in the prior as a numerical artifact.

B. Comparison to Training Data

In this section, we compare the training data used to calibrate the model to the model's predictions at the same operating conditions. Doing so is an assessment of goodness of fit—determining whether the model is sufficiently complex to reproduce the experimental results. We summarize the error of the posterior predictions relative to the training data in Tab. 5. Here, we compute the error E as relative L_2 norms over all the data for the different quantities of interest, i.e.,

$$E_{I_d}(\Theta) = \sqrt{\frac{\sum_{j=1}^{n_{I_d}} \left(f_{I_d}(\mathbf{d}_{I_d,j};\Theta) - y_{I_d,j}\right)^2}{\sum_{j=1}^{n_{I_d}} y_{I_d,j}^2}}, \tag{14a}$$

$$E_{T_c}(\Theta) = \sqrt{\frac{\sum_{j=1}^{n_{T_c}} \left(f_{T_c}(\mathbf{d}_{T_c,j};\Theta) - y_{T_c,j}\right)^2}{\sum_{j=1}^{n_{T_c}} y_{T_c,j}^2}}, \tag{14b}$$

$$E_{u_{ion}}(\Theta) = \sqrt{\frac{\sum_{j=1}^{n_{u_{ion}}} \sum_{l=1}^{m_j} \left(f_{u_{ion}}(\mathbf{d}_{u_{ion},j},\, z_l;\Theta) - y_{u_{ion},j,l}\right)^2}{\sum_{j=1}^{n_{u_{ion}}} \sum_{l=1}^{m_j} y_{u_{ion},j,l}^2}}, \tag{14c}$$

which follows the forms we used in defining the likelihood of Eq. (13). We compute each of these error metrics for every sample of the posterior. For each quantity of interest—each of the subequations of Eq. (14)—we then compute the mean (μ_E) and standard deviation (σ_E) over these samples. We tabulate these calculations in Tab. 5, where the

Table 5 Relative L_2 errors E computed over the posterior predictions for the different quantities of interest in the data.

Quantity	Thruster	γ [%]	Mean μ_E [%]	Std. Dev. σ_E [%]
T_c	**SPT-140 (this work)**	5.0	7.4	1.6
	SPT-100 ([6])	1.0	3.5	0.5
	H9 ([6])	-	-	-
I_d	**SPT-140 (this work)**	5.0	15.0	2.4
	SPT-100 ([6])	10.0	3.9	1.4
	H9 ([6])	10.0	4.3	1.1
u_{ion}	**SPT-140 (this work)**	5.0	13.2	2.3
	SPT-100 ([6])	5.0	13.8	1.3
	H9 ([6])	5.0	5.3	1.0

horizontal lines separate calculations for the different predicted quantities. For comparison, we also tabulate the same calculation for the posterior predictions on the SPT-100 and H9 thruster in our previous work with the PEM [6], and we provide the uncertainty assumed in the data γ as a reference.

The mean error varies between the three variables, from 7% for thrust to 15% for discharge current. We find the error in the ion velocity predictions comparable to our previous results on the SPT-100, but the thrust and discharge current predictions show increased error against the training dataset. The thrust and discharge current errors are larger because here we have trained and predicted against multiple different power levels, where previously we had trained at only a single power. Thrust and discharge current vary more strongly as a function of power than do the ion velocity profiles. We highlight these trends in more detail in the following sections.

1. Thrust and Discharge Current

We plot the training data for thrust from [22] as black circles in Fig. 3. The bars indicate the experimentally-reported measurement error (very small on the scale of the figure). The training data represent multiple data series, including two pressure studies conducted at 1.5 and 4.5 kW discharge (two horizontal lines of circles) and a list of points at different powers but all at the minimum achievable operating pressure (slanted series of circles toward the left). See also Tab. 2. These data demonstrate that increased background pressure in the facility causes an increase in thrust, attributed to ingestion of facility gas as propellant and increased collimation of the ion beam [22, 31, 32].

We contrast the training data with the predictions at the same operating conditions rendered by the trained model, which we plot as red circles. We evaluated the model for each sample of the parameters drawn during model training (see Sec. II.C). The circle is the median prediction, while the error bars represent the 5% and 95% quantiles—a 90% credible interval. While the model is able to reproduce approximately one third of the thrust training data within uncertainty, we observe two key discrepancies. First, the model tends to under-predict thrust at higher power and over-predict thrust at lower power, and second, the model tends to under-predict the change in thrust observed as a function of background pressure. The latter trend is a likely consequence of the first—the inference prioritizes trying to match how thrust scales with power (a strong effect) over capturing how it scales with pressure (a weaker effect).

We find that this trend scaling with power is also evident in the model predictions for discharge current, which we render against the training data in Fig. 4. Note in the figure that the discharge current training data are flat with background pressure: during experiments the mass flow rate supplied to the thruster was adjusted as a function of background pressure to hold discharge current constant [22]. While we find that for the 4.5 kW data (15 A discharge current in the figure) the model predicts the discharge current within uncertainty, the 1.5 kW data (5 A discharge current) tend to be over-predicted by a factor 1.5. This discrepancy is driven by an overly large electron current predicted by the thruster model, which we revisit in Sec. IV.A.

2. Ion Velocity

We find that the model is generally able to reconstruct the training ion velocity profiles within uncertainty. As a representative illustration, we compare ion velocity data taken from [2] to model predictions in Fig. 5. These data are

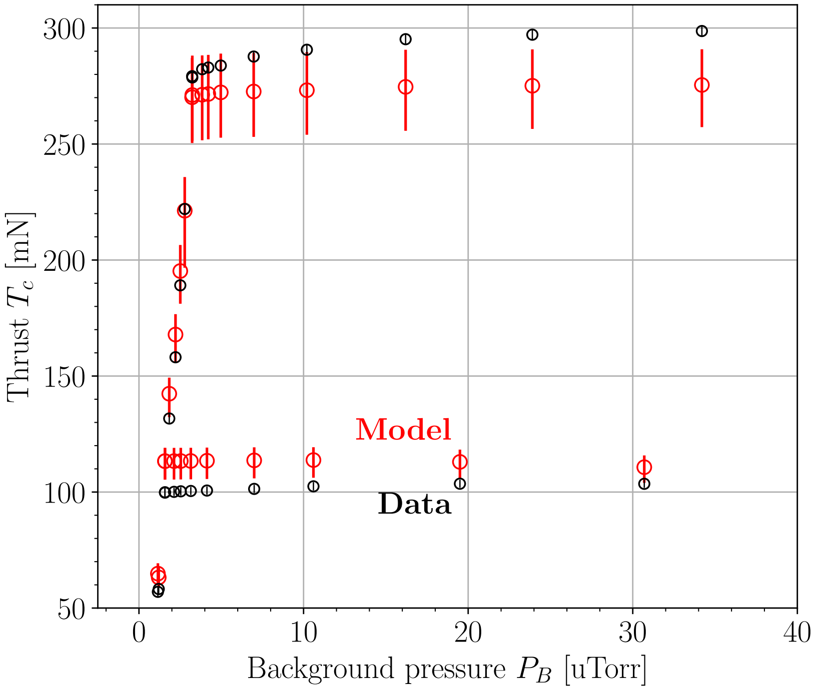

Fig. 3 Thrust training data (black circles) reported in [22] and model predictions (red circles) as a function of facility pressure.

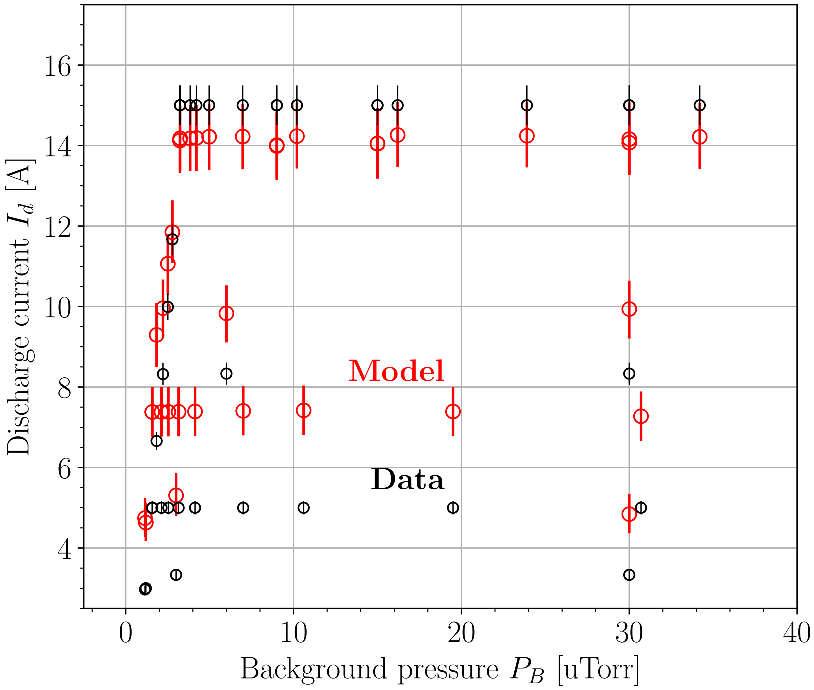

Fig. 4 Discharge current training data (black circles) reported in [22] and [2] model predictions (red circles) as a function of facility pressure.

for discharges at 4.5 kW and with a magnet current of 6 A. We render the training data as individual points and plot the medians and 90% credible intervals over our predictions as solid lines and shaded regions, respectively. Figure 5a renders a comparison for 15 µTorr background pressure over the entire domain of the simulations, while Fig. 5b focuses in on the acceleration region and renders multiple operating pressures.

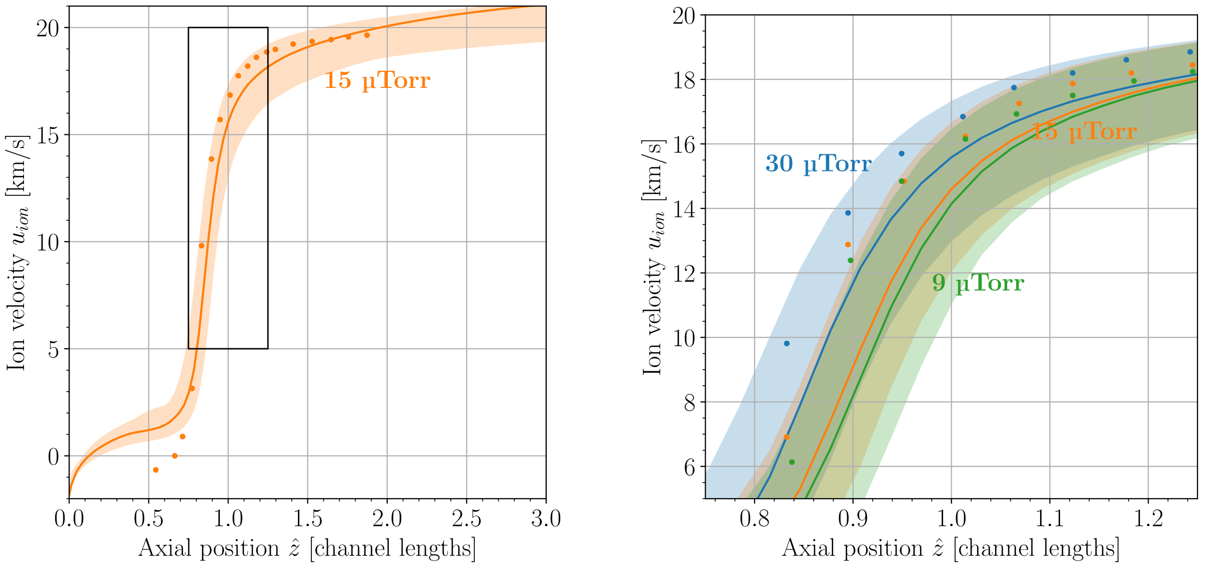

(a)

(b)

Fig. 5 Comparison of measured (points, [2]) and predicted (lines and shaded regions) ion velocity profiles at 4.5 kW and 6 A coil current: (a) single profile at 15 µTorr **background pressure; (b) highlighted region of (a) at all measured background pressures**

In Fig. 5a, we observe that the training data are reconstructed within uncertainty, accurately capturing the acceleration of the ion beam over a narrow region and the more gradual saturation of ion velocity thereafter. The only exception is the low-velocity data taken between 0.5 and 0.75 channel lengths downstream of the anode ($\hat{z} = 0$), which is commonly the case for the thruster model [6] and is likely related to the simplified form of the anomalous collision frequency in Eq. (4) (cf. the more complex, individually tuned profiles of [2]). Further, in Fig. 5b we observe that the data exhibits a trend where this acceleration region moves farther upstream as the background pressure rises. Consistent with Eq. (6), the model retrieves this trend, including that the shift between 9 µTorr and 15 µTorr is smaller in magnitude than the shift between 15 µTorr and 30 µTorr.

The recession of the acceleration region tends to better collimate the ion beam, decreasing divergence losses according to Eq. (9) and increasing thrust, which we see reflected in the predictions of Fig. 3. The location of the transport barrier's nadir $\hat{z}_{anom}$ controls where the acceleration is manifest, because the increased resistance to electron transport is associated with a strong gradient in potential which localizes the acceleration of the ions [3]. That the model correctly localizes the transport barrier (as judged by the ion velocity profiles) but incompletely captures the magnitude of electron current (as judged by the discharge current) suggests that it is the magnitude of the transport barrier which varies as a function of power.

3. Cathode Coupling and Plume

For completeness' sake, we render in Fig. 6 a plot of the predicted cathode coupling voltage for each operating condition modeled as a function of pressure. We observe that indeed the model predicts only minute changes in cathode coupling voltage with pressure, as we had observed in Sec. III.A. We also find that all predictions (even differing powers and magnetic field configurations) track along a single line. This is because the cathode coupling model, Eq. (1), only includes a pressure dependency, and does not model a dependence on the different flow rates the result from different powers.

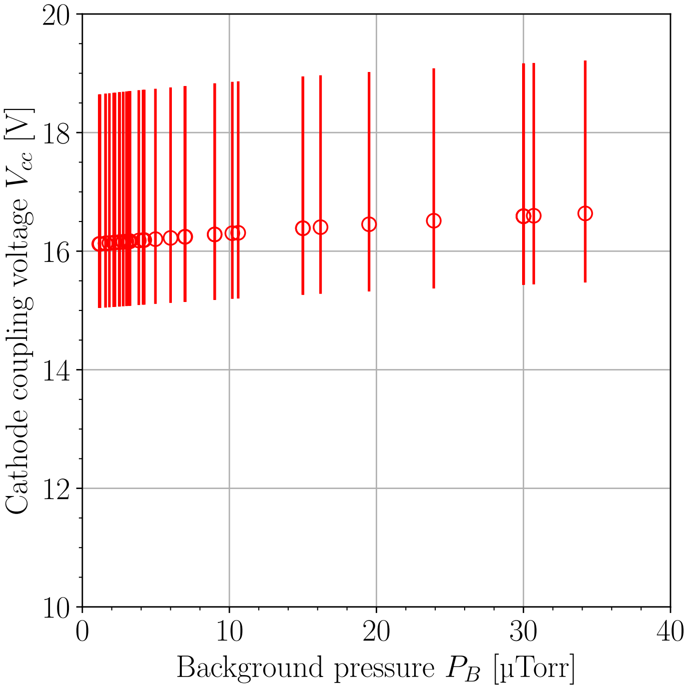

Fig. 6 Predicted cathode coupling voltage for the training data as a function of pressure; consistent with previous plots, we render the predictions as red circles, and we have no experimental data to compare against.

Similarly, we include a plot illustrating the predicted pressure dependence on the ion current density profiles in Fig. 7. Here we render three different pressure conditions all at 4.5 kW discharge and 6 A magnet current. Comparable to our ion velocity profiles, the solid lines denote the medians while the shaded regions are 90% credible intervals. All of these profiles are resolved at a distance 1 m from the thruster exit. The characteristic broadening of the beam with decreasing background pressure is evident here by the shift to larger θ of the main lobe of the distributions. We also observe that the asymptotic floor of the distribution—which is formed by the hemispherical expansion of charge exchange ions per Eq. (7d)—lessens as the background gas thins and the corresponding neutral density n_n decreases. Consistent with our observations in Sec. III.A, we see that the credible interval places great uncertainty on the precise shape of the beam. Considering the 3 µTorr data, we see that the centerline ion current density varies between 20 to 100 A m^{-2}, and the final attenuation of the main and scattered beams to the charge exchange background could vary anywhere from about 40 to 80 degrees off centerline. However, the curves tend to coalesce approximately 20 degrees off centerline, reflecting a key constraint on the beam: the integrated current must sum to the beam current per Eqs. (7). Consequently, the additional current content of a fatter tail must be offset by lower current content in the main beam, and these countervailing forces tend to constrain the intermediate region at the "shoulder" of the beam.

C. Extrapolation to Space

The availability of thruster measurements in space from Psyche is a key potential validator of the model. Consequently, we predicted performance using the trained model for the space operating conditions of [11] consistent with the methods of Sec. III.B.1. We predicted the model at the four operating conditions listed in Tab. 2, for 750 posterior samples of the model parameters Θ, generating predictions with uncertainty. Crucially, the model was not trained on these data. We render a response plot of our predictions against the space data in Fig. 8. The error bars along the predicted thrust axis represent a 90% credible interval, and the dashed line tracks a perfect response (predicted thrust matches measured thrust).

For the measured thrust, the dominant source of uncertainty was in the reproducibility of the system itself. Psyche has four different SPT-140 thrusters. During calibration activities on-orbit, accurate measurements of thrust (error

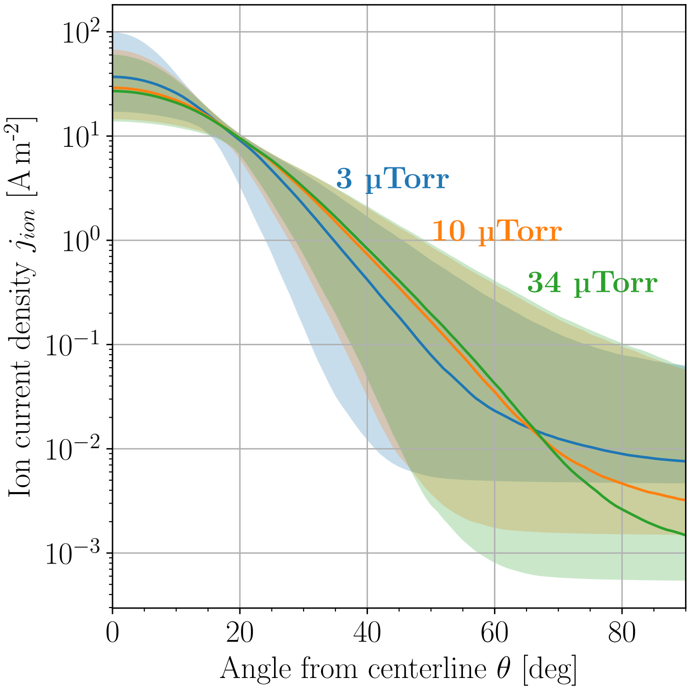

Fig. 7 Predicted ion current density profiles for 4.5 kW, 6 A magnet current operations at different pressures; solid lines are the median predictions, and shaded regions contain a 90% credible interval.

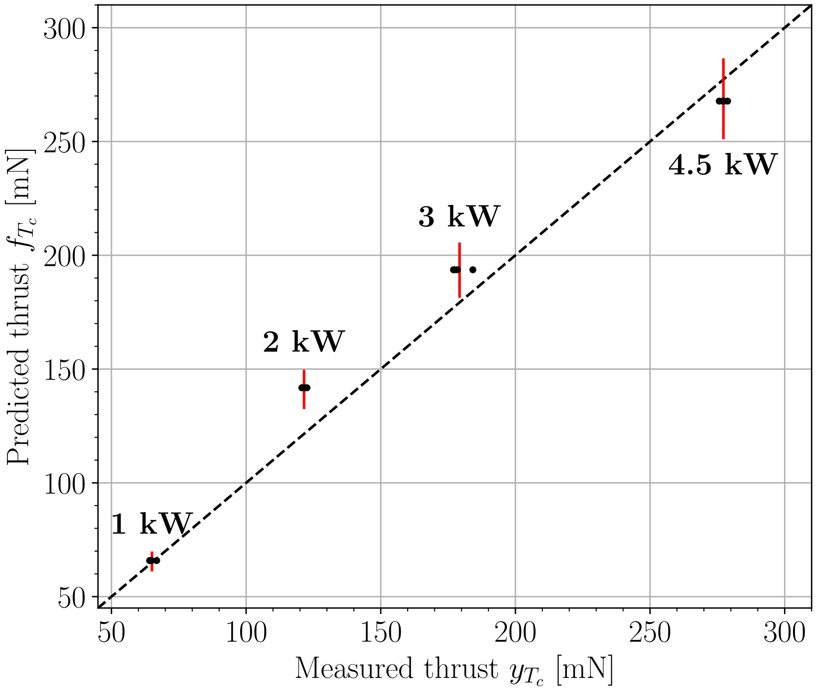

Fig. 8 Response plot of predicted space performance to on-orbit data of [22]; measurements of individual thrusters are rendered as red dots and used to compute the horizontal error bars.

of order 0.2 mN) were obtained for each thruster individually. However, these measurements revealed variations in performance between the four thrusters of order 1-2 mN, likely the result of small differences in manufacturing and fabrication [11]. Consequently we represent this error in Fig. 8 by plotting the individual measurements for each thruster as black dots. Overall, we observe fair agreement with the space data, with two of the four conditions matching within 90% uncertainty. This is consistent with our predictions against the training data in Fig. 3, where the model was able to more accurately reproduce data at lower background pressure.

IV. Discussion

We reexamine our results in this section, discussing their implications for the underlying physics, broader development efforts, and future work. In Sec. IV.A, we consider how additional scaling in our anomalous mobility model might improve predictive power and reconcile our predictions better with data. In Sec. IV.B, we recalibrate the model on only higher pressure data to simulate challenges that might be encountered for testing high power systems. Finally, in Sec. IV.C, we compute model predictions that invert the dependency between discharge current and mass flow rate to more closely correspond to how experimental data are collected.

A. Anomalous Mobility Power Scaling

In Sec. III.B.1, we had observed that the model tended to under-predict thrust and discharge current at higher power and over-predict thrust and discharge current at lower power (see also Figs. 3 and 4). On this basis, we hypothesized that this discrepancy was rooted in how the thruster model treats electron mobility via an anomalous collision frequency—Eq. (4). The discharge current is the sum of the ion current exiting the thruster and the parasitic electron current drawn from the cathode to the anode. The maximum possible ion current is limited by the propellant mass flow rate, thus setting a lower bound on the electron current. For the 1.5 kW tests of Fig. 4 (lower horizontal data series), for example, the anode mass flow rate was approximately 6.5 mg/s, which for the xenon used would correspond to an ion current of 4.8 A assuming singly charged ions and full ionization of the propellant. Thus, the predicted value of 7.8 A must be composed of at least 3 A of electron current.

The electron current is chiefly driven by the anomalous electron collisions of Eq. (4), with the magnitude at the nadir of the profile (i.e., $\beta_{anom}\alpha_{anom}$ as indicated in Fig. 2) most primarily serving to throttle this current and determine its magnitude. We observe that the discrepancy between the data and model predictions for discharge current in Fig. 4 is nearly a constant factor across data taken at the same power. Combined with our conclusion that the shift in location of the acceleration region is well captured by the model (see Sec. III.B.2), this suggests there exists some functional dependence of the scale factors for the anomalous mobility model on power that is not presently included in Eq. (4).

Some dependence is captured implicitly through the magnetic field scaling of Eq. (3). The coil current set point of the SPT-140 for Psyche increases along with power, from 2.75 A of coil current at 1 kW power to 6 A at 4.5 kW [22]. The concomitant increase in magnetic field discourages cross-field electron transport, decreasing discharge current for a fixed mass flow rate. Fig. 4 indicates, then, that this scaling alone is incomplete, with mobility too low at higher powers (too little discharge current) and mobility too high at lower powers (too much discharge current).

That this degree of tailoring the anomalous collision frequency is necessary is unsurprising. Such anomalous collision frequency profiles are often tuned to ion velocity data for individual operating conditions [1, 33], rather than trying to infer some trend. Indeed, that is exactly what [2] did in making their predictions with the high fidelity Hall2De code, which we highlight against ours and the data in Fig. 9. Here, we have adopted the convention that the data the model was trained on are rendered in circles, while validation data not used to train the model are rendered as squares. The tack of tuning the anomalous collision profile to each simulation individually has in part provided the Hall2De predictions close correspondence to the data.

These individually calibrated profiles demonstrate the trend encoded by Eq. (6) that the acceleration region moves upstream with increased background pressure. More critically, these tuned profiles also show that the magnitude of the anomalous collision frequency scales with power as well. We focus our attention on the minimum anomalous collision frequency in these curves—the equivalent of the nadir of the Gaussian dip in Eq. (4) and Fig. 2. We plot in Fig. 10 the minimum value of anomalous collision frequency in the tuned profiles of [2] for different operating conditions as a function of coil current at that operating condition. We find there to be only minor differences in magnitude between different background pressures at the same discharge power and coil current settings, so we consider only the 30 µTorr results (which is the only background pressure tested at all powers).

Fig. 10 demonstrates that the magnitude of the anomalous collision frequency is in fact linear with coil current (and hence magnetic field). We include the corresponding line of best fit—in a least squares sense—in the figure alongside

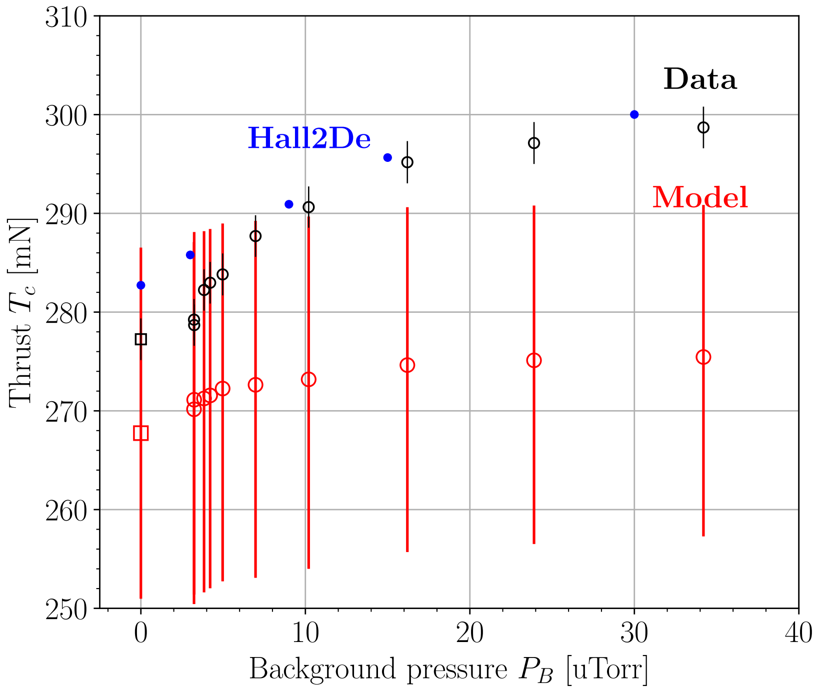

Fig. 9 Model predictions (red) and data from [22] and [11] (black) for thrust at 4.5 kW, including both training (open circle) and validation (square) datasets; the solid blue points indicate Hall2De simulations from [2] that are calibrated on individual operating conditions

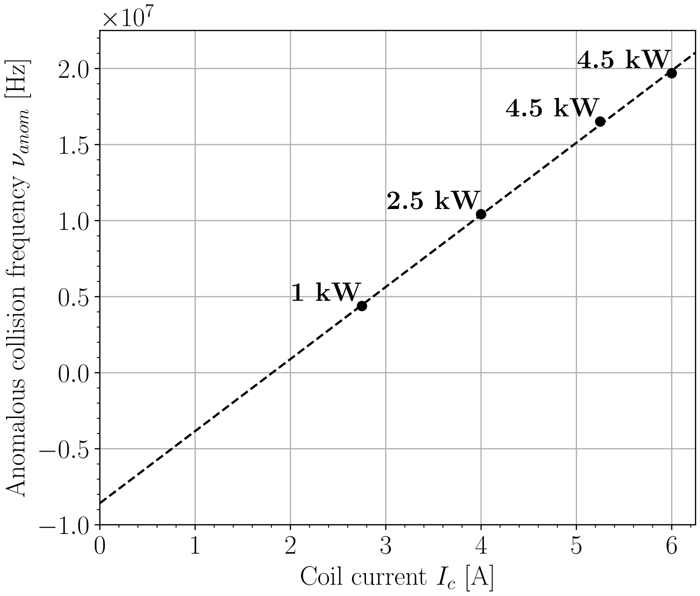

Fig. 10 Minimum anomalous collision frequency from the calibrated profiles of [2] at 30 μTorr background pressure (points) and a linear fit to those data (dashed line).

the data. This is puzzling, because it suggests that the scaling of our anomalous collision frequency model Eq. (4) should be able to capture this trend, with the values of α_{anom} and β_{anom} calibrating to match the magnitude and then the ω_{ce} scaling capturing the increase with magnetic field. We find resolution in that the best-fit line of Fig. 10 has a substantial offset. This offset is equivalent to a constant negative term added to the collision frequency which does not scale with magnetic field strength, explaining why Eq. (4)—which lacks such a term—was not able to capture it.

We hypothesize that adding a constant term to Eq. (4) that could be calibrated from data would be able to reconcile the discrepancy in discharge current over the data set seen in Fig. 4. Subtracting a constant from the anomalous collision frequency would tend to shift all the discharge current predictions down—because electron mobility would be reduced according to Eq. (5). Subsequently, the value of α_{anom} or β_{anom} could be calibrated to rescale the discharge current to match. We aim to incorporate this expanded scaling into `HallThruster.jl`.

B. Calibration Restricted to High Pressures

One of the central reasons why JANUS seeks to quantify and predict facility effects is the development of higher power, next generation propulsion systems. High power systems (order 100 kW) produce additional strain on vacuum test facilities. In particular, the higher flow rates of propellant cannot as adequately be evacuated, leading to higher background pressures (compare with Fig. 3, where the leftmost data illustrate the increase in minimum operating pressure as power increases). Consequently, a key concern is that the pressure may not be low enough to resolve the pressure dependence of performance, preventing a confident extrapolation to space. Consider, for example, the nonlinearity in the thrust measurements of Fig. 3, where data at higher pressure are nearly constant within experimental uncertainty; if only these data were available, the model might have greater uncertainty in extrapolating to the sharper change in thrust at lower pressures. To simulate this potential inadequacy of data, we performed two additional calibrations of the model. We followed the same procedure as in Sec. II.C, except that one was calibrated only on experimental data taken at background pressures above 5 µTorr and the other only on experimental data taken at background pressures above 20 µTorr.

We plot the thrust data and predictions for this exercise in Fig. 11. As before, black are measurements, red are predictions, circles denote the data the model was trained on, and squares denote data the model was not trained on; we emphasize the pressure cutoffs for training with vertical dashed lines. We also present in Fig. 12 ion velocity predictions for the 4.5 kW, 6 A magnet current power operating condition as a function of pressure, as in Fig. 5b. We have limited the plots to the same domain, but instead of rendering the 8 µTorr condition we present black curves of our predictions

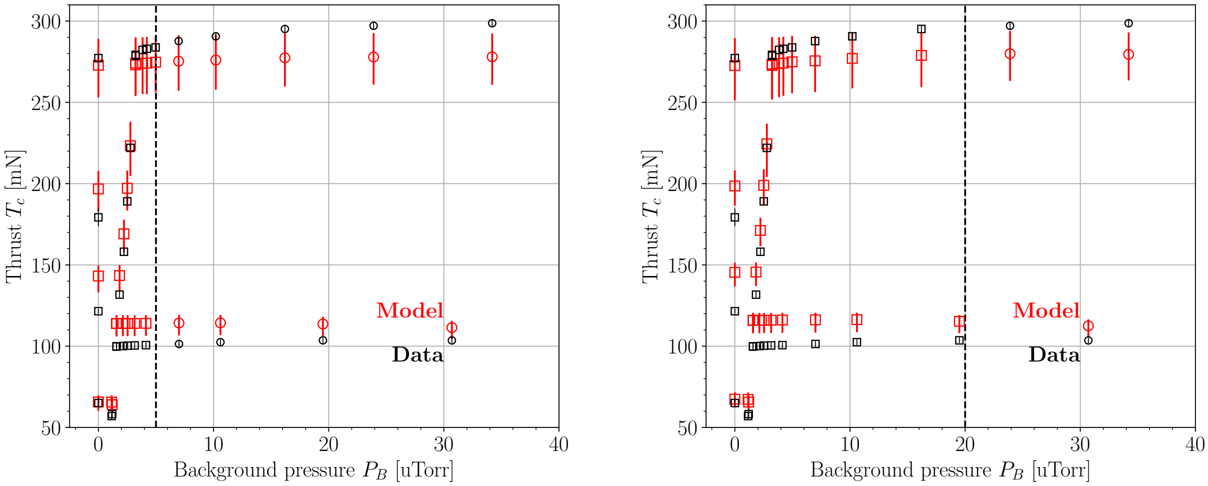

(a)

(b)

Fig. 11 Model predictions (red) and data from [22] and [11] (black) for thrust made using restricted training datasets. Circles are the operating conditions the model was trained on, and squares those it was not. (a) Model calibrated only on data above 5 µTorr background pressure, (b) model calibrated only on data above 20 µTorr background pressure.

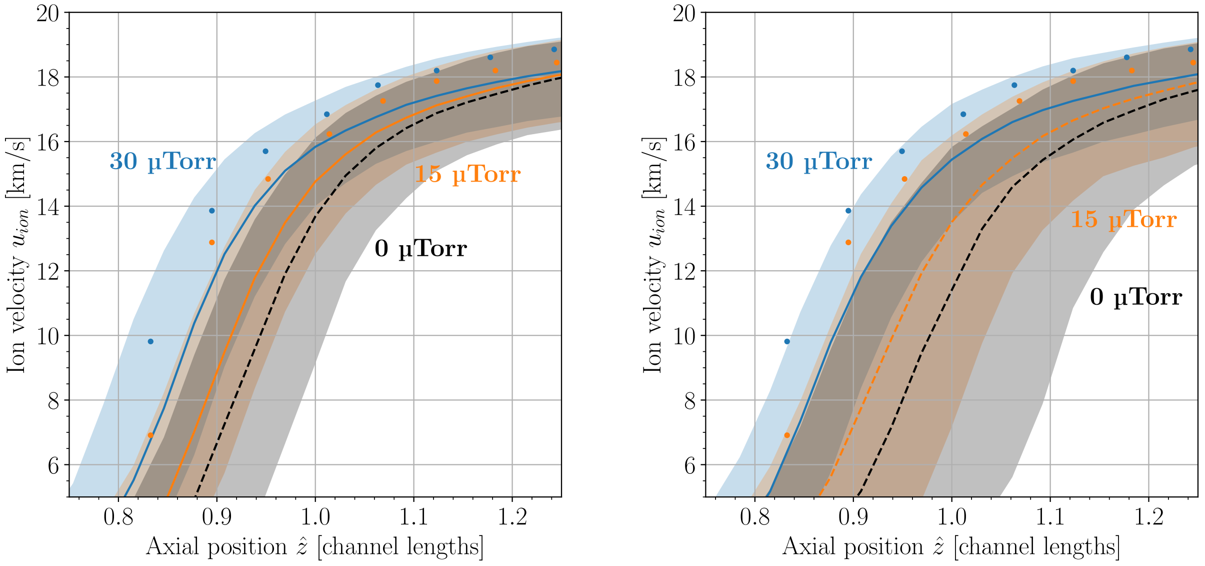

(a)

(b)

Fig. 12 Model predictions for (lines and shaded regions) and data (points, [2]) for ion velocity profiles at 4.5 kW, 6 A magnet current operation. Dashed lines denote operating conditions the model was not trained on. (a) Model calibrated only on data above 5 µTorr background pressure, (b) model calibrated only on data above 20 µTorr.

to space (0 µTorr). The dashed lines here denote operating conditions not in the training dataset.

Contrasting the recalibrated thrust with Fig. 3, we observe that the predictions have not changed much, despite having trained on substantially fewer data. The primary difference is that the median thrust prediction is modestly greater (less than 3%) for every operating condition. Importantly, the uncertainty in the predictions is approximately the same magnitude as for the full training dataset, even when extrapolating down to zero pressure. That the restricted sets of data are still informative enough to constrain the model and retrieve comparable thrust predictions suggests sufficient physics (even empirical physics) is incorporated within the PEM to extrapolate key trends in thrust. That is, the model "knows" that there should be a change in thrust consistent with the background pressure scalings of Eqs. (1), (6), and (8), so it is only necessary to provide a few data to constrain the scale of the system for these scaling laws. In the context of testing high power systems, this is a positive result: it suggests that even given only high pressure data—which may be the only attainable data—a useful prediction is possible. In particular, even the comparatively insensitive change in thrust with pressure at high pressures is sufficient.

For the ion velocity predictions, we find that both the 5 µTorr and 20 µTorr cutoffs are able to predict the 30 µTorr ion velocity profile comparably to the full dataset, with similar median predictions and an uncertainty in the position of the acceleration region of approximately 0.1 channel lengths. However, we see a qualitative difference between the two different restricted calibrations when extrapolating to lower pressures. For the 5 µTorr cutoff, the extrapolation to space has comparable uncertainty in the position and shape of the ion velocity profile to the predictions at higher pressures, consistent with the trends of Fig. 5b. However, the extrapolation to space for the 20 µTorr cutoff has substantially greater uncertainty in the position of the acceleration region, about 0.25 channel lengths. Greater uncertainty when extrapolating to lower pressures suggests that resolving the changes in the internal dynamics of the thruster still requires a comparatively rich data set of several different background pressures. Consequently, predicting the internal discharge dynamics—which are key to estimating thruster lifetime due to erosion of the channel [2]—likely remains challenging for high power systems if only high-pressure data are available.

C. Inverting for Mass Flow Rate

At a few points in the text we noted that while experimentally the discharge current is treated as an independent variable and mass flow rate as a dependent variable, in the model it is the opposite. Truly in experiments mass flow rate

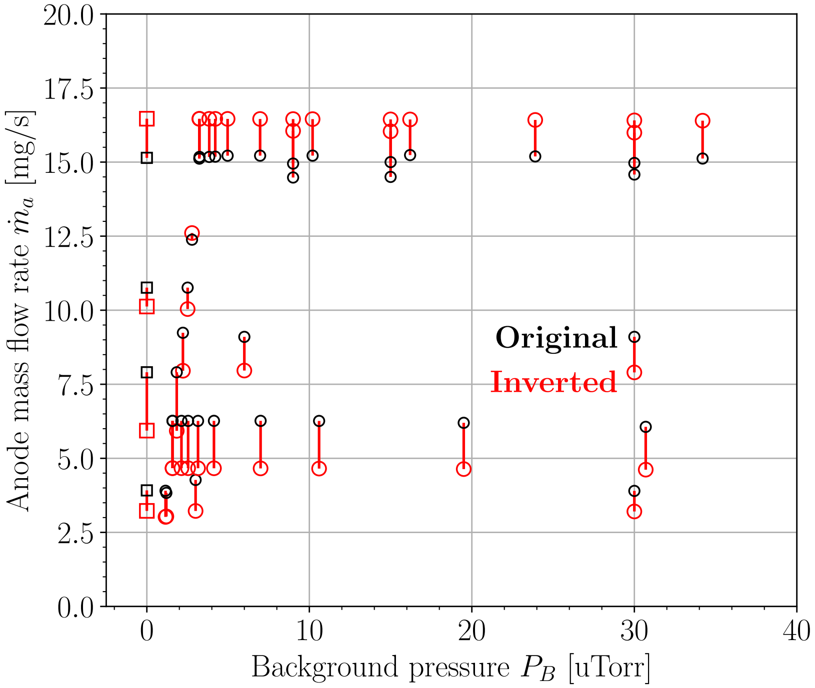

Fig. 13 Mass flow rate predictions of the inverted model (red) against the original data (black) [2, 11, 22]; circles and squares differentiate training and validation data sets, and the errorbars show the change in mass flow rate.

is indeed the variable controlled directly. Power supplies are typically operated in a voltage-controlled mode and a fixed mass flow fed into the system via a mass flow controller—the discharge current is a natural consequence determined by the plasma that forms. However, practically operators iteratively adjust the input mass flow rate until they achieve the discharge current they desire, inverting this dependency.

Inverting the model to reflect this actuality would tend to change our predictions. In particular, the model would be constrained to match the discharge current, with the corresponding shifts in mass flow rate propagating into other variables (e.g., thrust). As an illustration of how our predictions would differ with an inverted model, we wrapped the model in a nonlinear solver, such that mass flow rate was adjusted iteratively until the discharge current matched the experimental measurements. This is an expensive operation—it requires about ten regular model evaluations to converge, an order of magnitude increase in cost. Consequently, rather than recalibrating the inverted model, we simple use the MAP parameters of Tab. 4 to predict at each operating condition in both the training [2, 22] and validation [11] data sets. These parameters are not exactly what an inverted and fully recalibrated model would yield, but are a decent first estimate.

In Fig. 13, we render the re-tuned mass flow rates. Consistent with previous plots, the black points reflect data—the mass flow rates measured or assumed when originally calibrating the model—and the red reflect model predictions—the predicted mass flow rate from the inverted model. The vertical bars indicate the change in mass flow rate of each datum, and as before the circles denote the training operating conditions while the squares represent the validation operating conditions. We see that the changes in mass flow rate closely track the difference between the measured and modeled discharge current in Fig. 4. Essentially, the discrepancies in discharge current have been reconciled entirely by a change in ion current; less mass flowed in produces fewer ions at an approximately fixed ratio—the mass utilization efficiency. We find that changes in the mass utilization itself and in the electron current driven by anomalous mobility are more second-order effects.

The changes in ion current from this inversion have a corresponding effect on thrust, which we show in Fig. 14. This plot follows the symbology of Fig. 13, where the error bars indicate the change in thrust from the predictions of Fig. 3 resulting from inverting the model to match discharge current. Much as a changes in mass flow rate have produced proportional changes in ion current, changes in ion current have produced proportional changes in thrust. On the whole,

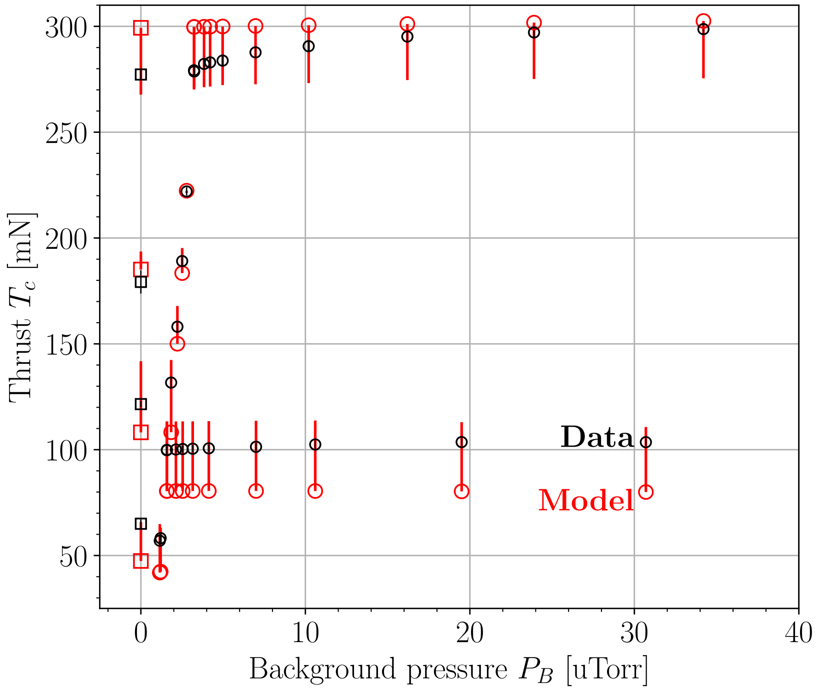

Fig. 14 Thrust predictions of the inverted model (red) against the original data (black) [11, 22]; circles and squares differentiate training and validation data sets, and the error bars show the change in thrust.

this has neither bettered nor worsened the predictions for thrust—some predictions more closely correspond to data, some are farther off. Essentially, in inverting the model we have traded matching mass flow rate exactly and discharge current only approximately to matching discharge current exactly and mass flow rate only approximately. That both parameters are not fully reconciled is a continuing consequence of the anomalous mobility physics we discussed in Sec. IV.A.

Nevertheless, inverting the model remains a useful exercise, because it permits us the flexibility to treat data more closely to how they are generated experimentally. Iteratively resolving the model from the beginning within a nonlinear solver to enforce the constraint on discharge current, however, has high computational expense—the 1-D thruster model must be initialized and allowed to converge fully each iteration. A more economical solution is to enforce the constraint on discharge current within the 1-D solver itself (i.e., by allowing the mass flow rate boundary condition to be adjusted as the solver iterates in order to match the specified discharge current). This approach has been pursued before in other codes and is an active area of development for the PEM.

V. Conclusion

A fundamental limitation in electric propulsion is that systems must be qualified in ground based vacuum test facilities that incompletely simulate the space environment, in particular effects caused by the finite partial pressure of residual gas. It is necessary to extrapolate these results to on-orbit conditions, and doing so is the central goal of NASA's JANUS, which has been working to develop a predictive engineering model for Hall thrusters. In this work, we successfully deployed the PEM on the SPT-140 system used as part of NASA's Psyche mission, generating a posterior distribution over the model parameters, comparing the model predictions to the training dataset, and extrapolating the model to the space environment.

We implemented a key modification to the PEM compared to previous work by exposing a linear scale factor to the radial magnetic field profile of the thruster, allowing the model to capture adjustments to the magnetic field that are made to tailor performance at different operating conditions. The calibrated Hall thruster component model parameters we showed to be largely the same as previous work (overlapping median parameter values for our prior calibrations on the SPT-100 and H9 systems), indicating some of the high-level physics are fundamentally transportable between system. Comparing the model predictions to the training dataset, we find that error in the thrust and ion velocity predictions—7% and 13% respectively—are comparable to our prior studies.

We saw a substantial increase in discharge current error compared to prior work, however (15% vs. 4%). The cause of this discrepancy was that the PEM has an anomalous electron collision frequency model which is not adjusted at different power levels. We contrasted these results with predictions of NASA JPL's 2-D Hall2De code, for which the anomalous collision frequency profile was tuned to individual operating conditions. These hand-tuned profiles suggested that a comparatively simple addition to the PEM's anomalous collision frequency model—a constant offset added to its standard cyclotron scaling—could allow it to represent a broad range of power conditions without needing to calibrate to individual operating conditions.

Finally, we examined recalibrating the model only on data above 20 µTorr (removing the region most sensitive to changes in background pressure). We found that in extrapolating to space, the thrust predictions differed from the fuller training dataset within 3% and had comparable uncertainty. This result suggests that there is sufficient prior intuition and physics built into the PEM that even comparatively few data collected at high power could be used to make predictions to the space environment. However, ion velocity profiles were less reliable. The position of the acceleration region extrapolated to space was uncertain to 0.25 channel lengths, compared to the uncertainty of 0.1 predicted with the fuller training dataset. The ion velocity profiles are representative of the states internal to the Hall thruster discharge and are crucial to assessing erosion by the plasma and thus device lifetime. These data may still require more detailed characterization at lower pressures to be identified.

Acknowledgments

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

References

[1] Ortega, A. L., Jorns, B., Mikellides, I. G., and Hofer, R. R., "Numerical simulations of the XR-5 Hall thruster for life assessment at different operating conditions," *51st AIAA/SAE/ASEE Joint Propulsion Conference*, 2018. https://doi.org/10.2514/6.2015-4008.

[2] Lopez Ortega, A., Mikellides, I. G., Chaplin, V. H., Snyder, J. S., and Lenguito, G., "Facility pressure effects on a Hall thruster with an external cathode: I. Numerical simulations," *Plasma Sources Sci. Technol.*, Vol. 29, No. 3, 2020, p. 035011. https://doi.org/10.1088/1361-6595/ab6c7e.

[3] Marks, T. A., and Jorns, B. A., "Challenges with the Self-Consistent Implementation of Closure Models for Anomalous Electron Transport in Fluid Simulations of Hall Thrusters," *Plasma Sources Science and Technology*, Vol. 32, No. 4, 2023, p. 045016. https://doi.org/10.1088/1361-6595/accd18.

[4] Eckels, J. D., Marks, T. A., Allen, M. G., Jorns, B. A., and Gorodetsky, A. A., "Hall thruster model improvement by multidisciplinary uncertainty quantification," *Journal of Electric Propulsion*, Vol. 3, No. 19, 2024. https://doi.org/10.1007/s44205-024-00079-w.

[5] Walker, M. L. R., Lev, D., Saeedifard, M., Jorns, B. A., Foster, J., Gallimore, A. D., Gorodetsky, A. A., Rovey, J. L., Chew, H. B., Levin, D., Williams, J. D., Yalin, A., Wirz, R. E., Marian, J., Boyd, I., Hara, K., and Lemmer, K., "Overview of the Joint AdvaNced PropUlsion InStitute (JANUS)," *Proc. of the 37th International Electric Propulsion Conference, Cambridge, MA, USA*, 2022. IEPC-2022-156.

[6] Marks, T. A., Eckels, J. D., Mora, G. E., and Gorodetsky, A. A., "Uncertainty quantification of a multi-component Hall thruster model at varying facility pressures," *J. Appl. Phys.*, Vol. 138, No. 15, 2025. https://doi.org/10.1063/5.0283796.

[7] Snyder, J. S., Sereno, V., Kerl, T., and Li, J., "Electric Propulsion for the Psyche Mission: One Year to Launch," *AIAA Propulsion and Energy 2021 Forum*, 2021. https://doi.org/10.2514/6.2021-3426, URL https://arc.aiaa.org/doi/abs/10.2514/6.2021-3426.

[8] Jorns, B. A., and Byrne, M. P., "Model for the Dependence of Cathode Voltage in a Hall Thruster on Facility Pressure," *Plasma Sources Science and Technology*, Vol. 30, No. 1, 2021, p. 015012. https://doi.org/10.1088/1361-6595/abd3b6.

[9] Marks, T., Schedler, P., and Jorns, B., "HallThruster.Jl: A Julia Package for 1D Hall Thruster Discharge Simulation," *Journal of Open Source Software*, Vol. 8, No. 86, 2023, p. 4672. https://doi.org/10.21105/joss.04672.

[10] Allen, M. G., Eckels, J. D., Byrne, M. P., Gorodetsky, A. A., and Jorns, B. A., "Application of Optimal Experimental Design to Characterize Pressure Related Facility Effects in a Hall Thruster," *Proc. of the 37th International Electric Propulsion Conference*, 2022.

[11] Snyder, J. S., Kelly, C. L., Garner, C., Bradley, N., Johnson, I., Corey, R., Ream, J. B., and Weiss, B. P., "Initial checkout of the Psyche electric propulsion system after launch," *38th International Electric Propulsion Conference*, ???? IEPC-2024-161.

[12] Mikellides, I. G., and Katz, I., "Numerical Simulations of Hall-effect Plasma Accelerators on a Magnetic-Field-Aligned Mesh," *Physical Review E*, Vol. 86, No. 4, 2012, p. 046703. https://doi.org/10.1103/PhysRevE.86.046703.

[13] Boeuf, J.-P., "Tutorial: Physics and Modeling of Hall Thrusters," *Journal of Applied Physics*, Vol. 121, No. 1, 2017, p. 011101. https://doi.org/10.1063/1.4972269.

[14] Hara, K., and Mikellides, I. G., "Characterization of low frequency ionization oscillations in Hall thrusters using a one-dimensional fluid model," *2018 Joint Propulsion Conference, Cincinatti, Ohio, USA*, 2018. https://doi.org/10.2514/6.2018-4904, AIAA 2018-4904.

[15] Sahu, R., Mansour, A. R., and Hara, K., "Full Fluid Moment Model for Low Temperature Magnetized Plasmas," *Physics of Plasmas*, Vol. 27, No. 11, 2020, p. 113505. https://doi.org/10.1063/5.0021474.

[16] Marks, T. A., and Jorns, B. A., "Evaluation of algebraic models of anomalous transport in a multi-fluid Hall thruster code," *Journal of Applied Physics*, Vol. 134, No. 15, 2023, p. 153301. https://doi.org/10.1063/5.0171824.

[17] Mikellides, I. G., Lopez Ortega, A., and Chaplin, V. H., "Theory of the anomalous momentum exchange from wave-particle interactions in Hall-effect ion accelerators and comparisons with measurements," *Phys. Fluids*, Vol. 36, No. 7, 2024, p. 074121. https://doi.org/10.1063/5.0213605.

[18] Brick, D. G., Roberts, P. J., and Jorns, B., "Numerical Investigation of Electron Energy Transport in Hall Thrusters," *AIAA SCITECH 2025 Forum*, 2025. https://doi.org/10.2514/6.2025-0298, URL https://arc.aiaa.org/doi/abs/10.2514/6.2025-0298.

[19] Hofer, R., Katz, I., Goebel, D., Jameson, K., Sullivan, R., Johnson, L., and Mikellides, I., "Efficacy of Electron Mobility Models in Hybrid-PIC Hall Thruster Simulations," *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, American Institute of Aeronautics and Astronautics, 2008. https://doi.org/10.2514/6.2008-4924.

[20] Macdonald-Tenenbaum, N., Pratt, Q., Nakles, M., Pilgram, N., Holmes, M., and Hargus, W., "Background Pressure Effects on Ion Velocity Distributions in an SPT-100 Hall Thruster," *Journal of Propulsion and Power*, Vol. 35, 2019, pp. 1–10. https://doi.org/10.2514/1.B37133.

[21] Brown, D. L., Walker, M. L. R., Szabo, J., Huang, W., and Foster, J. E., "Recommended Practice for Use of Faraday Probes in Electric Propulsion Testing," *AIAA Journal of Propulsion and Power*, Vol. 33, No. 3, 2017. https://doi.org/10.2514/1.B35696.

[22] Snyder, J. S., Lenguito, G., Frieman, J. D., Haag, T. W., and Mackey, J. A., "Effects of Background Pressure on SPT-140 Hall Thruster Performance," *Journal of Propulsion and Power*, Vol. 36, No. 5, 2020, pp. 668–676. https://doi.org/10.2514/1.B37702, URL https://doi.org/10.2514/1.B37702.

[23] Arkhipov, B., Bober, A., Day, M., Gnizdor, R., Kozubsky, K., and Maslennikov, N., "Extending the range of SPT operation - Development status of 300 and 4500 W thruster," *32nd Joint Propulsion Conference and Exhibit*, 1996. https://doi.org/10.2514/6.1996-2708, URL https://arc.aiaa.org/doi/abs/10.2514/6.1996-2708.

[24] Manzella, D., Manzella, D., Hamley, J., Miller, J., Clauss, C., Kozubsky, K., Gnizdor, R., Hamley, J., Miller, J., Clauss, C., Kozubsky, K., and Gnizdor, R., "Operational characteristics of the SPT-140 Hall thruster," *33rd Joint Propulsion Conference and Exhibit*, 1997. https://doi.org/10.2514/6.1997-2919, URL https://arc.aiaa.org/doi/abs/10.2514/6.1997-2919.

[25] Manzella, D., Sarmiento, C., Sankovic, J., and Haag, T., "Performance evaluation of the SPT-140," *27th International Electric Propulsion Conference, Cleveland, Ohio*, 1997. IEPC-097-59.

[26] W. Hargus, J., Fife, J., McFall, K., Jankovsky, R., Mason, L., Snyder, J., Malone, S., Haas, J., Gallimore, A., and Bauer, N., "Status of U.S. testing of the High Performance Hall System SPT-140 Hall thruster," *38th Aerospace Sciences Meeting and Exhibit*, 2000. https://doi.org/10.2514/6.2000-1053, URL https://arc.aiaa.org/doi/abs/10.2514/6.2000-1053.

[27] Delgado, J. J., "Qualification of the SPT-140 for use on Western Spacecraft," *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, 2014. https://doi.org/10.2514/6.2014-3606, URL https://arc.aiaa.org/doi/abs/10.2514/6.2014-3606.

[28] Mikellides, I. G., Katz, I., Hofer, R. R., and Goebel, D. M., "Magnetic shielding of a laboratory Hall thruster. I. Theory and validation," *J. Appl. Phys.*, Vol. 115, No. 4, 2014, p. 043303. https://doi.org/10.1063/1.4862313.

[29] Hofer, R. R., Goebel, D. M., Mikellides, I. G., and Katz, I., “Magnetic shielding of a laboratory Hall thruster. II. Experiments,” *J. Appl. Phys.*, Vol. 115, No. 4, 2014, p. 043304. https://doi.org/10.1063/1.4862314.

[30] Hofer, R. R., Cusson, S. E., Lobbia, R. R., and Gallimore, A. D., “The H9 Magnetically Shielded Hall thruster,” *Proc. of the 35th International Electric Propulsion Conference, Atlanta, GA, USA*, 2017. IEPC-2017-232.

[31] Myers, R. M., and Manzella, D. H., “Stationary Plasma Thruster Plume Characteristics,” *Proc. of the 23rd International Electric Propulsion Conference, Seattle, Washington, USA*, 1993. IEPC-1993-096.

[32] Walker, M. L. R., Victor, A. L., Hofer, R. R., and Gallimore, A. D., “Effect of Backpressure on Ion Current Density Measurements in Hall Thruster Plumes,” *Journal of Propulsion and Power*, Vol. 21, 2012. https://doi.org/10.2514/1.7713.

[33] Su, L. L., Marks, T. A., and Jorns, B. A., “Trends in mass utilization of a magnetically shielded Hall thruster operating on xenon and krypton,” *Plasma Sources Science and Technology*, Vol. 33, 2024, p. 065008. https://doi.org/10.1088/1361-6595/ad52be.