Downloaded by University of Michigan on January 30, 2024 | http://arc.aiaa.org | DOI: 10.2514/6.2024-2164

Optimal Experimental Design for Inferring Anomalous Electron Transport in a Hall Thruster

Madison G. Allen^{*} and Thomas A. Marks[†] and Joshua Eckels[‡] and Alex Gorodetsky[§] and Benjamin A. Jorns[¶] University of Michigan, Ann Arbor, MI, 48109

Optimal experimental design (OED) is applied to reduce the number of laser-induced fluorescence measurements required to calibrate the parameters of an anomalous electron transport model in a multi-fluid, 1D Hall thruster model. The time-averaged velocity data from the channel centerline of the SPT-100 Hall effect thruster is used as the dataset for calibrating a multi-zone Bohm model for electron transport. This dataset is then sub-sampled with an OED algorithm employing a batch method based on the expected information gain metric. It is shown that a batch with only 25% of the data points can yield model calibrations that match the median calibration from a full dataset, though the variance in the batch result is higher. This finding is discussed in the context of designing experimental efforts where time and resources are limited.

I. Nomenclature

θ	=	Model parameter
d	=	Data
Р	=	Probability
Z	=	Experimental design (Axial location)
<i>u</i> _i	=	Ion Velocity
<i>n</i> and <i>m</i>	=	Number of samples
v_{AN}	=	Anomalous electron collision frequency
c_i	=	Anomalous Coefficients
ω_{ce}	=	Electron cyclotron frequency

II. Introduction

Hall effect thrusters are one of the leading forms of electric propulsion due in large part to their high specific impulse (> 1500 s) and relatively high thrust density.[1, 2] The widespread success of these devices for cislunar applications has led to their consideration for missions beyond their historical operating envelope with increased powers (> 13 kW) and lifetimes (> 5 kh). Modeling and simulation tools are becoming increasingly important for enabling this paradigm of next-generation Hall thruster technology. These tools can help guide design and augment testing and qualification efforts. With that said, despite the high fidelity of existing numerical tools for Hall thrusters, there remain major challenges related to their predictive capabilities.[3, 4]

The predominant obstacle with Hall thruster modeling to date stems from ambiguity about a key aspect of their operation: the non-classical diffusion of electrons.[5, 6] The physical processes driving this enhanced diffusion are poorly understood, and as a result, it has not been possible to self-consistently predict its effects in the types of multi-fluid based models that are commonly employed for engineering of Hall effect thrusters. Instead, it is a common practice to introduce free parameters in the model, typically represented as a spatially-varying transport coefficient, that then can be inferred by iteratively running the model and comparing its output to experimentally measurable quantities like the acceleration of ions along channel centerline and the thrust. This approach to modeling has been highly successful in

^{*}Ph.D. Candidate, Department of Aerospace Engineering, mgallen@umich.edu

[†]Postdoctoral Student, Department of Aerospace Engineering, marksta@umich.edu

[‡]Ph.D. Candidate, Department of Aerospace Engineering, eckelsjd@umich.edu

[§]Assistant Professor, Department of Aerospace Engineering, goroda@umich.edu

[¶]Associate Professor, Department of Aerospace Engineering, bjorns@umich.edu, AIAA Associate Fellow

guiding thruster design, extracting new insight into the physics of these devices, and informing near term projections for lifetime [7, 8].

One limitation for this current method of calibration is that the method for acquiring the experimental data can be expensive and time-consuming. This is a particular concern for generating datasets of the acceleration of ions with the method of laser-induced fluorescence (LIF). In some cases, collecting LIF data for a single thruster operating condition can require more than a day of experimental time.[9] This invites the question as to whether it is possible to reduce the number of LIF measurements that must be taken to inform model coefficients.

The purpose of this work is to explore this possibility through the lens of optimal experimental design. To this end, this paper is organized in the following way. In the first section, we overview a 1D multi-fluid model for Hall thrusters, how anomalous transport is parameterized in this model, and how experimental data informs the calibration. In the second section, we introduce the method of a batch-based optimal experimental design. In the third section, we present a proof of concept for this method based on a published dataset for the SPT-100 thruster. We conclude with a discussion of the limitations and implications of our results.

III. Hall thruster model and calibration

In this section, we will briefly overview the model we investigated for this effort, the parameterization of electron transport in this model, and the standard method employed for model calibration.

A. Hallthruster.jl

We first provide the following figure for the simulation region for a Hall thruster.



Fig. 1 A schematic of a Hall thruster

The model implemented for this work is an open source code developed by the University of Michigan, HallThruster.jl.[3] Here we account for a one-dimensional region to propagate across in the axial direction, perpendicular to the thruster face. The model assumes a multi-fluid system including neutrals, ions of multiple species, and electrons. For each fluid, we solve the equations of motion- continuity and momentum. For electrons, we solve a generalized Ohm's law. By including the assumption of a quasineutral plasma, we solve for the electric potential as well. [10] If we suppose there exists a closure model for the electron dynamics, then with these tools we simulate the behavior of a Hall thruster for a given geometry and condition, such as mass flow rate, discharge voltage, and magnetic field. We are then able to produce a time averaged solution and output specifically the mean ion velocity profile for each ion species along channel centerline, as well as other metrics such as discharge current and thrust.

The model input variables for this study are shown in Table 1. Due to the intricacy of HallThruster.jl, the computational run times can be about 1 to 2 minutes for 2 ms of a Hall thruster simulation. This can be prohibitively long for the method of OED outlined in the following sections. To reduce computational time, we developed a surrogate model for the HallThruster.jl outputs.[11] Similar to fitting a model to data, the surrogate is fit by adjusting hyperparameters until the original high fidelity model, HallThruster.jl, matches. This approach, utilizing the work by Jakeman et al., adaptively generates a data-driven map between model parameters and the ion velocity along the centerline, $M_{sug.}(\theta) \rightarrow u_i^M(z_j; \vec{\theta})$. To determine the accuracy of the surrogate, we convergence of the model must be met, specifically convergence to the high fidelity output.

B. Parameterization of electron transport

As with all multi-fluid Hall thruster models, the cross-field electron transport is modeled by introducing a parameterized model for the transport coefficient. In this case, this is represented by prescribing an "anomalous collision

frequency," v_{AN} , as a function of axial position in the thruster. Variations to this effective collision frequency, which greatly exceeds classically-induced collisional transport, can promote electron currents commensurate with experimental measurement. In our case, we adopt a multi-zone Bohm model:

$$v_{AN} = \begin{cases} c_1 \omega_{ce} & \text{if } z < z_{start} \\ c_2 \omega_{ce} & \text{if } z > z_{end} \end{cases}$$
(1)

where $c_1, c_2, z_{start}, z_{end}$ are free parameters, ω_{ce} is the electron cyclotron frequency, and z is the axial distance from the anode. Physically, this represents Bohm diffusion of the plasma, or that the anomalous collision frequency follows Bohm scaling laws such that the anomalous diffusion is related to the varying strength of the magnetic field lines. [7] In this work, we apply the two zones seen in Eq. 1. The region between the two nodes, z_{start} and z_{end} , is a logarithmic linear interpolation from c_1 to c_2 for a continuous transition. By including the two spatial free parameters, we are able to shift where the ions experience a large acceleration out of the channel. To calibrate the model, we must tune the free parameters until the model aligns with the experimental dataset. There are many methods of calibration, but here we focus on parameter estimation using Bayes' Rule.

C. Model calibration

As outlined in the preceding, the parameterization of the anomalous collision frequency introduces four unknown coefficients, $c_1, c_2, z_{start}, z_{end}$, and these must be inferred via experimental data. Directly measuring the effective transport coefficient in these discharges has historically proven to be nearly intractable. Instead, it is common practice to use another, more experimentally accessible, prediction from the model, e.g. the ion velocity at multiple spatial locations, for comparison. We represent the model prediction for ion velocity as $u_i^M(z_j; \vec{\theta})$ where z_j denotes the spatial location of the prediction and $\vec{\theta} = (c_1, c_2, z_{start}, z_{end})$ are the values of model parameters selected to run the simulation.

We adopt a Bayesian framework to infer the values of the model parameters from a set of experimental measurements, $\vec{d} = (z_1, u_{i(1)}), (z_2, u_{i(2)}), ...(z_m, u_{i(m)})$. Specifically, the Bayesian approach seeks a posterior probability distribution $P(\vec{\theta} \mid \vec{d})$ over the uncertain parameters $\vec{\theta}$ through the application of Bayes' rule

$$P(\vec{\theta}|\vec{d}) \propto \mathcal{L}(\vec{d}|\vec{\theta}) P(\vec{\theta}), \tag{2}$$

where $P(\vec{\theta})$ denotes the prior probability distribution function for the model parameters and $\mathcal{L}(\vec{d}|\vec{\theta})$ denotes the probability density of the experimental dataset under the given model and parameters. Specifically, the prior distribution encapsulates information about the parameters prior to obtaining data, e.g., knowledge about the range or expected statistics. The likelihood is chosen to arise as a normally-distributed function due to a Gaussian assumption on the measurement noise.

$$\mathcal{L}(\vec{d}|\vec{\theta}) = \exp\left[-\sum_{j=1}^{m} \frac{1}{2\sigma_j^2} \left(u_{i(j)} - u_i^M(z_j;\vec{\theta})\right)^2\right],\tag{3}$$

where σ_j is the experimental uncertainty in the measurement. For nonlinear models u_i^M , there is generally no analytical solution to Bayes' rule, and numerical approaches are required. A standard approach is to generate samples from the posterior via Markov Chain Monte Carlo. Using samples, we can generate statistical estimates of the median and credible intervals for the values of model parameters. A major challenge in this approach, however, is that running the model to completion can be computationally expensive. In our setting, HallThruster.jl is a relatively inexpensive

model with a 1-2 minute run time. Yet, this time can still be prohibitive for obtaining the number of samples (> 10,000) required to accurately estimate statistics of interest. This motivates the use of the surrogate described in the preceding section.

IV. Optimal Experimental Design

The uncertainty represented by the posterior distribution of the parameters is directly a function of the choice of prior distribution and the quality of the data. In the experimental context, it is therefore critical to choose highly informative data to rapidly reduce the uncertainty in the unknown parameters. In practice, it is common to attempt to take a sufficient number of measurements such that the full acceleration of the ions can be captured in the thruster–typically with 1-5 mm resolution. For most systems this corresponds to over twenty data points, a dataset size that can be prohibitive or time consuming from an experimental standpoint.

In this section, we describe an optimal experimental design (OED) framework that can leverage a model to determine a smaller subset of experimental measurements that could in principle yield the same effective calibration for the model with a standard, larger dataset. The premise of model (or simulation) based OED is to leverage the model to search a given "design space"—in this case possible locations for performing an ion velocity experimental measurement—to determine which experiments would most likely reduce uncertainty in the parameters. This reduction in uncertainty would then signify swift convergence of the model parameters.

The method we apply in this work is based on the Expected Information Gain, which we apply to design a batch of experiments to conduct.[12, 13] To motivate this approach, we first introduce the concept of information gain. We denote each batch as k experimental conditions, in our case k locations where ion velocity data could be obtained. For example, for the j^{th} batch we have $\vec{b}^{(j)} = (z_1^{(j)}, z_2^{(j)}, ..., z_k^{(j)})$. We denote the resulting velocity measured from this batch as $\vec{d}^{(j)} = (d_1^{(j)}, d_2^{(j)}, ..., d_k^{(j)})$. The information gain in our probabilistic representation of the model parameters yielded by measuring at these locations is quantified by the difference between prior and posterior parameter distributions. This is formulated in terms of the KL divergence:

$$D_{KL}(P(\vec{\theta}|\vec{d}^{(j)}), P(\theta)) = \int P(\vec{\theta}|\vec{d}^{(j)}) \ln \frac{P(\vec{\theta}|\vec{d}^{(j)})}{P(\theta)} d\theta.$$
(4)

This prescription is unambiguous for comparing a prior and posterior provided we have done the measurement. The challenge in OED, however, is that we have not yet performed the experiment. We instead need to make an educated guess about which measurements have the highest likelihood, given our current understanding of the model, to improve our knowledge of the model parameters. To represent this rigorously, we introduce the expected information gain as the average of the information gain over all possible realizations of what the measurements at locations $\vec{b}^{(j)}$ could be

$$EIG(\vec{b}^{(j)}) = \int P(\vec{d}^{(j)})(D_{KL}(P(\vec{\theta}|\vec{d}^{(j)}), P(\theta))d_1^{(j)}d_2^{(j)}...d_k^{(j)}$$
(5)

The integral in this result is not tractable. As an alternative, it is a common practice to approximate it with a nested sampling scheme [12]:

$$EIG(\vec{b}^{(j)}) = \frac{1}{n} \sum_{i=1}^{n} \left(\ln P(\vec{d}^{(i)} \mid \theta^{(i)}) - \ln \frac{1}{m} \sum_{j=1}^{m} \left(P(\vec{d}^{(i)} \mid \theta^{(i,j)}) \right) \right), \tag{6}$$

where $P(\vec{d} \mid \theta)$ is the same form as the likelihood given in Eq. 3. [14, 15] The following algorithm describes how we evaluated this expression: We note in this algorithm that the variance in sampling is assumed to be the typical experimental error. For our evaluations we used typical nesting sampled values of n = 2000 and m = 3000.

The above approach provides a framework for determining a priori the best set of locations to measure, assuming the number of measurements is constrained to a set of n points. To evaluate this approach, we describe in the following program a test case based on an existing data set for a Hall effect thruster.

Algorithm 1: Expected Information Gain for one batch of k designs

for *i* in *n* do Sample $\theta^{(i)} \sim P(\theta)$; Evaluate $M(\theta^{(i)}, d_{0...k})$; Sample $d_{0...k}^{(i)} \sim \mathcal{N}(M(\theta^{(i)}, d_{0...k}), \sigma^2)$; Evaluate likelihood $P(d^{(i)} | \theta^{(i)})$; for *j* in *m* do Sample $\theta^{(i,j)} \sim P(\theta)$; Evaluate the likelihood $P(d_{0...k}^{(i)} | \theta^{(i,j)})$; end end



Fig. 2 Axial ion velocity profile shown with errorbars and the normalized magnetic field along channel centerline for an SPT-100 with the peak at the exit plane

V. Test case

We considered in this work open source data from the SPT-100 Hall effect thruster. This device is $259 \times 164 \times 120$ mm Hall effect thruster with a channel length of 25 mm that operates on xenon at nominally 300V and 4.5A. The dataset used for the test case was published by MacDonald-Tenenbaum et al as a part of a background pressure study on ion velocity distributions.[16] Here we consider the dataset at the lowest background pressure recorded, 1.7E - 05 Torr, with a discharge current of 4.25A. The axial, time-averaged LIF measurements were taken at 22 locations along channel centerline from 9 mm to 75 mm. The most probable velocity was calculated for each distribution to arrive at the ion velocity profile shown in Fig. 2. We also show in Fig. 2 the normalized magnetic field along the channel that we used to determine the cyclotron frequency. We formulated a model for the SPT-100 in HallThruster.jl using the reported flow rates, magnetic field distribution, and geometry. Our simulation consisted of a 200 point discretization.

We used the existing 22-point dataset as a proof of concept for batching to inform model calibration in the following way. We begin with no experimental measurements and assume that only the axial position of these 22 locations as possible regions where we can perform measurements. We then restrict our batches dataset to only five of these locations. We evaluate the EIG for all possible five point configurations and determine the combination that yields the highest value. We then use the actual measurements from the dataset to perform Bayesian inference to determine the model parameters and the subsequent comparison of model velocity to the measured velocity. We contrast this result with the model calibration yielded from the full 22 point dataset.

VI. Results

In this section, we first present the EIG for different five-point batch combinations. We then show the results of the model calibration comparing the results from the best batch and from using the full dataset for calibration.

A. Expected information gain from batching

We determined the optimal set of 5 designs by comparing all combinations of 5 positions across the 22 available positions provided by the SPT-100 dataset. This resulted in over 26,000 evaluations for EIG to compare and find a maximum as the optimal set. We show in Fig. 3 an illustrative example of how the EIG varied for four different batch configurations. These include sequential groups of 5 at the beginning of the dataset, the middle, and the end. The fourth dataset is the argument for the maximum EIG that we identified from all roughly 26,000 possible combinations, i.e. the optimal design.



Fig. 3 Hall thruster diagram with example batches of axial locations for laser induced fluorescence measurements: a.) first 5 locations, b.) middle 5 locations, c.) end 5 locations, d.) optimal 5 locations

Figure 4 shows the expected information gain for each set shown Fig. 3. It is evident from these results that datasets that are grouped spatially are less informative for updating the model. This is intuitive as these measurements do not provide a diversity of data—the expected change in ion velocity in this region is moderate. In contrast, the optimal dataset has measurement locations distributed across the possible locations. This is a qualitative indication of the effectiveness of the OED in guiding the batch decision. It is also physically intuitive as an experimentalist would likely opt to take measurements spaced across the domain of interest.



Fig. 4 Expected information gain for the example comparison of sets of axial locations shown in Fig. 3 in order a.) to d.) where d.) is optimal

B. Results from model calibration

Armed with the optimal design identified from the previous section, we show in Fig. 5 the resulting posterior distributions for the two anomalous coefficients c_1 and c_2 as well as the two spatial parameters z_1 and z_2 from calibrating with the optimal batch identified from the previous section. For comparison, we also show the marginalized posterior distribution from the calibration performed with the full dataset.

There are two salient features from these results. First, we note that the median values for both distributions are approximately the same, though the batch result is slightly lower for each coefficient. This qualitatively suggests that the batch largely is able to capture the same median value inferred from a full dataset. The second feature is that the variance of the distribution with the batch is twice that of the full dataset. This latter result is expected given that certainty improves with more model data. This result speaks to a larger and more general trend, which is that batching likely can yield the same "best fit" but the relative confidence in this fit will suffer.



Fig. 5 Posterior distributions for the transport coefficients using a.) the optimal dataset against b.) the full dataset

We illustrate the consequences of this uncertainty graphically by showing in Fig. 6 the median ion velocity prediction compared to data generated from the model parameters inferred with the full dataset and with the ideal batch. We also show the credible intervals for both fits. As can be seen, the median values for the models align over most of the experimental domain with a slight deviation in the downstream region. However, the credible intervals for the batch result are wider near the acceleration region and further downstream, a direct result of the inherently larger uncertainty in this regression. More generally, we also remark that both the full data set and batch data calibrations exhibit major departures from the experimental data in the upstream region, overpredicting the ion velocity. This result suggests a deficiency in the model itself rather than the calibration method. We also note the broad distribution at the first location for the full data set fit. This is possibly due to the fitting mechanisms using Markov Chain Monte Carlo with the large uncertainty in the data at that region.

With that said, in the context of this specific model, the results for the median comparison of the batch technique to the full dataset are compelling, suggesting that this is potentially a viable technique for designing more sparse experimental campaigns. We expand on this finding in the following section.

VII. Discussion

We have demonstrated using an optimal experimental design framework that we can produce a model calibration with only about 25% of the dataset. The uncertainty in the coefficient estimations shown in the width of the parameter posterior distributions is larger for the reduced, optimal calibration than the full calibration. This can be expected since incorporating data typically tends to increase the confidence in a model fit. To decrease the error in the coefficients, more data may be necessary.

To this end, larger batches of possibly 10 optimal designs, for example, could improve this error. Alternatively, a higher fidelity parameterized model may need fewer data points to reach the same level of model confidence. We note here as well that our approach in this work was based on expected information gain, a metric to try to explore parameter space for calibrating the model coefficients. There may be other metrics specifically targeted toward reducing model uncertainty that could be more parsimonious with point selection



Fig. 6 Ion velocity profile from the optimal set calibration in red compared to the full set calibration in black with 5% and 95% credible intervals

In terms of the practical implications for this result, we recall that a motivating factor in this study was to determine a priori if there strategies for reducing the number of necessary experimental points. Our results suggest that the best strategy is to select a diversity of datasets over the experimental domain–a result consistent with experimental intuition. More broadly, by applying a similar approach on different test articles, we may be able to arrive at some consensus about which diverse datapoints for a typical thruster are the most value for informing LIF data, which in turn may be a precursor to best practices. This is a particularly crucial consideration given the expense and challenge of testing higher power thrusters.

With that said, we do note that the computational run times for this process can be extensive depending on the size of the batch of designs requested. If OED is executed in advance of the test, its application may be worth the savings in expense for setup and implementation of test. However, the computational test likely precludes the use of this method as an attractive alternative for information real-time selection of experimental data points.

VIII. Conclusion

The development of fully predictive engineering models for Hall effect thrusters is hindered by the phenomena of anomalous electron transport. The current method to improve Hall thruster model fidelity is to calibrate the anomalous collision frequency to data collected by laser-induced fluorescence. This diagnostic is expensive and time-intensive to implement.

Motivated by this challenge, we investigated in this work the possibility of using optimal experimental design to determine if it is possible to calibrate a model with fewer ion velocity measurements. By utilizing the function Expected Information Gain in a batch method, we were able to show that we could achieve the same median predictions for a calibrated 1D Hall thruster model with only 25% of the nominal data set. This serves as positive proof of concept for guiding future experimental efforts. Indeed, this work could be leveraged as guide for designing best practices for designing new ion velocimetry experiments in Hall thrusters. This is a particularly critical consideration for testing next-generation systems, where the expense and challenge of implementing diagnostics increases.

IX. Acknowledgments

This work is supported by the NASA Space Technology Graduate Research Opportunity 80NSSC22K1172 and the Joint Advanced Propulsion Institute (JANUS), a NASA Research Institute.

References

 Goebel, D. M., Jameson, K. K., Katz, I., and Mikellides, I. G., "Potential fluctuations and energetic ion production in hollow cathode discharges," *Physics of Plasmas*, Vol. 14, 2007, pp. 1–15. https://doi.org/10.1063/1.2784460.

- [2] Simmonds, J., Raitses, Y., and Smolyakov, A., "A theoretical thrust density limit for Hall thrusters," *Journal of Electric Propulsion*, Vol. 2, 2023, p. 12. https://doi.org/10.1007/s44205-023-00048-9.
- [3] 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, 2023, p. 4672. https://doi.org/10.21105/joss.04672.
- [4] Ortega, A. L., Katz, I., Mikellides, I. G., and Goebel, D. M., "Self-Consistent Model of a High-Power Hall Thruster Plume," *IEEE Transactions on Plasma Science*, Vol. 43, 2015, pp. 2875–2886. https://doi.org/10.1109/TPS.2015.2446411.
- [5] Meezan, N. B., Hargus, W. A., and Cappelli, M. A., "Anomalous electron mobility in a coaxial Hall discharge plasma," *Physical Review E*, Vol. 63, 2001, p. 026410. https://doi.org/10.1103/PhysRevE.63.026410.
- [6] Lafleur, T., Baalrud, S., and Chabert, P., "Theory for the anomalous electron transport in Hall effect thrusters. I. Insights from particle-in-cell simulations," *Physics of Plasmas*, Vol. 23, No. 5, 2016.
- [7] Mikellides, I. G., Katz, I., Hofer, R. R., and Goebel, D. M., "Hall-effect thruster simulations with 2-D electron transport and hydrodynamic ions," *Internationel Electric Propulsion Conference*, 2009.
- [8] 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.
- [9] Su, L. L., Marks, T. A., and Jorns, B. A., "Investigation into the Efficiency Gap between Krypton and Xenon Operation on a Magnetically Shielded Hall Thruster," 2022.
- [10] Hara, K., "Non-oscillatory quasineutral fluid model of cross-field discharge plasmas," *Physics of Plasmas*, Vol. 25, No. 12, 2018, p. 123508. https://doi.org/10.1063/1.5055750, URL https://doi.org/10.1063/1.5055750.
- [11] Jakeman, J. D., Friedman, S., Eldred, M. S., Tamellini, L., Gorodetsky, A. A., and Allaire, D., "Adaptive experimental design for multi-fidelity surrogate modeling of multi-disciplinary systems," *International Journal for Numerical Methods in Engineering*, Vol. 123, 2022, pp. 2760–2790. https://doi.org/10.1002/nme.6958.
- [12] Ryan, K. J., "Estimating Expected Information Gains for Experimental Designs With Application to the Random Fatigue-Limit Model," *Journal of Computational and Graphical Statistics*, Vol. 12, 2003, pp. 585–603. https://doi.org/10.1198/1061860032012.
- [13] Huan, X., and Marzouk, Y. M., "Simulation-based optimal Bayesian experimental design for nonlinear systems," *Journal of Computational Physics*, Vol. 232, 2013, pp. 288–317. https://doi.org/10.1016/j.jcp.2012.08.013.
- [14] Lindley, D. V., "On a measure of the information provided by an experiment," *The Annals of Mathematical Statistics*, Vol. 27, No. 4, 1956, pp. 986–1005.
- [15] Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N., "A review of modern computational algorithms for Bayesian optimal design," *International Statistical Review*, Vol. 84, No. 1, 2016, pp. 128–154.
- [16] 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. 403–412. https://doi.org/10.2514/1.B37133.