

Model Inference from Electrospray Thruster Array Tests

Collin B. Whittaker*, Alex A. Gorodetsky†, and Benjamin A. Jorns‡
University of Michigan, Ann Arbor, Michigan, 48109

A model for the current emitted by a porous conical-type electrospray array is formulated by considering the operation of individual emitters in parallel. The individual emitter behavior is determined from five interrelated models: one predicting the current sourced by an individual emission site based on the feeding conditions of the substrate and the applied electric field, another predicting the number and size of such sites for a given emitter, a model which calculates the electric field at the emitter from the applied voltage, and lastly models which predict the characteristic fluid flow impedance and the electric field as a function of device geometry. A Bayesian inference problem over the associated model parameters is constructed, incorporating an approximate pseudomarginalization over the device geometry and other nuisance parameters of the model to propagate their uncertainty. This inference is performed considering experimental data and thruster geometry for a 576 emitter array developed by the Air Force Research Laboratory. Samples from the resulting posterior distribution over parameters are drawn and subsequently used to make predictions for the device behavior, which is contrasted with that observed experimentally. It is determined that if the variability in emitter geometry is neglected, the model is unable to reproduce the experimental behavior; by accounting for this effect, however, the model matches experimental observations. This contrast in model agreement is discussed in the context of the role of manufacturing tolerance on emitter performance and in the design of similar systems in the literature. Limitations on the extensibility of the ionic emission model are also examined in light of the fact that there is a possible correlation between the model parameters and geometry.

I. Nomenclature

R_c	=	emitter tip radius of curvature
α	=	emitter cone half-angle
h	=	emitter height
d	=	tip-to-extractor distance
R_A	=	aperture radius
r_p	=	emitter pore radius
i	=	emitter index
I	=	current (various)
j	=	emission site index (per emitter)
N_s	=	number of emission sites
Q	=	volumetric flow rate (various)
ρ	=	propellant density
$\langle \frac{q}{m} \rangle$	=	bulk charge-to-mass-ratio
p_r	=	reservoir pressure
p_b	=	Taylor cone base pressure
R_H	=	hydraulic impedance (various)
γ	=	propellant surface tension
$r_{p,r}$	=	characteristic reservoir pore radius
ζ	=	ionic emission pressure dependence
E	=	applied electric field

*PhD Candidate, Department of Aerospace Engineering, cbwhitt@umich.edu, Student Member AIAA.

†Assistant Professor, Department of Aerospace Engineering, Member AIAA.

‡Assistant Professor, Department of Aerospace Engineering, Associate Fellow AIAA.

E_0	=	onset electric field
ε_0	=	vacuum permittivity
E_c	=	characteristic electric field
p_c	=	characteristic capillary pressure
r_0	=	characteristic emission site size
\hat{I}	=	dimensionless emission site current
K	=	propellant conductivity
\hat{E}	=	dimensionless applied field
\hat{E}_0	=	dimensionless onset field
\hat{p}_r	=	dimensionless reservoir pressure
C_R	=	dimensionless hydraulic impedance
ζ_1	=	ionic emission offset coefficient
ζ_2	=	ionic emission slope coefficient
μ	=	propellant viscosity
κ	=	emitter substrate permeability
ϕ	=	emitter substrate porosity
b_0	=	pooling radius
N_{max}	=	maximum possible number of emission sites
r_{min}	=	smallest active feature size
V	=	voltage (various)
$\frac{ E }{V}$	=	voltage-normalized electric field magnitude
N_e	=	number of emitters
N_d	=	number of data
X	=	collection of data
k	=	datum index
x	=	individual datum
θ	=	collection of model parameters
$P()$	=	probability
G	=	nuisance parameters
F	=	model prediction
σ	=	standard deviation, experimental noise
$\hat{P}()$	=	Monte-Carlo estimate to probability
N_r	=	number of realizations
l	=	Monte-Carlo realization index
X^*	=	collection of predicted data
\mathcal{N}	=	normal (Gaussian) distribution
\mathcal{U}	=	uniform distribution, utility function
$O()$	=	order (Big O notation)
N_f	=	number of samples
ζ_1^*	=	alternative ionic emission offset coefficient

II. Introduction

ELECTROSPRAY thrusters are a form of electric propulsion utilizing electrohydrodynamic processes to extract charged droplets or ions from an electrified liquid meniscus. These particles are then accelerated through an applied electric field to achieve high exit velocities and produce thrust. This principle of operation leads to several potential advantages compared to more conventional electric propulsion (EP) technologies like Hall effect thrusters (HETs) and gridded ion thrusters [1]. First, propulsive electrosprays operate typically on liquid metals or ionic liquids, room-temperature molten salts with negligible vapor pressures. This prevents the propellant from boiling off in the vacuum of space and allows it to be stored more compactly and at lower pressure than gaseous propellants. Second, since the propellant does not need to be ionized or confined, the electrospray thruster avoids inherent volume to area losses associated with forms of EP that rely on impact ionization to energize their propellant. Third, the thrust generation process for electrosprays

occurs at the micro- to nanoscale. Individual electrosprays thus have the highest thrust densities of any form of EP [1]. This property in principle could be leveraged to create ultra compact, high power systems.

The major limitation of electrospray thrusters stems from the fact that individual emitters generate low levels of thrust—typically less than $\sim 1 \mu\text{N}$. This thrust level, though useful for precise attitude control [2], is impractical for providing spacecraft primary propulsion. In order to realize the potential of this technology for higher power missions and larger spacecraft, it is necessary to create arrays with hundreds or thousands of emitters [3–5]. To this end, several architectures for scaling electrospray emitters to larger arrays have been proposed to date, including capillary emitters which wick propellant to the tip of a capillary needle [2] and externally-wetted emitters which transport propellant along a roughened exterior [6]. Of particular interest, however, are porous electrosprays, which operate by conducting the propellant through a porous substrate. In these systems, the entire body of the emitter serves as a propellant reservoir, the character of the emission can be altered by the choice of pore size and porosity, and emitters can be machined both by laser ablation and through conventional means [3, 7]. Additionally, these devices can be made to emit purely ionic species, achieving higher specific impulse and providing corresponding propellant savings. Together these properties have led to a growing interest in leveraging porous electrospray emitters for larger-scale arrays [1, 3, 7–9].

Creating large-scale (e.g. $> 10,000$ emitters) porous electrospray arrays has proved to be a major technical obstacle to date. Indeed, while it is possible to fabricate an individual emitter with high performance metrics, e.g. specific impulse $> 1000 \text{ s}$ and efficiency greater than 50%, uncertainty in manufacturing tolerance can lead to variability in emitter performance when this structure is duplicated thousands of times in an array. This will result in aggregate reductions to the array’s performance. As an extreme example, for an array intolerant to the failure of a single emitter (i.e. by arcing to the extractor), even one off-nominal emitter in an array of thousands can result in a precipitous drop in predicted device lifetime [10]. The challenge thus is ultimately to identify “robust” array designs capable of maintaining required performance metrics in spite of known manufacturing tolerances. With that said, the trade space of electrospray configurations is large. Possible design iterations include but are not limited to substrate material, porosity, propellant choice, emitter shape, emitter spacing, extractor distance, and operating voltage. Systematically investigating this wide parameter space for robust designs thus poses a major technical challenge.

In order to address this challenge, we have been pursuing an iterative framework based on reduced-fidelity modeling, rapid prototyping, and reinforced learning. As we recently presented [11], by employing computationally inexpensive models based on established fundamental scaling laws, we can numerically investigate the performance of millions of design combinations. We similarly are able to determine how robust the performance of the most promising designs are to manufacturing tolerance. The robust designs in turn inform the fabrication and testing of prototypes. The advantage of this approach is that it allows us—in principle—to converge more rapidly to new and optimal array configurations. With that said, while lower fidelity models are an enabling feature of this method, they are also its major limitation. Indeed, unlike full, direct numeric simulations, the scaling laws we employ rely on empirical model parameters that must be inferred from data. These parameters have their own inherent uncertainty (which impacts confident model predictions), and, since they are calibrated empirically, their use invites questions about the extensibility of the models beyond the datasets we used to calibrate them.

With these limitations in mind, a key element of our approach is to continue to update our reduced fidelity models as we design and build new prototypes. As we continue to generate more data from each new array, the fidelity of the model parameters will only improve—provided the underlying physics of the reduced models are valid. To this end, in our previous work [12], we demonstrated the ability to use data from an individual porous-based emitter with known geometry to perform inference for several scaling laws applicable to porous electrosprays. We in turn were able to characterize model-based uncertainty inherent to these scaling laws.

While the application of model inference to a single emitter is unambiguous, in practice, the data that is more commonly available for model inference (and that we anticipate generating from our prototyping) comes from arrays of 100s- 1000s of emitters. Since the scaling laws we are calibrating are derived for individual emitters, performing model inference on a data source that is a convolution of multiple emitters poses a unique challenge. Indeed, while we can argue that each individual emitter in an array is governed by the same scaling laws, variations in geometry of each emitter due to manufacturing tolerance will lead to different responses (e.g. emitted current versus voltage) for each emitter. The data that is available, however, is the summation of currents from each emitter. Thus, there exists the need for a method to analyze or deconvolve this data such that one can infer model parameters which describe the behavior of individual emission sites.

The goal of this work is to address this need by investigating a method for model inference of an array based on combining random sampling (to account for manufacturing tolerance) with Bayesian inference. To this end, this work is organized in the following way. In the first section, we describe the framework for modeling the current emitted by

an array of porous emitters as a function of voltage. We review the reduced-fidelity models for individual emitters and how we combine these models to represent an array of emitters. In the second section, we present our methods for performing model inference based on Bayesian regression. In the third section, we apply our approach to data extracted from measurements performed on real emitter arrays. In the fourth and final section, we discuss the physical significance of our results and extensions of our proposed methodology.

III. Reduced fidelity model for current from a porous emitter array

We motivate in this section a reduced fidelity model for the current as a function of voltage (I-V curve) emitted by an array of emitters fabricated on a porous substrate. We first describe the geometry for the system and its principle of operation. We then review the scaling laws that govern the current emitted by a single emitter with known geometry. We conclude by describing a methodology for combining these scaling laws to represent a full array.

A. Principle of operation

Fig. 1 shows the geometry and coordinate convention we adopt for modeling a porous conical emitter. The element has a cylindrical symmetry in a conical shape with key features including the radius of curvature of the tip, R_c , the cone half angle, α , the height of the emitter from the underlying substrate, h , the distance of the emitter tip to the extractor plate, d , the radius of the extractor aperture, R_A , and the typical radius of pores on the surface, r_p . During operation, ionic liquid is embedded in the porous substrate and drawn by capillary action to the surface. This liquid in turn pools in the pores of the emitter, each pool of liquid on the surface forming into a meniscus. A strong electric field is applied between the extractor plate and the emitter, which leads to an electrohydrodynamic force deforming the menisci subject to sufficiently strong field (those typically at the end of the emitter) into elongated cones. The actual contour of this shape is dictated by the balance between internal pressure in the droplet and the electrostatic pressure applied by the external field. As the field continues to increase, the electric field at the tip of the cone intensifies until eventually ions can be stripped directly from the liquid. These ions then are accelerated by the applied field downstream and out of the extractor aperture. If the electric field continues to increase, additional current can be extracted from the emitting cone. At the same time, other cones from adjacent pores continue to elongate until they too begin to emit. In this way, a single emitter can support multiple emission sites.




Fig. 1 Schematic for porous emitter illustrating key dimensions, the coordinate convention, and quantities relevant to the emitter model

B. Models related to current from a single emitter

We describe in this section models related to key processes that govern ionic emission. These include a model for how ionic current scales with electric field and a model for the onset of multiple emission sites.

1. Ionic Emission

The scaling law we employ for ionic emission is a modification of the work from Coffman et al. [13] who considered the response of a meniscus on a flat plate. We briefly review here the key steps of this model as well as the modifications we have introduced to better represent emission from a porous cone. To this end, we start by considering the case where ionic emission is already occurring from a population of emission sites on the i^{th} emitter of an emitter array. By continuity, the amount of current emitted from each site must be balanced by flow passing through the base of the emitting cone:

$$I_i^{(j)} = Q_i^{(j)} \rho \left\langle \frac{q}{m} \right\rangle, \quad (1)$$

where $I_i^{(j)}$ is the ionic emission from the j^{th} emission site of the i^{th} emitter, $j = 1, \dots, N_{s_i}$, $Q_i^{(j)}$ is the volumetric flow rate through the emission site cone, $\left\langle \frac{q}{m} \right\rangle$ is the charge to mass ratio of the emitted ions, and ρ is the mass density of the ionic fluid.

We in turn can relate the volumetric flow to the properties of the emitter. To this end, we denote the pressure inside the bulk of the substrate under the emitters as a “reservoir” pressure, p_r . We in turn denote the pressure at the base of the meniscus structure (i.e. a Taylor cone) $p_{b_i}^{(j)}$. The volumetric flow rate through the site is thus given through an effective equivalent circuit representation:

$$Q_i^{(j)} = \frac{p_r - p_{b_i}^{(j)}}{R_{H_i}^{(j)}}, \quad (2)$$

where $R_{H_i}^{(j)}$ denotes a characteristic hydraulic impedance for the flow path from the reservoir to the base of the emission site. Physically, this expression indicates that as the pressure differential increases across the flow path through the emitter, more flow can be accommodated. For a porous substrate the reservoir pressure is given from the characteristic pore radius of the reservoir substrate $r_{p,r}$,

$$p_r = -\frac{2\gamma}{r_{p,r}}, \quad (3)$$

where γ is the surface tension in the liquid, consistent with the treatment of [13, 14]. This expression would seem to suggest that there would be negative flow to the emission site; however by convention the electrostatic traction is taken to provide a negative internal pressure; following the work of Coffman et al, we represent this relationship symbolically as

$$-p_{b_i}^{(j)} = \zeta(E_i^{(j)}), \quad (4)$$

where $\zeta()$ is a monotonically increasing function of applied electric field. Physically, this expression suggests that as the electric field increases, the internal pressure decreases, facilitating more flow to the emitter. This is in keeping with the electrohydrodynamic nature of these devices. While Coffman [13] explored numerically the form of $\zeta()$, they also motivated a simple scaling law by expanding it around an onset electric field below which emission does not occur:

$$\zeta(E_i^{(j)}) \approx \zeta(E_{0_i}^{(j)}) + \left. \frac{\partial \zeta}{\partial E_i^{(j)}} \right|_{E_{0_i}^{(j)}} (E_i^{(j)} - E_{0_i}^{(j)}), \quad (5)$$

which, when combined with Eqs. (1,2) yields

$$I_i^{(j)} = \begin{cases} 0, & E_i^{(j)} < E_{0_i}^{(j)} \\ \frac{\rho \left\langle \frac{q}{m} \right\rangle}{R_{H_i}^{(j)}} \left(\zeta(E_{0_i}^{(j)}) + \left. \frac{\partial \zeta}{\partial E_i^{(j)}} \right|_{E_{0_i}^{(j)}} (E_i^{(j)} - E_{0_i}^{(j)}) + p_r \right), & \text{otherwise.} \end{cases} \quad (6)$$

The parameters of the equation above, and in particular the expansion of the pressure curve, are dimensional quantities entwined with other physics. Motivated by the desire to develop a form for this ionic emission model which is universal for emission sites, Coffman presents a nondimensionalization that we mirror here. It is useful to first introduce the following scale parameters,

$$\frac{1}{2}\varepsilon_0 \left(E_{c_i}^{(j)}\right)^2 = p_{c_i}^{(j)} = \frac{2\gamma}{r_{0_i}^{(j)}}, \quad (7)$$

where ε_0 is the permittivity of free space, γ is the propellant surface tension, and r_0 is a characteristic size of the emission site (i.e. the Taylor cone base radius). The quantity $p_{c_i}^{(j)}$ is thus a characteristic capillary pressure for the emission site, and $E_{c_i}^{(j)}$ is the characteristic electric field that produces an electric pressure equal to $p_{c_i}^{(j)}$. In general, r_0 may vary between sites, so we retain the indices here. Though we have assumed γ invariant over emitters and emission sites here, it could vary in principle (e.g. as a result of local temperature perturbations); however, it is expected from the results of [15] that these perturbations should be negligible. We these scale parameters identified, we proceed to defining dimensionless forms for the quantities of (6).

$\hat{I}_i^{(j)}$ is the dimensionless current emitted from the site,

$$\hat{I}_i^{(j)} = \frac{I_i^{(j)}}{KE_{c_i}^{(j)} \left(r_{0_i}^{(j)}\right)^2}, \quad (8)$$

where K is the conductivity of the propellant, which we also assume constant across emitters. $\hat{E}_i^{(j)}$ is the dimensionless applied background field, that is, the field strength at the emission site unperturbed by the presence of the meniscus,

$$\hat{E}_i^{(j)} = \frac{E_i^{(j)}}{E_{c_i}^{(j)}}. \quad (9)$$

Similarly, the term $\hat{E}_{0_i}^{(j)}$ is a dimensionless onset or starting voltage for the emission sites, at which the meniscus destabilizes into a Taylor cone structure and spray begins,

$$\hat{E}_{0_i}^{(j)} = \frac{E_{0_i}^{(j)}}{E_{c_i}^{(j)}}. \quad (10)$$

The dimensionless reservoir pressure $\hat{p}_{r_i}^{(j)}$ is given as

$$\hat{p}_{r_i}^{(j)} = \frac{p_r^{(j)}}{p_{c_i}^{(j)}}, \quad (11)$$

where p_r is the reservoir pressure feeding the emission site, taken here to be the same for all emission sites. $C_{R_i}^{(j)}$ is a dimensionless hydraulic impedance to the emission site,

$$C_{R_i}^{(j)} = \frac{KE_{c_i}^{(j)} \left(r_{0_i}^{(j)}\right)^2}{p_{c_i}^{(j)} \rho \langle q/m \rangle} R_{H_i}^{(j)}. \quad (12)$$

Here, ρ and $\langle q/m \rangle$ are both taken as uniform across the array. Finally, we note that the pressure dependence of Eq. 4 admits an equivalent dimensionless representation

$$-\hat{p}_{b_i}^{(j)} = p_{c_i}^{(j)} \hat{\zeta}(\hat{E}_i^{(j)}), \quad (13)$$

$$= \hat{\zeta}(\hat{E}_{0_i}^{(j)}) + \left. \frac{\partial \hat{\zeta}}{\partial \hat{E}_i^{(j)}} \right|_{\hat{E}_{0_i}^{(j)}} \left(\hat{E}_i^{(j)} - \hat{E}_{0_i}^{(j)} \right), \quad (14)$$

$$= \zeta_1 + \zeta_2 \left(\hat{E}_i^{(j)} - \hat{E}_{0_i}^{(j)} \right), \quad (15)$$

where we have collapsed the terms of the expansion into the model parameters ζ_1 and ζ_2 , which [13] suggests should be universal. Combined, these definitions yield a dimensionless form of the ionic emission model,

$$\hat{f}_i^{(j)} = \begin{cases} 0, & \hat{E}_i^{(j)} < \hat{E}_{0_i}^{(j)} \\ \frac{\zeta_1 + \zeta_2 (\hat{E}_i^{(j)} - \hat{E}_{0_i}^{(j)}) + \hat{p}_{r_i}^{(j)}}{C_{R_i}^{(j)}}, & \text{otherwise.} \end{cases} \quad (16)$$

2. Onset criterion

The emission model requires an onset criterion for computing $\hat{E}_{0_i}^{(j)}$. We implement here an onset criterion where the electric pressure is balanced with a characteristic capillary pressure of the emission site and with the internal pressure from the reservoir,

$$\frac{1}{2} \epsilon_0 \left(E_{0_i}^{(j)} \right)^2 = p_{c_i}^{(j)} - p_r. \quad (17)$$

That is to say, the emission site will deform into a meniscus capable of shedding current when the electrostatic pressure exceeds the restorative force provided by the capillary pressure and the reservoir, an extension of the classic onset criterion consistent with the observations of [14].

3. Hydraulic impedance

The emission model also requires a description of the hydraulic impedance. We suppose that the hydraulic impedance to each emission site is related to the total hydraulic impedance of the emitter R_{H_i} through their parallel sum,

$$R_{H_i} = \frac{1}{\sum_{j=1}^{N_{s_i}} \frac{1}{R_{H_i}^{(j)}}}, \quad (18)$$

where we estimate the emitter hydraulic impedance consistent with the treatment of Courtney [16],

$$R_{H_i} = \frac{\mu}{2\pi\kappa} \frac{1}{1 - \cos \alpha_i} \left[\frac{\tan \alpha_i}{R_{c_i}} - \frac{\cos \alpha_i}{h_i} \right], \quad (19)$$

where μ is the propellant viscosity, assumed uniform, α_i is the emitter cone half-angle, h_i is the emitter height (taken from the base of the cone to the virtual apex of the conical section). These quantities are again intrinsic to the emitter. κ is the substrate absolute permeability. The permeability can be estimated assuming spherical particles (see [16, 17]),

$$\frac{r_p^2}{60 (1 - \phi)^2}, \quad (20)$$

where r_p is the effective pore radius of the substrate, and ϕ its porosity (percent void volume). Depending on the packing of the substrate, the relationship between the effective pore radius and the distribution of pore radii in the substrate is nontrivial. Though these vary locally, we suppose that on the scale of the emitter they take some bulk value. Where the permeability is known directly, this estimation is unnecessary. The impedance estimate Eq. (19) models a radially inward flow across the porous cone, from a radius coincident with the cone base $\cos \alpha_i / h_i$ to the spherical cap of the emitter $\tan \alpha_i / R_{c_i}$, where the hydraulic impedance of the substrate between the emission site and the edge of the spherical cap are assumed to be negligible.

4. Multi-Site Emission Onset

The characteristic size of the emission sites must still be determined to close the emission model, and the model provides no prescription for the number of emission sites that develop. To estimate the size and number of emission sites for each emitter, we incorporate here an empirical model for the onset of multiple emission sites presented by St. Peter et al. [18] and implemented in the Electrospray Propulsion Engineering Toolkit (ESPET),

$$r_{0_i}^{(j)} = b_0 \left[1 - \frac{j-1}{N_{max_i}} \right]^{1/4} + r_{p_i}. \quad (21)$$

Here, r_{p_i} is a local pore size. While for estimating the hydraulic impedance of the emitter we suppose it sufficient to invoke some bulk measure of the pore radius, here we retain the emitter index i , indicating that the size of pores localized to the emitter tip could vary from emitter to emitter as an abstraction of the distribution of pore sizes in the substrate. In this model, b_0 is a “pooling” radius and N_{max_i} the maximum number of emission sites the emitter can support. This model indicates that as additional emission sites form on an emitter, the characteristic feature size of new emission sites decreases, eventually approaching the local pore size. Within this empirical abstraction, this may correspond to a scenario where propellant pools outside of pores instead of being anchored within them, or could correspond to the presence of larger pores on the substrate surface than the internal characteristic hydrodynamic pore size. In any case, the effect of Eq. (21) is to increase the emission site size beyond the nominal pore size r_{p_i} . b_0 is a parameter that must be calibrated from data, but N_{max_i} is estimated by comparing the area available for emission (taken to be the area of the spherical emitter tip) to the area of a site with radius an average of the first and ultimate sites, that is

$$N_{max_i} = \frac{2\pi R_{c_i}^2 (1 - \sin \alpha_i)}{\pi \left(\frac{(b_0 + r_{p_i}) + r_{p_i}}{2} \right)^2} \quad (22)$$

While Eq. (21) yields the size of the j^{th} site, it does not explicitly give the number of emission sites N_{s_i} . Instead, we compute the number of sites using the onset criterion of Eq. (17). That is, we suppose that each site having a characteristic feature size that satisfies the onset criterion will emit current. Inasmuch as Eq. (21) indicates the j^{th} emission site will have a characteristic radius r_{o_i} that monotonically decreases with j , the total number of emission sites N_{s_i} is equivalent to the largest index $j \leq N_{max_i}$ for which the onset criterion is satisfied. Thus, if the onset criterion is not satisfied for a site of size $b_0 + r_{p_i}$, the largest, then there will be no active sites. If the onset criterion is satisfied for a site of size r_{p_i} , the smallest, then there will be $\lfloor N_{max_i} \rfloor$ sites, where $\lfloor \cdot \rfloor$ is the floor function. In between these two edge cases, we can invert the onset criterion Eq. (17) to solve for the smallest feature size that could be activated r_{min_i} ,

$$r_{min_i} = \frac{2\gamma}{\frac{1}{2}\epsilon_0 E_i^2 + p_r}. \quad (23)$$

Thus, N_{s_i} is given symbolically

$$N_{s_i} = \max j = \begin{cases} 0, & \text{if } \frac{1}{2}\epsilon_0 E_i^2 < \frac{2\gamma}{b_0 + r_{p_i}} - p_r, \\ \lfloor N_{max_i} \rfloor, & \text{if } \frac{1}{2}\epsilon_0 E_i^2 > \frac{2\gamma}{r_{p_i}} - p_r, \\ \left\lfloor 1 + (N_{max_i} - 1) \left(1 - \frac{r_{min_i} - r_{p_i}}{b_0} \right)^4 \right\rfloor, & \text{otherwise.} \end{cases} \quad (24)$$

As the applied field becomes stronger, progressively smaller features activate. This expression thus predicts that as the emitter voltage is increased the number of active emission sites increases, as has been observed experimentally even for discrete emitters [7, 16, 18]. We illustrate this schematically for pores of different sizes in Fig. 2.

5. Electrostatic emitter model

The last model treated here is the electrostatic mapping between the voltage of the emitter V_i to the corresponding electric field. A classic approximation to electrospray emitter geometries is the hyperboloidal geometry of [19]. However, for the emitter geometries considered here, the gap distance between the electrodes approaches zero, for which the hyperboloidal approximation diverges, since it does not account for the presence of the aperture in the extractor electrode. To determine this mapping, then, we implement here an electrostatic simulation to determine exactly the field magnitude for the idealized emitter geometry of Fig. 1 as a function of the 5 relevant geometric parameters: d , R_c , α , h , and R_A . We display in Fig. 3 the axisymmetric simulation domain and a scalar field representing the electric field magnitude normalized by the voltage between the electrodes $\frac{|E|}{V}$ (units m^{-1}) for a nominal emitter design (described later). The domain extends from the centerline of the emitter to one-half the emitter pitch in the radial direction and from the basal plane of the emitter chip to one-half the emitter height beyond the far plane of the extractor in the axial direction. The emitter structure and basal plane were held at unity emitter potential relative to the grounded extractor geometry. In constructing the mesh we targeted a characteristic cell size of 1 micron at the emitter and extractor surfaces, but permitted this to relax to tens of microns in size in between them. As expected, for this design a high-field region is localized to the emitter tip, with the peak electric field strength on centerline.




Fig. 2 A notional mechanism underpinning the empirical scaling of (24), wherein the increasing electric field activates increasingly small pores in the substrate

C. Model for emitter array

We now seek to construct a model for an array of N_e emitters. We index these emitters $i = 1, 2, \dots, N_e$. The i^{th} emitter is biased to a potential V_i relative to the extractor electrode and emits a current I_i . We suppose that each emitter operates at the same voltage, the array voltage, $V_i = V$ (say, by construction of the distal electrode). The total current emitted by the array is thus given simply as a sum over the emitter currents,

$$I = \sum_{i=1}^{N_e} I_i. \quad (25)$$

Each of these unique emitters may host multiple active emission sites (that is, distinct Taylor cone structures), and this number N_{s_i} may differ from emitter to emitter, as predicted by the multi-site onset model. These individual sites are indexed $j = 1 \dots N_{s_i}$. The current sourced by each emitter is similarly the sum of the current emitted by each of its emission sites,

$$I_i = \sum_{j=1}^{N_{s_i}} I_i^{(j)}. \quad (26)$$

Finally, we note the implicit assumption that all N_e emitters are operated in parallel but independently, such that the behavior of one emitter does not influence the other emitters in the array.

We lastly simplify the analysis in two respects. Firstly, lacking a model for the radial position of emission sites, we suppose that each emission site is localized to the very tip of the emitter, or at least is close enough to the tip for there not to be an appreciable change in the field magnitude

$$E_i^{(j)} = E_i, \quad (27)$$

where E_i is this the tip field. Relatedly, we also suppose that the hydraulic impedance of each emission site is the same, such that, subject to Eq. (18),

$$R_{H_i}^{(j)} = N_{s_i} R_{H_i}. \quad (28)$$

These assumptions and the proceeding models close the model for the array to be only a function of design and operational parameters. That is, given the emitter voltage V , the substrate properties $(r_{p,r}, r_p, \phi)$, the propellant properties (γ, K, ρ, μ) , the beam properties $(\langle q/m \rangle)$, and the relevant emitter geometries $(d_i, R_{c_i}, R_{A,i}, \alpha_i, h_i)$, the current emitted from each emission site of each emitter is determined.




Fig. 3 Electrostatic simulation for a nominal emitter design; left: simulation domain and mesh; right: normalized electric field magnitude $\frac{|E|}{V}$

IV. Methodology for performing model inference

We next formulate an inference problem over the model parameters ζ_1 and ζ_2 of the emission model and parameter b_0 of the multi-site onset model. These expressions predict the current sourced by a single emitter. However, we may wish to learn these parameters from data captured on an array of emitters. In principle, doing so could be accomplished by a simple scaling down of the current of the array by N_e , the number of emitters in the array. However, no two emitters in the array will be functionally identical, a result of the finite precision with which they can be manufactured. Within this uncertainty, the nonlinear behavior admitted by the preceding models indicates the mean behavior over the array may not be simply related to the behavior of the nominal design. This motivates an inference framework that can rigorously describe our state of knowledge in the parameters of interest given the other uncertain parameters of the model.

A. Bayesian Inference Framework

We begin by considering that we have a set of N_d data X , which are indexed $k = 1, 2, \dots, N_d$, such that $X = \{x^{(k)}\} = \{x^{(1)}, x^{(2)}, \dots, x^{(k)}\}$ (not to be confused with the index j , which indexes emission sites). Each datum is a pair of voltage and current measurements, $x^{(k)} = (V^{(k)}, I^{(k)})$, where the voltage data are taken as an independent variable (as they are typically prescribed and known very precisely) and the current data are taken as a dependent variable. We wish to infer the collection of parameters $\theta = (\zeta_1, \zeta_2, b_0)$ given the data X . Bayes' theorem provides a framework for estimating the probability distributions for these models parameters in light of the experimental data:

$$P(\theta | X) \equiv \frac{P(X | \theta) \times P(\theta)}{P(X)}, \quad (29)$$

where here P denotes probability. We are interested in computing the posterior probability $P(\theta | X)$, which expresses our state of knowledge in the parameters θ given our observed data X . Since our state of knowledge is rarely certain, this and the other terms take the form of probability distributions (i.e. a varying degree in belief in different values of the parameters). For our purposes, the evidence, $P(X)$, can be discarded as a normalizing factor and need not be computed, yielding

$$P(\theta | X) \propto P(X | \theta) \times P(\theta). \quad (30)$$

The prior probability, $P(\theta)$, encodes our state of knowledge in the parameters prior to observing the data X . In absence of previous calibration data, this may be chosen to be comparatively uninformative. The likelihood, $P(X | \theta)$, encodes the probability of observing the data X given the model parameters were known and practically manifests as a noise model for how we believe the experimental data should be distributed about our model predictions, supposing the model is correct. Bayes' theorem thus provides a method of updating our understanding in the model parameters in light of our observations.

B. Pseudomarginalization of the Likelihood

The prior distribution over parameters is typically either chosen or provided from previous analysis, and so the key task in applying Bayes' theorem is the computation of the likelihood, and hence, the posterior. In our case, the likelihood cannot be computed directly; in addition to the parameters θ , the emission and multi-site onset models are also functions of the substrate properties $(r_{p,r}, r_p, \phi)$, the propellant properties (γ, K, ρ, μ) , the beam properties $(\langle q/m \rangle)$, and the relevant emitter geometries for each emitter $(d_i, R_{c_i}, R_{A_i}, \alpha_i, h_i)$. That is to say, we can only compute the likelihood $P(X | \theta, G)$, where G is collection of these other quantities,

$$G = ((r_{p,r}, r_p, \phi), (\gamma, K, \rho, \mu), (\langle q/m \rangle), (R_{c_1}, d_1, R_{A_1}, \alpha_1, h_1, r_{p_1}), \dots, (R_{c_{N_e}}, d_{N_e}, R_{A_{N_e}}, \alpha_{N_e}, h_{N_e}, r_{p_{N_e}})), \quad (31)$$

These are nuisance parameters; that is, within our inferential context we do not care, ultimately, about their value, but they are necessary to compute our model predictions and their uncertainty will muddy our understanding of the model parameters. We must obtain the likelihood needed for Eq. (30), then, by marginalizing over these nuisance parameters,

$$P(X | \theta) = \int P(X, G | \theta) dG = \int P(X | \theta, G) \times P(G | \theta) dG. \quad (32)$$

Thus, we have substituted evaluating the integral above over directly calculating the likelihood.

To compute the probability $P(X | \theta, G)$, we assume that each current datum $I^{(k)}$ is normally distributed about the model prediction $F(V^{(k)}; \theta, G)$ with standard deviation $\sigma^{(k)}$, the noise in the k^{th} datum, which is taken to be the uncertainty reported from experiment. Assuming the noise in each datum is independent, then, the likelihood is calculated as

$$P(X | \theta, G) = \prod_{k=1}^{N_d} \frac{1}{\sigma^{(k)} \sqrt{2\pi}} \exp \left[-\frac{(F(V^{(k)}; \theta, G) - I^{(k)})^2}{2(\sigma^{(k)})^2} \right]. \quad (33)$$

If the quantities of G were known exactly (i.e., distributed as a delta function), the marginalization integral of Eq. (32) would be trivial. However, since these parameters may be appreciably uncertain, this integral is nontrivial and, given the complexity of the models, analytically intractable. Indeed, for an array of N_e emitters, the integral as written is over an $(8 + 6N_e)$ -dimensional parameter space, which for many electrospray arrays would be a few thousand.

To simplify this integral we assume that the nuisance parameters G are independent from the model parameters θ ,

$$P(G | \theta) = P(G). \quad (34)$$

That is, knowledge of the model parameters θ does not modify our understanding of the nuisance parameters, which is characterized by other measurements (like microscopy to determine emitter geometries or spectrometric measurements of the beam composition). The integral Eq. (32) is still largely intractable despite this simplification, and so a practical solution we adopt here is to instead form a Monte-Carlo estimator that approximates the integral,

$$P(X | \theta) \approx \hat{P}(X | \theta) = \frac{1}{N_r} \sum_{l=1}^{N_r} P(X | \theta, G^{(l)}), \quad G^{(l)} \sim P(G). \quad (35)$$

Here, $\{G^{(l)}\}$ is a set of $l = 1, \dots, N_r$ samples from the distribution over nuisance parameters $P(G)$. Functionally, this means we randomly generate N_r arrays of fixed geometry (and other parameters), compute the likelihood of the training data for that array, and then average the estimates for the likelihood, as is shown schematically in Fig. 4. In the limit $N_r \rightarrow \infty$, this is identical to the true value (provided the variance is finite). We now have the means to calculate (or approximately calculate) the posterior as a function of the values of the parameters θ , provided the distribution over nuisance parameters $P(G)$ is specified.




Fig. 4 Schematic diagram of pseudomarginalization; the groups of the emitters at the bottom are different samples from the geometry, $G^{(l)}$; each emitter is predicted to produce an associated current $I_i^{(l)}$ (small blue arrows), which combined predicts a total current $I^{(l)}$ for the array (larger blue arrows); each of these predictions for the different geometry is compared with the measured datum $I^{(k)}$ (red arrows) to determine its likelihood; these likelihoods are then averaged to form the pseudomarginalization

In principle, we could fully characterize the posterior by evaluating it over all parameter space, but doing so to sufficient resolution is computationally burdensome, as, even after marginalizing the likelihood over the nuisance parameters, the posterior resides in a 3-dimensional space. For example, resolving even 1000 points along each dimension would require 10^9 evaluations of the posterior. To avoid this burden, we instead seek to sample from the posterior, that is, to generate a population of sets of the model parameters that is statistically representative.

C. Sampling from the Posterior

A sampling methodology particularly well-suited to Bayesian inference problems is Metropolis-Hastings Markov Chain Monte Carlo (MH-MCMC). We consider here an MH-MCMC sampler with three adaptations, a delayed rejection (DR) scheme, an adaptive Metropolis (AM) scheme, and compatibility with pseudomarginalization (PM) over the nuisance parameters. The corresponding pseudomarginal delayed rejection adaptive Metropolis algorithm (PM-DRAM) is described in more detail in the appendix.

D. Generating posterior predictions

After sampling from the posterior, we seek to use this knowledge to make probabilistic performance predictions with credible intervals. To do so, we form the posterior predictive distribution,

$$\begin{aligned}
 P(X^* | X) &= \int P(X^*, \theta | X) d\theta, \\
 &= \int P(X^* | \theta, X) \times P(\theta | X) d\theta, \\
 &= \int P(X^* | \theta) \times P(\theta | X) d\theta.
 \end{aligned} \tag{36}$$

Here, X^* represents new data that we have not observed and $P(\theta | X)$ is simply the posterior. $P(X^* | \theta)$ is the likelihood (i.e. our noise model), but computed for the new data, where the elimination of X is because the new data are




Fig. 5 Schematic diagram of sampling from the posterior predictive distribution; we first make a model prediction for each sample from the posterior to capture uncertainty in the model parameters; then, to simulate measuring these data with a real instrument, we corrupt these predictions by the experimental noise

conditionally independent from the observed data given the model parameters θ . Given samples from the posterior, we can sample from the posterior predictive by drawing from $P(X^* | \theta)$ conditional on those samples; however, in our case, we cannot sample from $P(X^* | \theta)$ directly as a result of G . Instead, maintaining the pseudomarginalization from earlier, we express

$$\begin{aligned}
 P(X^* | \theta) &= \int P(X^* | G, \theta) \times P(G | \theta) dG, \\
 &= \int P(X^* | G, \theta) \times P(G) dG, \\
 \approx \hat{P}(X^* | \theta) &= \frac{1}{N_r} \sum_{l=1}^{N_r} P(X^* | \theta, G^{(l)}), \quad G^{(l)} \sim P(G),
 \end{aligned} \tag{37}$$

where the intermediate equality follows by assumption of independence between the model and nuisance parameters. Practically, we can draw samples from $\hat{P}(X^* | \theta)$ by first selecting from one of the N_r $G^{(l)}$'s at random and then drawing from the noise model $P(X^* | \theta)$ (i.e. corrupting the model prediction by the experimental noise). We present a simplified representation of this process in Fig. 5. Thus, subject to the pseudomarginal approximation, we can draw samples from the posterior predictive from which to form credible intervals.

V. Description of the dataset

We perform our model inference using data published by Natisin et al for a unit of the Air Force Research Laboratory's AFET-2 thruster, a picture of which is included from [3] in Fig. 6. The design of the AFET-2 is described more thoroughly in the listed reference, but we provide here an overview. The primary component of the AFET-2 is a chip of 576 nearly conical emitters machined from porous borosilicate glass in a rectilinear grid. This emitter array is interfaced with a porous borosilicate reservoir and embedded in an isolated propellant module such that the chip and reservoir are retained within the module by a distal electrode. Placed above the emitter chip is an extractor electrode consisting of two pieces, a thin metallic sheet with a circular aperture for each emitter bonded onto a thicker frame to provide structural support.

In addition to publishing current and voltage data for the thruster, Ref. [3] contains a detailed description of the thruster manufacturing processes, which we use to inform the nuisance parameters G . The substrate parameters and local pore radius uncertainties are taken from pore size distributions quoted by the manufacturer (ROBU Glasfilter-Geräte GmbH, P4 and P5). The propellant property uncertainties are retrieved from the Electrospray Propulsion Engineering Toolkit propellant database [18] and correspond to a 292-298 K temperature window. The charge-to-mass ratio of the spray is taken from TOF spectra reported by Natisin et al alongside the current and voltage data [3]; note that this value corresponds to positive thruster polarity. Lastly, the remaining emitter geometries are determined from Ref. [3]'s description of the manufacturing process and published optical micrographs of the emitter array.




Fig. 6 A picture of an AFET-2 tested at the AFRL, with overall dimensions labeled, from [3]

We collect in Table 1 the components of $P(G)$ for each parameter, which we approximated as independent. While this is not strictly true, it is a useful expedience, as the covariance between the emitter geometries, for example, is determined by multiple interdependent physical processes related to their manufacture and would require a more detailed characterization of the emitter array than that available here to assess.

Nuisance Parameter	Distribution $P(G)$	Units
$r_{p,r}$	$\mathcal{U}[5.000 \times 10^{-6}, 8.000 \times 10^{-6}]$	m
κ	$\mathcal{N}(1.510 \times 10^{-13}, 6.040 \times 10^{-15})$	m^{-2}
K	$\mathcal{U}[1.147, 1.390]$	S m^{-1}
γ	$\mathcal{U}[5.003 \times 10^{-2}, 5.045 \times 10^{-2}]$	N m^{-1}
ρ	$\mathcal{U}[1.280 \times 10^3, 1.284 \times 10^3]$	kg m^{-3}
μ	$\mathcal{U}[2.612 \times 10^{-2}, 3.416 \times 10^{-2}]$	N s m^{-2}
$\langle q/m \rangle$	$\mathcal{N}(5.500 \times 10^5, 1.003 \times 10^4)$	C kg^{-1}
R_{C_i}	$\mathcal{U}[1.000 \times 10^{-5}, 2.000 \times 10^{-5}]$	m
d_i	$\mathcal{N}(3.000 \times 10^{-6}, 5.230 \times 10^{-6})$	m
R_{A_i}	$\mathcal{N}(2.486 \times 10^{-4}, 3.596 \times 10^{-6})$	m
α_i	$\mathcal{N}(2.678 \times 10^{-1}, 4.000 \times 10^{-3})$	rad
h_i	$\mathcal{N}(3.018 \times 10^{-4}, 5.130 \times 10^{-6})$	m
r_{p_i}	$\mathcal{N}(6.500 \times 10^{-7}, 7.500 \times 10^{-8})$	m

Table 1 Nominal uncertainties for nuisance parameters in G , $P(G)$, assumed independent. For uniform distributions, the two values are the upper and lower bounds; for normal distributions, the first value is the mean while the second is the standard deviation.

To encode our comparative ignorance in the parameters $\theta = (\zeta_1, \zeta_2, b_0)$ before learning from the data, we choose uniform prior distributions, except that the parameters are constrained to be positive, which we formalize in Table 2. Practically, then, the priors can be absorbed into the proportionality of Eq. (30), except that nonpositive parameter values were disallowed by the sampler (i.e. are assigned arbitrarily low likelihood).

VI. Results and Analysis

A. Perfect emitter Inference

As an illustrative comparison, we first considered a reduced inference problem over a "perfect" array; that is, we sought a fit to the parameters θ ignoring uncertainty in the nuisance parameters G and instead assuming each to be

Parameter	Prior Distribution $P(\theta)$
ζ_1	$\mathcal{U}(0, \infty)$
ζ_2	$\mathcal{U}(0, \infty)$
b_0	$\mathcal{U}(0, \infty)$

Table 2 Prior distributions considered for the model parameters $\theta = (\zeta_1, \zeta_2, b_0)$

equal to the mean of its respective distribution. In this way, the array is effectively treated as having been manufactured to exact specification, and there is no variability in performance between emitters. Thus, it is equivalent to considering a single emitter whose behavior is multiplied $N_e = 576$ times over. Mathematically, G approaches the limit of a delta function, and the marginalization collapses and becomes unnecessary.

For this inference problem, we identified the maximum a-posteriori (MAP) set of parameters to be $(\zeta_1, \zeta_2, b_0) = (1.18, 7.35 \times 10^{-3}, 7.00 \times 10^{-6} \text{ m})$. The corresponding model prediction for the array is displayed in Fig. 7 alongside experimental data from [3]. The behavior of the prediction is consistent with onset of each emitter within the array at




Fig. 7 Predicted performance for a perfect emitter array; the dashed black is the current predicted using the MAP model parameters for an array of 576 emitters of the mean geometry of Table 1; in blue are data reported by [3]

just above 1700 V and a linear increase in the current with voltage, consistent with the onset criterion (17) and the ionic emission model (16). While these predictions are able to capture the approximately linear behavior of the thruster in the high-voltage regime, they cannot do so for the lower voltage set points. In reality, there exist emitters (i.e. those with smaller R_c and correspondingly higher tip fields that), other things equal, would be predicted by the model to start at lower voltages. As a concrete example, the electrostatic simulation depicted in 3, for the mean geometry of Table 1 ($R_c = 1.5 \times 10^{-5}$), produces a normalized field of $4.64 \times 10^4 \text{ m}^{-1}$. Keeping other parameters the same but decreasing R_c in the electrostatic model to $1 \times 10^{-5} \text{ m}$, the smallest considered, results in a field per voltage of $6.17 \times 10^4 \text{ m}^{-1}$, an additional 33% increase that would reduce the onset to 1270 V for the same model parameters.

B. Approximately pseudomarginal inference for a single realization

To rigorously account for the uncertainty in the nuisance parameters, then, we extended the inference problem to include the pseudomarginalization of Sec IV.B. Subject to our prior observations considering the performance of the perfect array, we utilized the electrostatic simulation to calculate the normalized electric field for each geometry.

However, the electrostatic model is comparatively expensive to evaluate, $O(1)$ s of CPU time. Thus, to draw N_f samples from the posterior would require $O(N_r \times N_e \times N_f)$ s of CPU time, independent of other overhead, which is impractically long for our purposes, where even for $N_r = 1$ this would be $O(10^8)$. To alleviate this computational burden, we instead implemented an approximate pseudomarginalization over the nuisance parameters; that is, instead of sampling N_r realizations of G and solving the electrostatic simulation for each geometry each time the likelihood was evaluated, we sampled N_r realizations of G before drawing from the posterior and used this same population of samples each time we needed to calculate the likelihood. This eliminates the scaling with N_f , but also introduces some bias to the posterior, since we are fixing the $G^{(l)}$'s.

We began by drawing samples from this approximately pseudomarginal posterior taking $N_r = 1$. The maximum identified by initialization with the optimizer was $(\zeta_1, \zeta_2, b_0) = (2.78, 1.96 \times 10^{-2}, 1.94 \times 10^{-5} \text{ m})$, which is substantially different than that identified from the perfect array. We present a corner plot of the samples in Fig. 8. This plot consists of the lower triangle of a 3×3 grid and mimics the structure of a covariance matrix between the 3 parameters. The diagonal entries are 1-d histograms of the marginal distribution over the corresponding parameter,

$$P(\zeta_1 | X) = \int P(\zeta_1, \zeta_2, b_0 | X) d\zeta_2 db_0, \quad (38)$$

and similarly for the other two parameters. That is, these histograms show the distribution of the parameter irrespective of the values of the others. The off-diagonal entries are 2-d histograms of the corresponding 2-d marginals,

$$P(\zeta_1, \zeta_2 | X) = \int P(\zeta_1, \zeta_2, b_0 | X) db_0, \quad (39)$$

and so forth. Thus, these entries display the correlation between each pair of parameters. Our uncertainty in the parameters is expressed through the characteristic width in these distributions. However, this case does not include uncertainty in the parameters resulting from uncertainty in G because it uses a single, fixed sample. Nevertheless, it is illustrative in that it represents the emitter array as a group of emitters with distinct geometries. Given a population of samples, the marginalizations of Eqs. (38,39) are performed trivially by ignoring the unrepresented parameter(s) for each sample. The distributions are peaked at a point slightly different than the maximum identified by the initializer, with multimodality in the distributions for ζ_1 and b_0 . Most notably, there is a strong positive correlation between ζ_1 and b_0 , though aside from the secondary mode the samples are narrowly distributed around the peak.

Taking every tenth sample from the posterior, we then generated samples from the corresponding posterior predictive distribution consistent with the description of Sec. IV.D. We predicted at the same locations as the measured data, for which the noise is taken from experiment. We display in Fig. 9 the envelope of these predictions alongside the training data. The predicted behavior largely agrees with the experimental data, producing a characteristic inflection in the current response. The predictions appear to disagree most substantially with the experimental data in the intermediate voltage range between 1100 and 1400 V, where the data appear outside of the uncertainty in the predictions. However, as previously mentioned the uncertainty in this prediction is not representative of the full problem as a result of having selected only $N_r = 1$.

C. Approximately pseudomarginal inference for 100 realizations

Correspondingly, we then performed the same inference problem but instead taking $N_r = 100$, that is, generating 100 samples from the nuisance parameters to marginalize over. The initial sample identified by the optimizer was $(\zeta_1, \zeta_2, b_0) = (2.71, 1.94 \times 10^{-2}, 1.93 \times 10^{-5} \text{ m})$, which is close to the same optimum identified for the single-realization case on the scale of the variance in its distribution. We show in Fig. 10 the corresponding corner plot. The character of the posterior is largely the same, with a few key differences. Firstly, the peak of the posterior is noticeably different than that for the single realization, and the posterior is substantially broader on the scale of the parameters. Indeed, it seems that the $\zeta_1 - b_0$ subspace of the $N_r = 1$ case is a subset of that of the $N_r = 100$ case. Additionally, the multimodality observed earlier appears also in this case, and in approximately the same region of parameter space, though with a different relative posterior probability between the modes. Lastly, while ζ_1 and b_0 maintain their high correlation, ζ_2 is now virtually independent from the other parameters.

We then drew samples from the posterior predictive distribution for the $N_r = 100$ case. A plot of the predictions is shown in Fig. 11. As is evident in the figure, the inclusion of more than one realization of the nuisance parameters has resulted in a dramatic increase in the uncertainty in the predictions. Indeed, it indicates that there exist model predictions within significant uncertainty that yield a current two orders of magnitude greater than that sourced by the real device,




Fig. 8 Corner plot of the posterior for the approximately pseudomarginal inference problem with $N_r = 1$; the diagonal entries are the 1d marginal distributions of each parameter, while the off-diagonal entries are 2d marginal distributions showing correlation between the corresponding pairs of parameters

which seems nonphysical. Upon detailed inspection of the model for these values of the parameters, we found this to be caused by the reservoir pressure p_r . Here, we have treated the reservoir pressure as uncertain to represent our lack of knowledge in its exact value, though we expect that its magnitude should be approximately determined by the characteristic pore diameter of the reservoir substrate, consistent with (3). However, the parameter ζ_1 was determined to be highly correlated with this reservoir pressure, with the model therefore highly sensitive to deviations from this correlation. As a consequence, within our inference framework the assumption $P(G | \theta) = P(G)$ is not satisfied when considering the reservoir pressure, resulting in the behavior of Fig. 11.

D. Approximately pseudomarginal inference for 100 realizations excluding reservoir pressure

To combat this problem, we performed the approximately pseudomarginal inference again for $N_r = 100$ but considering the limit where the reservoir pore radius takes a deterministic value $r_{p,r} = 8 \times 10^{-6}$ m, that is, the largest characteristic pore size of the P4 substrate considered here. We drew samples from the posterior and posterior predictive distributions as before. The matching corner plot and prediction plots are shown in Figs. 12 and 13, respectively.

In this case, the initial point determined by the maximizer was $(\zeta_1, \zeta_2, b_0) = (2.57, 1.69 \times 10^{-2}, 2.00 \times 10^{-5})$ m. The posterior distribution is largely similar to the previous $N_r = 100$. However, the bimodality evident previously has narrowed, being mostly replaced with a tail in ζ_1 and b_0 . Additionally, the MAP occurs in a segment of the $\zeta_1 - b_0$ subspace at lower magnitude than observed previously. This is consistent with the positive correlation between p_r and these parameters, since p_r has been decreased by taking the larger limit on size of the reservoir pores. ζ_2 is still largely independent from the other parameters and takes values modestly smaller than previously.




Fig. 9 Posterior predictive distribution at the positions of the training data for the approximately pseudomarginal inference with $N_r = 1$; in blue are the data of [3]; the solid black is the median prediction, while the black dashed lines contain 90% of all predictions

The corresponding posterior predictions are able to retrieve the behavior as before, but without the nonphysical behavior observed earlier. Accounting for the uncertainty in the nuisance parameters has otherwise broadened the predictive envelope compared with that of the $N_r = 1$ case, such that all the training data are found within the 90% credible interval. The last feature of note is the broadening of the credible interval at the final datum, which indicates that for this voltage there is a relatively small population of samples with a substantially higher predicted current (as might correspond to onset of an additional emission sites for many emitters in the array).

E. Posterior predictions for negative polarity

Lastly for the inference problem, we sought to predict the behavior of the same thruster, but operating in the negative polarity mode. Doing so represents a small test of the extensibility of the parameters learned previously to data conditions outside the training set. We formed a posterior predictive distribution as before, but predicting for the negative voltage and current data reported by [3]. We reuse the same samples from the nuisance parameters, except that we drew new samples for the beam properties (the charge to mass ratio $\langle \frac{q}{m} \rangle$) to reflect the ionic content of the negative polarity spray. The corresponding distribution is $\langle \frac{q}{m} \rangle \sim \mathcal{N}(6.2787 \times 10^5, 1.4371 \times 10^4) \frac{\text{C}}{\text{kg}}$. We collect the posterior predictions made for these data in Fig. 14. The predictions indicate the model parameters have captured the behavior of the thruster when operating in the negative emission mode, within the uncertainty in the model parameters and the associated experimental noise. This is consistent with the scaling of Eq. (16), which predicts that negative emission mode should differ only from the positive mode by the difference in their charge-to-mass ratios.

VII. Discussion

A. Array uncertainty and thruster scaling

We begin our discussion by examining the disparate results for the model behavior between treating the emitter array as certain (i.e. a "perfect" array) and considering it as uncertain (i.e. a "real" array). Assuming a nominal emitter geometry, predicting the behavior of that geometry, and then scaling this prediction by the number of emitters in the system is a simple and natural strategy. However, we observed, as evidenced by Fig. 7, that this type of scaling was unable to capture the measured behavior of the array, at least for the modeling framework and thruster treated here. In contrast, as shown in Fig. 9, accounting for differences in emitter geometry achieved both qualitative and quantitative




Fig. 10 Corner plot of the posterior for the approximately pseudomarginal inference problem with $N_r = 100$; the diagonal entries are the 1d marginal distributions of each parameter, while the off-diagonal entries are 2d marginal distributions showing correlation between the corresponding pairs of parameters

agreement with experimental results over the domain.

This is not an unexpected result in that it follows directly from tolerances in the emitter geometry being significant on the scale of the problem. This is perhaps most easily illustrated by the disparate magnitudes in the electric field we calculated by electrostatic simulation for two emitters differing only in their radius of curvature, which implied correspondingly disparate onset voltages. However, it is an interesting result inasmuch as it suggests that the behavior of these devices may fundamentally be predicated on this uncertainty. We note that there exists no nonlinear mechanism in the ionic emission scaling Eq. (16). Although the multi-site onset model Eq. (24) predicts that additional emission sites will develop on a single emitter at nonlinear intervals, the predicted behavior between these discrete jumps is otherwise linear, as has been observed in experiments with single electrospray emitters [7, 20]. Within this framework, then, the only physical mechanism admitting the smooth, nonlinear current response observed is that it is formed from the superposition of a large population of emitters, each of which turns on at a different voltage and scales differently from there. Indeed, we can observe from the data of [3] that the thruster behavior becomes approximately linear in the high-voltage regime, as is evident elsewhere [5], and might be a result of a saturation point where all of the emitters have already onset and are sourcing additional current consistent with (16). Given these insights, we conclude that, while in the limit of perfect manufacturing this process would be expected to disappear, for real systems (cf. [5, 8, 14]) it is a nonnegligible effect.

Additionally, accounting for the uncertainty in the nuisance parameters is vital to rigorously quantifying uncertainty in the model parameters and thus in using them to form predictions for other designs. As an example, we highlight the differences between the posterior predictions of Fig. 9, for which $N_r = 1$, and Fig. 13, for which $N_r = 100$. While both




Fig. 11 Posterior predictive distribution at the positions of the training data for the approximately pseudo-marginal inference with $N_r = 100$; in blue are the data of [3]; the solid black is the median prediction, while the black dashed lines contain 90% of all predictions

of these inferences reproduce the behavior of the data, because the $N_r = 1$ implicitly takes the nuisance parameters to be deterministic, the corresponding uncertainty in its model predictions is substantially lower, misrepresenting our true state of knowledge in the model parameters, which is diluted by our limited understanding of the nuisance parameters. Nevertheless, we have been able to demonstrate that, properly treated, thruster measurements can be used to effectively draw inferences about the model parameters. Within the context of a broader design and inference loop, for example the robust design problem, this analysis thus provides a method of incorporating data from thruster tests as they are performed. It additionally serves as a tool for making strategic decisions within an associated optimal experimental design problem. That is, in seeking to calibrate these models further and extend their results to a broader class of thrusters, we wish to identify what experiments we should perform to most efficiently calibrate these models (i.e. to determine an experimental configuration η that maximizes some utility function $\mathcal{U}(\eta)$, which takes the form of a reduction in the parameter variance). By leveraging the analysis presented here, we can assess the comparative utility of thruster designs that we may seek to manufacture, subject to the finite tolerance with which we can manufacture them.

B. Implications for emitter model development

We now discuss the difficulties uncovered in performing our analysis and their implications for further model development. Firstly, our analysis suggests that the choice of electrostatic model is an integral part of modeling the behavior of these systems. For the models considered here, the field strength at the emitter serves both to dictate the onset of one or more emission sites and modulate the emitted current. Within the context of variable emitters discussed previously, it is important to capture the dependence of the field on the uncertain emitter geometries. However, as we noted earlier, the classic hyperboloidal approximation loses validity for the geometry examined here. While we were able to substitute an electrostatic simulation to determine the field exactly within the limit of an abstraction of the emitter geometry, we did by making concessions elsewhere (i.e. introducing bias to the pseudomarginalization). This trade may not be sufficient in other contexts, however, including in leveraging these models to search for robust emitter designs, where one must consider uncertainty at every potential point in design space, as opposed to at just one. Problems such as these, then, would benefit from a new scaling law for the electric field that accommodates these nontraditional electrode configurations.

In performing our analysis we also observed a strong correlation between the offset parameter of the ionic emission scaling law ζ_1 and the reservoir pressure p_r . This suggests that the parameter ζ_1 , which arises from a linear expansion of the pressure-electric field dependence of an emission site, should itself be some function of the (dimensionless)




Fig. 12 Corner plot of the posterior for the approximately pseudomarginal inference problem with $N_r = 100$ and the reservoir pore radius taken to be $r_{p,r} = 8 \times 10^{-6}$ m; the diagonal entries are the 1d marginal distributions of each parameter, while the off-diagonal entries are 2d marginal distributions showing correlation between the corresponding pairs of parameters

reservoir pressure $\zeta_1 = \zeta_1(\hat{p}_r)$, warranting a reparameterization of the emission model. A simple such choice would be to remove a factor \hat{p}_r from ζ_1 , such that $\zeta_1 = -\zeta_1^* \hat{p}_r$ and therefore

$$\hat{I} = \frac{(1 - \zeta_1^*) \hat{p}_r + \zeta_2 (\hat{E} - \hat{E}_0)}{C_R}, \quad (40)$$

which would allow a fixed value of ζ_1^* to have much the same behavior as a value of ζ_1 correlated with \hat{p}_r would. This dependence likely arises from the pressure dependence expansion, which is performed about the onset voltage. This onset voltage scales with the reservoir pressure (cf. [14, 15, 21]) and, thus, ζ_1 may only succeed as a dimensionless representation within its dependence on the dimensionless reservoir pressure \hat{p}_r . As additional motivation for this scaling, we note our previous work considering model inference for single-emitter tests reported by Perez-Martinez [7, 20], where we could not consolidate the behavior between the emitters which did not consider as dimensionless a representation of the current emission scaling as that presented here.

VIII. Conclusion

In this work, we presented a model for the current emitted by an array of porous conical electrospray emitters. We then formulated a Bayesian inference problem over the model parameters which rigorously considers uncertainty in the emitter geometry and other nuisance parameters which appear in the modeling framework. Using data from




Fig. 13 Posterior predictive distribution at the positions of the training data for the approximately pseudo-marginal inference with $N_r = 100$ and the reservoir pore radius taken to be $r_{p,r} = 8 \times 10^{-6}$; in blue are the data of [3]; the solid black is the median prediction, while the black dashed lines contain 90% of all predictions

the Air Force Research Laboratory’s AFET-2 thruster, we sampled from the corresponding posterior distribution over parameters. In doing so, we determined that representing the emitter array as a collection of emitters each of the nominal design of the array fundamentally failed to reproduce the observed behavior. Conversely, considering uncertainty in the emitter geometries was able to match the observed behavior and suggests that the nonlinear emission behavior of such systems in low-voltage regimes may be predicated on the associated lack of uniformity in onset and emission.

Additionally, by adopting this marginal description over uncertain model parameters, we have been able to more rigorously describe uncertainty in our understanding of the model parameters, of key importance both in making predictions for designs outside of the dataset, as we did for the opposite polarity emission mode of the thruster, or for determination of an optimal experimental design. We also uncovered through this analysis a dependence between the reservoir pressure of the array and the offset parameter of the ionic emission scaling law, which motivates a reparameterization of the model to disentangle the correlated physics.

Acknowledgments

This work was supported by an Early Stage Innovations grant from NASA’s Space Technology Research Grants Program and a NASA Space Technology Graduate Research Opportunity fellowship. The authors would like to thank Joshua Eckels for constructing the electrostatic simulation used in this work and also Dr. Rainer Dressler and Mr. Benjamin St. Peter of Spectral Sciences Inc., Dr. John Yim of the NASA Glenn Research Center, and Dr. Michael Natisin and Dr. Dan Eckhardt of the Air Force Research Laboratory for their useful discussions.

Appendix: PM-DRAM Sampler

A. Sampler Nomenclature

- f_θ = unnormalized target distribution
- s = sample
- f_1 = primary proposal distribution
- \bar{s}_1 = primary proposed sample

A_1	=	primary proposal acceptance probability
n	=	sampler iteration index
\tilde{s}_2	=	secondary proposed sample
f_2	=	secondary proposal distribution
A_2	=	secondary proposal acceptance probability
n_0	=	adaptive Metropolis burn-in
\tilde{f}_1	=	adaptive primary proposal
\tilde{f}_2	=	adaptive secondary proposal
z	=	stored target distribution evaluations
\tilde{z}_1	=	primary proposal target evaluation
u_1	=	primary proposal acceptance probability sample
\tilde{z}_2	=	secondary proposal target evaluation
u_2	=	secondary proposal acceptance probability sample
S	=	proposal covariance
C	=	sample covariance
s_d	=	sampling dimensional scale factor
ξ_{AM}	=	adaptive Metropolis conditioning factor
D	=	sample dimensionality
I	=	identity matrix
γ_{DR}	=	delayed rejection scale factor

B. Sampling Methodology

The pseudomarginal delayed rejection adaptive Metropolis sampler implemented in this work is summarized in Algorithm 1.

This algorithm samples from a target distribution f_θ that need only be known within a multiplicative constant, making it attractive for Bayesian parameter estimation (where the evidence is often discarded as a normalizing constant). It begins with an initial sample $s^{(0)}$. At each step, it proposes a new sample \tilde{s}_1 from the proposal distribution f_1 , and accepts this sample with probability

$$A_1(s, \tilde{s}_1) = \min \left(1, \frac{f_\theta(\tilde{s}_1) f_1(s | \tilde{s}_1)}{f_\theta(s) f_1(\tilde{s}_1 | s)} \right). \quad (41)$$

The probabilistic acceptance or rejection of the sample is practically implemented by drawing from a uniform distribution (denoted \mathcal{U}). Computing this acceptance probability requires an evaluation of the target distribution. To ensure compatibility with a pseudomarginal estimate of the target, it is necessary to store the computed realization [22], such that only one realization of the distribution is used for each sample, accepted or proposed. This is of little practical consequence, as it is desirable to store the evaluation regardless in order to avoid computational overhead. For clarity of what constitutes an evaluation of the target distribution, we retain the notation of Eq. (41).

In ordinary MH-MCMC, if the proposed sample was rejected, the sample at iteration n would be a duplicate of that at $n - 1$. With a DR scheme, however, when the first proposal is rejected, instead of duplicating the previous sample a second proposal \tilde{s}_2 is instead made from an alternative proposal f_2 , which is accepted with probability

$$A_2(s, \tilde{s}_1, \tilde{s}_2) = \min \left(1, \frac{f_\theta(\tilde{s}_2) f_1(\tilde{s}_1 | \tilde{s}_2) f_2(s | \tilde{s}_2, \tilde{s}_1) (1 - A_1(\tilde{s}_2, \tilde{s}_1))}{f_\theta(s) f_1(\tilde{s}_1 | s) f_2(\tilde{s}_2 | \tilde{s}_1, s) (1 - A_1(s, \tilde{s}_1))} \right). \quad (42)$$

This scheme can be extended to arbitrarily many proposals; if the final proposal is rejected, the sample is a duplicate of the previous as normal [23]. We implement here only up to a secondary proposal.

We utilize for these proposals an AM scheme on a Gaussian random walk. That is,

$$f_1(\tilde{s}_1 | s^{(n-1)}) = \mathcal{N}(s^{(n-1)}, S_{n-1}), \quad (43)$$

where S_{n-1} is a covariance that varies with the iteration number,

$$S_{n-1} = \begin{cases} S_0, & \text{if } n < n_0 \\ s_d C_{n-1} + s_d \xi_{AM} I, & \text{otherwise} \end{cases}, \quad s_d = \frac{2.4^2}{D}, \quad C_{n-1} = \text{Cov}(s^{(0)}, \dots, s^{(n-1)}). \quad (44)$$




Fig. 14 Posterior predictive distribution at the positions of negative polarity data for the approximately pseudomarginal inference with $N_r = 100$ and the reservoir pore radius taken to be $r_{p,r} = 8 \times 10^{-6}$; in blue are the data of [3]; the solid black is the median prediction, while the black dashed lines contain 90% of all predictions

The proposal is a normal distribution centered at the location of the current sample. In this way, the sampler walks randomly across parameter space. s_d is an approximately optimal scale factor based on the dimensionality of the samples D , \mathcal{I} is the identity matrix, and ξ_{AM} is a small factor to ensure the empirical covariances C are well-conditioned. The nature of the AM modification is that the covariance of the distribution is not fixed, but rather is adapted over time based on previously-accepted samples, beginning after a "burn-in" period $n < n_0$. Before this period is complete, the sampler uses some initialized covariance S_0 instead. Effectively, the algorithm maintains a separate proposal \tilde{f}_1 that it constantly updates, and once it has finished its burn-in period it switches to using the adaptive proposal.

The secondary proposal is simply the primary proposal with a smaller covariance,

$$f_2(\tilde{s}_2 | s_1, s) = f_2(\tilde{s}_2 | s) = \mathcal{N}(s, \gamma_{DR} S_{k-1}), \quad \gamma_{DR} < 1. \quad (45)$$

For the Gaussian random walk, this means the secondary proposal has a smaller step size, such that the is still likely to move around the domain little by little even if the comparatively large steps of the initial proposal, which could more rapidly traverse parameter space, are rejected.

Here, we initialize S_0 by fitting a normal distribution to the curvature of the posterior at the location of maximum likelihood (determined by using a canned minimizer). In this work, we use SciPy's `scipy.optimize.minimizer` Nelder-Mead solver to find a local maximum of the posterior to serve as the initial sample, and then initialize the covariance by a Laplace approximation of the posterior. It was necessary to add a small diagonal matrix to the corresponding initial covariance estimate to ensure it was well-conditioned. For all samples from the posterior, we drew 10^5 samples, chose an AM burn-in of $n_0 = 1000$, and chose a secondary proposal $\gamma_{DR} = 1 \times 10^{-2}$. For the AM proposal, $\xi_{AM} = 1 \times 10^{-20}$ was sufficient for all cases.

References

- [1] Lozano, P. C., Wardle, B. L., Moloney, P., and Rawal, S., "Nanoengineered thrusters for the next giant leap in space exploration," *MRS Bulletin*, Vol. 40, No. 10, 2015, p. 842–849.
- [2] Ziemer, J., Marrese-Reading, C., Dunn, C., Romero-Wolf, A., Cutler, C., Javidnia, S., Le, T., Li, I., Franklin, G., Barela, P., Hsu, O., Maghami, P., O'Donnell, J., Slutsky, J., Thorpe, J. I., Demmons, N., and Hruby, V., "Colloid Microthruster Flight Performance Results from Space Technology 7 Disturbance Reduction System," *International Electric Propulsion Conference*, 2017, p. 578.

Algorithm 1: PM-DRAM

Input: target f_θ , initial primary proposal f_1 , initial secondary proposal f_2 , number of samples N_f , burn-in n_0
Output: samples $[s^{(0)}, \dots, s^{(N_f)}]$
initialize $s^{(0)}$;
initialize AM proposals \tilde{f}_1, \tilde{f}_2 ;
 $z^{(0)} = f_\theta(s^{(0)})$;
for $n = 1$ **to** N_f **do**
 sample $\tilde{s}_1 \sim f_1(\tilde{s}_1 \mid s^{(n-1)})$;
 $\tilde{z}_1 = f_\theta(\tilde{s}_1)$;
 compute $A_1(s^{(n-1)}, \tilde{s}_1) = \min\left(1, \frac{\tilde{z}_1}{z^{(n-1)}} \frac{f_1(s^{(n-1)} \mid \tilde{s}_1)}{f_1(\tilde{s}_1 \mid s^{(n-1)})}\right)$;
 sample $u_1 \sim \mathcal{U}[0, 1)$;
 if $u_1 < A_1$;
 then
 $s^{(n)} = \tilde{s}_1$;
 $z^{(n)} = \tilde{z}_1$;
 else
 sample $\tilde{s}_2 \sim f_2(\tilde{s}_2 \mid \tilde{s}_1, s^{(n-1)})$;
 $\tilde{z}_2 = f_\theta(\tilde{s}_2)$;
 compute $A_2(s^{(n-1)}, \tilde{s}_1, \tilde{s}_2) = \min\left(1, \frac{\tilde{z}_2}{z^{(n-1)}} \frac{f_1(\tilde{s}_1 \mid \tilde{s}_2)}{f_1(\tilde{s}_1 \mid s^{(n-1)})} \frac{f_2(s^{(n-1)} \mid \tilde{s}_2, \tilde{s}_1)}{f_2(\tilde{s}_2 \mid \tilde{s}_1, s^{(n-1)})} \frac{1 - A_1(\tilde{s}_2, \tilde{s}_1)}{1 - A_1(s^{(n-1)}, \tilde{s}_1)}\right)$;
 sample $u_2 \sim \mathcal{U}[0, 1)$;
 if $u_2 < A_2$ **then**
 $s^{(n)} = \tilde{s}_2$;
 else
 $s^{(n)} = s^{(n-1)}$;
 end
 end
 end
 update \tilde{f}_1, \tilde{f}_2 with $s^{(n)}$;
 if $n = n_0$ **then**
 set $f_1 = \tilde{f}_1, f_2 = \tilde{f}_2$;
 end
end

- [3] Natisin, M. R., Zamora, H. L., McGehee, W. A., Arnold, N. I., Holley, Z. A., Holmes, M. R., and Eckhardt, D., “Fabrication and characterization of a full conventionally machined high-performance porous-media electrospray thruster,” *Journal of Micromechanics and Microengineering*, Vol. 30, 2020, p. 115021.
- [4] Legge, R. S., and Lozano, P. C., “Electrospray Propulsion Based on Emitters Microfabricated in Porous Metals,” *Journal of Propulsion and Power*, Vol. 27, No. 2, 2011, pp. 485–495.
- [5] Krejci, D., Mier-Hicks, F., Fucetola, C., Lozano, P., Schouten, A. H., and Martel, F., “Design and Characterization of a Scalable Ion Electrospray Propulsion System,” *International Electric Propulsion Conference*, 2015, p. 149.
- [6] Marrese-Reading, C. M., Anderson, J., Jung-Kubink, C., Polk, J., Singh, V., Yee, K., White, V., Wilson, D., Bruneau, P., Rouhi, N., Greer, F., Borgonia, J.-P., Dickie, M., Muller, R., Swindlehurst, R., Niblett, N., Ziemer, J., Gray, A., and Mueller, J., “Electrospray Thruster Performance with Microfabricated Emitter Arrays and Indium Propellant,” *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, 2016, p. 4738.
- [7] Perez-Martinez, C., and Lozano, P., “Ion field-evaporation from ionic liquids infusing carbon xerogel microtips,” *Applied Physics Letters*, Vol. 107, No. 4, 2015, p. 043501.
- [8] Petro, E., Bruno, A., Lozano, P., Perna, L. E., and Freeman, D., “Characterization of the TILE Electrospray Emitters,” *AIAA Propulsion and Energy 2020 Forum*, 2020, p. 3612.

- [9] Schönherr, T., Little, B., Krejci, D., Reissner, A., and Seifert, B., “Development, Production, and Testing of the IFM FEEP Thruster,” *International Electric Propulsion Conference*, 2019, p. 362.
- [10] Jorns, B. A., Gorodetsky, A., Lasky, I., Kimber, A., and Dahl, P., “Uncertainty Quantification of Electro spray Thruster Array Lifetime,” *International Electric Propulsion Conference*, 2019, p. 317.
- [11] Gorodetsky, A., Whittaker, C. B., Szulman, A., and Jorns, B., “Robust Design of Electro spray Emitters,” *AIAA Propulsion and Energy 2021 Forum*, 2021, p. 3422.
- [12] Whittaker, C. B., Gorodetsky, A. A., and Jorns, B. A., “Quantifying Uncertainty in the Scaling Laws of Porous Electro spray Emitters,” *AIAA Propulsion and Energy Forum*, 2020, p. 3615.
- [13] Coffman, C., Martínez-Sánchez, M., Higuera, F., and Lozano, P. C., “Structure of the menisci of leaky dielectric liquids during electrically-assisted evaporation of ions,” *Applied Physics Letters*, Vol. 109, No. 23, 2016, p. 231602.
- [14] Courtney, D. G., and Shea, H., “Influences of porous reservoir Laplace pressure on emissions from passively fed ionic liquid electro spray sources,” *Applied Physics Letters*, Vol. 107, 2015, p. 103504.
- [15] Coffman, C. S., Martínez-Sánchez, M., and Lozano, P. C., “Electrohydrodynamics of an ionic liquid meniscus during evaporation of ions in a regime of high electric field,” *Physical Review E*, Vol. 99, 2019, p. 063108.
- [16] Courtney, D. G., “Ionic Liquid Ion Source Emitter Arrays Fabricated on Bulk Porous Substrates for Spacecraft Propulsion,” Ph.D. thesis, Massachusetts Institute of Technology, 2011.
- [17] Glover, P. W., and Walker, E., “Grain-size to effective pore-size transformation derived from electrokinetic theory,” *Geophysics*, Vol. 70, No. 1, 2009, pp. E17–E29.
- [18] St.Peter, B., Dressler, R. A., Chiu, Y.-H., and Fedkiw, T., “Electro spray Propulsion Engineering Toolkit (ESPET),” *Aerospace*, Vol. 7, No. 91, 2020.
- [19] Martínez-Sánchez, M., “Lecture 23-25: Colloidal Engines,” Web, 2004. URL https://dspace.mit.edu/bitstream/handle/1721.1/100855/16-522-spring-2004/contents/lecture-notes/lecture23_25.pdf, lecture notes.
- [20] Perez-Martinez, C. S., “Engineering Ionic Liquid Sources for Ion Beam Applications,” Ph.D. thesis, Massachusetts Institute of Technology, 2016.
- [21] Wright, P. L., and Wirz, R. E., “Multiplexed electro spray emission on a porous wedge,” *Physics of Fluids*, Vol. 33, 2021, p. 012003.
- [22] Andrieu, C., and Roberts, G. O., “The pseudo-marginal approach for efficient Monte Carlo computations,” *The Annals of Statistics*, Vol. 37, No. 2, 2009.
- [23] Mira, A., “On Metropolis-Hastings algorithms with delayed rejection,” *Metron*, Vol. 59, No. 3, 2001.