

Emitter Variability Inferred by Training Single-Emitter Models on Electrospray Array Tests

IEPC-2025-564

*Presented at the 39th International Electric Propulsion Conference, Imperial College London, London, United Kingdom
14-19 September 2025*

Collin B. Whittaker* and Benjamin A. Jorns†

University of Michigan, Ann Arbor, MI, 48104, United States

A study inferring variability among individual emitters in an electrospray array thruster is presented. A model to predict the emission current of a single porous conical type electrospray emitter is first rederived. Experiments on a 6102-emitter porous electrospray array thruster (the MEAT-1.2) are then reviewed, with the emission current sourced by the thruster varying from 22 μA at 650 V to 5.6 mA at 2000 V. Surface profiles are also resolved for 143 emitters in the array, with the conical emitters having tip radii approximately 25 μm , a body half angle around 0.3 rad, and a basal radius of 145 μm , though it is further shown that there are distinct populations of emitters that can be tied to different cutting tools used in the manufacture of the array. Subsequently, a Bayesian inference problem to learn the *ad hoc* scaling coefficients introduced by the model from these emission data and informed by emitter metrology is formulated and executed. The model is able to reconstruct the training data set within 90% confidence, and predictions of individual emitter behavior indicate that emission is highly varied across the array, producing the superlinear dependence seen in the macroscopic array current. A sensitivity study quantifies that below 1300 V, the dominant source of uncertainty in the predictions is emitter geometry (because few emitters are active), but above 1300 V uncertainty in propellant temperature dominates (because emission is highly sensitive to propellant viscosity). Finally, these uncertainty sources are discussed in the context of designing robustly performant systems.

* Assistant Research Scientist, Aerospace Engineering, cbwhitt@umich.edu

† Assistant Professor, Aerospace Engineering



I. Introduction

Electrospray array thrusters exhibit emergent phenomena with key consequences for system design and performance. The fundamental unit of these devices is the emitter, a needle-like electrode supplied with conducting liquid. When energized, the emitter concentrates electric fields at its apex, to the point that an energetic beam of charged particles is extracted from the fluid.^{1–4} The reaction force produced by the beam is used as a means of propulsion in space.^{5,6} To this end, devices operating on ionic liquid propellants and extracting primarily ionic species (ionic liquid ion sources, ILIS) have seen extensive development because their working fluid can be stored as a nonvolatile liquid at room temperature and is only tenuously conducting, such that it avoids space-charge limits inherent to similar liquid metal ion sources.^{7,8} These traits—among others—make electrosprays a potentially enabling technology for small-scale spacecraft, which operate at low powers where more mature, plasma-based electric rockets like Hall thrusters suffer major performance degradation.⁹

However, since individual electrospray emitters tend to produce less than a micronewton of force, even at small scale it is necessary to aggregate many such emitters in an array to form a useful thruster.^{5,10–12} These arrays are subject to variability between individual emitters—principally resulting from variations in geometry^{11,13,14} but also stemming from nonuniform fluidic boundary conditions¹⁵ or electrode misalignment¹⁶—and this implies the behavior of an array is not simply a scale multiple of that of a single emitter. Indeed, new behaviors can emerge as a result of this nonuniformity, for example that the total current emitted by an array tends to be a positively inflected function of voltage^{5,11,17} while individual emitters tend to have a linear current response.^{18–20}

A more pernicious emergent phenomenon is a decrease in expected device lifetime. The primary failure mechanism in electrospray is an electrical short between the emitter and its extraction electrode,^{5,6,21,22} which prevents the device from standing off voltage sufficient to induce emission. Since the emitters electrically tied together, the short of any one emitter shorts all emitters. As such, the mean time to failure of the array decreases as the number of emitters in it increases,^{21,23} a consequence of an increased probability that a marginal emitter shorts prematurely compared to the bulk population.

Predicting how an array’s performance arises from—and varies subject to—uncertainty in its constituent emitters is thus a fundamental task in electrospray design. Experimental efforts have characterized emission nonuniformity at system scale,^{15,16,24,25} but achieving sufficient signal to noise to resolve individual emitter beams remains a challenge,²⁶ especially considering the degree to which probes must perturb the device under test.^{16,27} This leaves modeling and simulation as recourse to bridge the gap between what is observable at the emitter scale (e.g., variance in emitter geometry) and what is observable at the array scale (e.g., emission current). However, electrosprays are fundamentally multiscale devices,^{19,28–30} such that modeling each of the hundreds or thousands of emitters in an array to capture the physics at high fidelity—while also accounting for uncertainty in emitter geometry and other parameters—can become impracticable.

To that end, in this work we instead leverage semi-empirical models trained on experiments of a porous electrospray array thruster to infer the behavior of individual emitters in the array. That is, we first model the emission of an individual emitter (Sec. II.A) such that we can predict array behavior as the sum over many emitters;. However, our model contains scaling parameters that must be calibrated from the data itself (Sec. II.B). As such, we formulate an inference (Sec. III) over these missing parameters accounting rigorously for uncertainty in the model inputs and data, including to infer manufacturing tolerances for the emitters from individual emitter measurements (Sec. III.D). After training the model in this way (Sec. IV), we propagate uncertainty to make probabilistic predictions of performance over the array (Sec. V). From these results, we then discuss the inferred behavior of individual emitters (Sec. V.A), which sources of uncertainty dominate the analysis (Sec. V.B), and what implications our results have for designing systems that are robust to these sources of uncertainty (Sec. V.C). We end by collecting and summarizing our conclusions (Sec. VI).

II. Model and Experimental Data

Our methodology for understanding emitter variability relies on modeling the behavior of individual emitters and then computing the composition over all emitters to predict the behavior of an array, such that we can calibrate it against experimental data. We focus our analysis on porous ionic liquid ion sources (ILIS), which have seen widespread development motivated by their comparative ease of manufacture^{10,31,32} and



ability to facilitate higher current densities.^{1,11,33} To model the emission of individual emitters within the array, we utilize the model we previously detailed in Ref. 34, summarized here (Sec. II.A). As a calibration data set, we consider experiments (Sec. II.B) to measure the emission current as a function of voltage for the MEAT-1.2,³⁵ a thruster consisting of over 6000 individual porous electrospray emitters and constructed as a testbed for exploring challenges in scaling electrosprays to higher power systems.

A. Single Emitter Emission Model

High fidelity electrohydrodynamic models have examined the physics underlying ILIS operation and motivated scaling laws governing both the onset of Taylor cones and the current these cones emit.^{8,19,28,30,36} However, these analyses often consider only a single, archetypal fluid meniscus of defined size and location. In practice, not fully captured in this analysis is that porous electrospray emitters host many menisci where Taylor cones could form, and they are observed to develop multiple active emission sites on a single emitter.^{15,18,27} Our model abstracts the formation of multiple emission sites by supposing that the emitter surface is populated by many menisci and that these scaling laws dictate the onset and behavior of individual emission sites on a meniscus by meniscus basis. The aggregate emitter behavior is then the sum over all menisci.

Particularly, we model the emitter as a spherically capped cone (see Fig. 1 for key dimensions or Ref. 37). We suppose that the surface of the emitter is populated by menisci of varying diameter p , with the distribution




Figure 1. Diagram of spherically capped cone geometry

over meniscus size encapsulated by a probability density function of phenomenological form

$$f(p) = \begin{cases} \frac{1-m}{p_2^{1-m} - p_1^{1-m}} \frac{1}{p^m}, & p \in [p_1, p_2], \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

This equation indicates that the menisci are distributed inversely proportional to diameter by some characteristic exponent m , between a minimum meniscus size p_1 and maximum meniscus size p_2 . The number density of menisci on the surface is then $n_p = \frac{\psi}{A_p}$, where the surface wetting coefficient, ψ , is a normalizing

factor indicating the proportion of the emitter's surface area that is covered by the menisci. \bar{A}_p is the mean meniscus area, $\bar{A}_p = \int \frac{\pi}{4} p^2 f(p) dp$. We further suppose that the maximum meniscus size is limited in the following way:

$$p_2 = \min(\bar{p}_2, 2R_c \cos \alpha), \quad (2)$$

where p_2 is the maximum meniscus size supported by an emitter, \bar{p}_2 is the maximum meniscus size supported by the substrate, R_c is the tip radius of the cone, and α is the half-angle of the cone body. Eq. (2) thus communicates that while the substrate might be capable of supporting menisci as large as \bar{p}_2 , the emitter cannot support menisci larger than the diameter of its tip, due to a lack of space—e.g., a meniscus of diameter 10 μm could not form on an emitter of diameter only 1 μm .

The first scaling law determining individual emission site behavior over this population is an onset criterion, 1_O , which dictates when a meniscus develops into a Taylor cone. It takes the form of an indicator function,

$$1_O = \begin{cases} 1, & E(s) \geq E_0 = \beta_0 \sqrt{\frac{8\gamma}{\varepsilon_0 p}} - \beta_1 \frac{P_r}{\sqrt{\frac{2\gamma\varepsilon_0}{p}}}, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

This relation captures that onset occurs when the electric pressure at the meniscus exceeds the restorative forces provided by capillary pressure and the hydrostatic pressure in the reservoir feeding the emitter. ε_0 is the permittivity of vacuum, γ the surface tension of the working fluid, P_r the reservoir pressure, $E(s)$ the electric field at the emitter surface, and E_0 the corresponding onset field strength. Crucially, this electric field $E = E(s)$ is treated as position-dependent, captured by a parameter of the surface, s . The quantities β_0 and β_1 are constants of order unity that dictate the relative strength of the capillary pressure and reservoir pressure effects.

The second scaling law is a single site emission model predicting the current emitted by an individual meniscus, i . After the treatment of Refs. 8,19,30, we express this as

$$i = \frac{\xi\rho}{r_h} \left[\zeta_0 \frac{4\gamma}{p} + \zeta_1 \sqrt{\frac{2\gamma\varepsilon_0}{p}} (E(s) - E_0) \right], \quad (4)$$

where ξ is the charge to mass ratio of the emitted species, ρ the density of the fluid, and r_h the hydraulic impedance to an individual meniscus from the reservoir. This expression represents a linearization of the current emitted by a meniscus by some characteristic slope, ζ_1 , and offset, ζ_0 , within a dimensionless space.¹⁹ The individual site impedance we take as

$$r_h = N_s \frac{\mu}{2\pi\kappa} \frac{1}{1 - \cos \alpha} \left[\frac{\tan \alpha}{R_c} - \frac{\sin \alpha}{r_b} \right], \quad (5)$$

$$N_s = \iint 1_O f(p) dp n_p \frac{dA}{ds} ds, \quad (6)$$

with N_s the total number of onset emission sites, μ the dynamic viscosity of the fluid, κ the permeability of the porous substrate, and r_b the base radius of the emitter (see Fig. 1). Eq. (5) expresses that the hydraulic impedance to each emission site is equal, and that their parallel sum must equal the total hydraulic impedance of the emitter, the form of which follows that of Courtney.³⁸

The total current emitted I takes the form of an integral over all menisci in the population and over all area of the emitter,

$$I = \iint i 1_O f(p) dp n_p \frac{dA}{ds} ds, \quad (7)$$

where $\frac{dA}{ds}$ is the parameterization of the surface area by s ,

$$\frac{dA}{ds} = \begin{cases} 2\pi R_c \sin(s/R_c), & s \in [0, s_c], \\ 2\pi [R_c \cos \alpha + (s - s_c) \sin \alpha], & s \in [s_c, s_c + s_b], \end{cases} \quad (8)$$



$$s_c = \left(\frac{\pi}{2} - \alpha \right) R_c, \quad (9)$$

$$s_b = r_b \csc \alpha - R_c \cot \alpha. \quad (10)$$

Here, s_c is the arc length from the apex of the emitter to the edge of the cap and s_b is the arc length along the cone body (see Fig. 1).

Practically for $E(s)$, since the field strength scales linearly with the voltage of the emitter, V , we express $E(s) = V \hat{E}(s)$, where $\hat{E}(s)$ is a reference field strength at unity emitter voltage. $\hat{E}(s)$ is itself only a function of the emitter and extractor geometry,

$$\hat{E}(s) = \hat{E}(s; R_c, \alpha, r_b, R_A, t_e, \Delta). \quad (11)$$

Here, R_A denotes the radius of the extractor aperture, t_e its thickness, and Δ the displacement of its top surface from the basal plane of the emitter. Modeling the dependence of Eq. (11) requires additional analysis, which we treat in Sec. III.F.

Equation (7) is a composition over several submodels and so is a function of many parameters. For our purposes, these can be organized into classes: the scaling parameters, $\mathbf{Pm} = \{\beta_0, \beta_1, \zeta_0, \zeta_1\}$, which are coefficients resulting from linearization of higher order physics; the surface properties, $\mathbf{Sf} = \{\psi, p_1, \bar{p}_2, m\}$, which describe the meniscus population on the surface; the fluidics properties, $\mathbf{Fl} = \{\kappa, P_r, \gamma, \rho, \mu\}$, which dictate the flow of propellant in the system; the beam property, $\mathbf{Be} = \{\xi\}$, which converts the mass flow rate of propellant being drawn through the emitter into the charge flow rate being sprayed from it; and the emitter and extractor geometries, $\mathbf{Em} = \{R_c, \alpha, r_b\}$ and $\mathbf{Ex} = \{R_A, t_e, \Delta\}$, which modulate the hydraulic impedance and the field strength on the emitter.

Subject to this classification, the predicted current of an individual emitter is the function

$$I = I(V; \mathbf{Pm}, \mathbf{Sf}, \mathbf{Fl}, \mathbf{Be}, \mathbf{Em}, \mathbf{Ex}). \quad (12)$$

This model is the basic building block for predicting the current from an entire array of emitters, like that described in the following section.

B. Training Data

We focus our study on the Michigan Electrospray Array Thruster Series 1 Version 2 (MEAT-1.2) family of systems.^{35,37} The MEAT-1.2—shown mounted for experiments in Fig. 2—is a 10 W power class thruster based on a porous conical emitter architecture. The emitter chip consists of a single piece of P5 (1 μm pores) borosilicate glass frit 70 mm in diameter, from which the emitters are subtractively machined on a CNC mill using a square tapered end mill, based on earlier methods of Natsin *et al.*¹⁰ The end mill makes cuts at full depth to remove material between emitters and takes a circular path around each emitter to leave a conical geometry behind.

The emitter chip is coupled to a P4 (10 μm pores) borosilicate glass frit that acts as a propellant reservoir, and both are retained within an isolated propellant module. The extractor chip is a sheet of MACOR ceramic metallized on one side with a silver film and into which are machined circular apertures to match the emitters of the emitter chip (Fig. 1). This thin film architecture is designed to provide resilience to shorts between the electrodes.³⁹ The extractor chip is then bonded into a steel support frame to provide rigidity. Finally, the propellant module and extractor module are mounted within an aluminum housing, which includes set screws to adjust their relative alignment.

The MEAT-1.2 system is constructed such that the emitter and extractor chips can vary from thruster to thruster, allowing it to serve as a design testbed; the chips used in Ref. 35 were serial number 202. The emitter chip consisted of 6102 emitters each of a nominal height 455 μm and hexagonally tiled with a pitch of 660 μm . The extractor chip had a thickness of 460 μm and an aperture diameter of 460 μm .

The ionic liquid propellant 1-ethyl-3-methylimidazolium bis(trifluoromethylsulfonyl)imide (EMIm-Im) was loaded into the thruster and the thruster discharged in vacuum for the experiments. The emission current of the thruster was mapped in negative polarity, spanning from only a few nA at 350 V up to 5.6 mA at 2000 V, with corresponding powers ranging from ~ 1 μW to 11.2 W. The very smallest powers were inferred from probes in the plume and did not necessarily measure the total emission current, so for our training data we take only readings obtained directly from the emitter power supply, which we collect in Tab. 1. We have included a null datum at 325 V to indicate that no current was measured at or below this voltage (including






Figure 2. Photograph of the Michigan Electro spray Array Thruster Series 1 Version 2 #202 (MEAT-1.2-202) prepared for testing

Voltage [V]	Current [mA]	Voltage [V]	Current [mA]
325	0	1250	1.22
650	0.0216	1300	1.35
700	0.0391	1350	1.53
750	0.0610	1400	1.76
800	0.100	1450	2.00
850	0.158	1500	2.25
900	0.234	1550	2.51
950	0.333	1600	2.73
1000	0.449	1700	3.31
1050	0.590	1800	3.89
1100	0.755	1900	4.67
1150	0.925	2000	5.59
1200	1.09		

Table 1. Emission current as a function of operating voltage training data derived from Ref. 35

by the probes). These data encode the current response of individual emitters, and so we leverage this I-V map as a powerful set of training data for our current model, Eq. (7). However, since the array current is the composition over individual emitter currents and application of the model requires knowledge of many incidental parameters, we dedicate the following section to discussing our inference methodology.

III. Model Inference

We train the model of Sec. II.A on the data of Sec. II.B. The model is fully defined by the emitter voltage V and the set of parameters $\theta = \mathbf{Pm} \cup \mathbf{Sf} \cup \mathbf{Fl} \cup \mathbf{Be} \cup \mathbf{Em} \cup \mathbf{Ex}$, so our state of knowledge in those parameters describes our state of knowledge in the model. Within a Bayesian formulation, this constitutes a parameter inference problem given our observation of the data, which—by application of Bayes’ rule—we express as

$$P(\theta | D) \propto P(D | \theta) P(\theta), \quad (13)$$

where $P(\cdot)$ denotes probability and D are the data. In this formalism, our state of knowledge in the parameters is described by a probability distribution function over values of θ , such that our degree of belief in any particular value of θ given the data, the posterior $P(\theta | D)$, is proportional to the product of the likelihood $P(D | \theta)$, the probability of observing the data given that value of the parameters, and the prior $P(\theta)$, our degree of belief in that value prior to incorporating knowledge of the data. The constant of proportionality is a normalizing factor that is important in other analyses but can be discarded here.

For our purposes, we must extend our model to the array context. We make a notational shift such that I now denotes the total array current and is given as the sum over all N_e individual emitters,

$$I = \sum_{j=1}^{N_e} I_j, \quad (14)$$

where the current from each emitter is then defined by its own set of parameters. That is, $I_j = I_j(V_j; \theta_j)$, where $\theta_j = \mathbf{Pm}_j \cup \mathbf{Sf}_j \cup \mathbf{Fl}_j \cup \mathbf{Be}_j \cup \mathbf{Em}_j \cup \mathbf{Ex}_j$. We assume that the emitter voltage, scaling parameters, surface properties, fluidic properties, and beam property are common among all emitters ($V_j = V$, $\mathbf{Pm}_j = \mathbf{Pm}$, $\mathbf{Sf}_j = \mathbf{Sf}$, etc.), but that the emitter and extractor geometry varies on an emitter by emitter basis. The corresponding set of parameters defining the array current model, $I = I(V; \theta)$, is

$$\theta = \bigcup_{j=1}^{N_e} \theta_j = \mathbf{Pm} \cup \mathbf{Sf} \cup \mathbf{Fl} \cup \mathbf{Be} \cup \left[\bigcup_{j=1}^{N_e} \mathbf{Em}_j \cup \mathbf{Ex}_j \right]. \quad (15)$$

Now that we have identified the set of parameters defining the array current model, we must define our corresponding likelihood and prior.

A. Likelihood

Within our Bayesian inference framework, the likelihood $P(D | \theta)$ is the means by which we infuse our experimental observations into the inference, taking the form of a probabilistic model of those observations. In our case, from the experiments we have N_d pairs of voltage-current data, $D^{(k)} = \{V^{(k)}, I^{(k)}\}$, which are measurements of the emitter voltage and emitter current, forming

$$D = \bigcup_{k=1}^{N_d} D^{(k)}. \quad (16)$$

The superscript in parentheses notation distinguishes between a quantity and a measurement of that quantity (which may differ as a result of error in the measurement) and will be used throughout the remainder of our analysis. We neglect error in the voltage measurements (often less than 0.1%), but we model that the current measurements are distributed about the model predictions with some Gaussian noise, $\eta_I^{(k)}$,

$$I^{(k)} = I(V^{(k)}; \theta) + \eta_I^{(k)}, \quad \eta_I^{(k)} \sim \mathcal{N}\left(0, \left(\sigma^{(k)}\right)^2\right), \quad (17)$$



where $\sigma^{(k)}$ is the standard deviation of the noise for datum k and $\mathcal{N}(a, b^2)$ denotes the normal distribution of mean a and variance b^2 .

Ideally, the $\sigma^{(k)}$'s are derived directly from specifications like the instrument resolution. However, the error in these data also stems from uncertain mechanisms in the emission behavior (unsteady emission, hysteresis, etc.).³⁵ Given this uncertainty in the precision of the measurements, we adopt an abstract model for the noise of the form

$$\sigma^{(k)} = \sigma_0 + \sigma_1 I^{(k)}. \quad (18)$$

That is, the error in a current datum is the combination of some noise floor, σ_0 , and an error proportional to the magnitude of that datum, σ_1 . The constant term models the finite resolution of the instrumentation, while the proportional term models the variability/noise inherent to the source—which scale with the amount of current it sources. The one exception to Eq. (18) we make for the null current datum of Tab. 1, $V^{(1)} = 325$ V and $I^{(1)} = 0$ μ A, which from the instrumentation is known with much higher precision, allowing us to assign instead simply $\sigma^{(1)} = 1$ nA.

The unknown noise terms σ_0 and σ_1 constitute a set of hyperparameters, $\theta' = \{\sigma_0, \sigma_1\}$. While we may have some prior intuition for these noise terms, we suspect the data will be much more informative, and so learn them alongside the model parameters,

$$P(\theta, \theta' | D) \propto P(D | \theta, \theta') P(\theta) P(\theta'). \quad (19)$$

Such that, assuming the error in each datum is independent, we achieve the final functional form of the likelihood,

$$P(D | \theta, \theta') = \prod_{k=1}^{N_d} \frac{1}{\sigma^{(k)} \sqrt{2\pi}} \exp \left[-\frac{(I(V^{(k)}; \theta) - I^{(k)})^2}{2(\sigma^{(k)})^2} \right]. \quad (20)$$

B. Prior

The prior distribution $P(\theta)$ encodes any assumptions or other relevant information about the parameters aside from the data that we wish to incorporate to our analysis. In our case, we utilize many different sources of information to motivate the various components of the model, and we summarize this prior prescription in Tab. 2. Mathematically, this represents a factorization of the prior as

$$\begin{aligned} P(\theta') P(\theta) &= P(\sigma_0) P(\sigma_1) P(\zeta_0) P(\zeta_1) P(\beta_0, \beta_1 | D_\beta) \\ &\times P(\psi) P(p_1) P(\bar{p}_2) P(m) \\ &\times P(\kappa) P(P_r | \gamma) P(\gamma, \rho, \mu | T) P(T) P(\xi) \\ &\times P(\cup_j \mathbf{E} \mathbf{m}_j | D_{\mathbf{E} \mathbf{m}}) \prod_j P(R_{A_j}) P(t_{e_j}) \\ &\times P(\cup_j \Delta_j | D_\Delta), \end{aligned} \quad (21)$$

which we must describe piece by piece.

For the noise parameters, we suppose that the constant noise term has a magnitude around 1 μ A—motivated by the capability of the power supply's current monitor—but ascribe 95% confidence to it being within one order of magnitude larger or smaller. For the proportional term, we model that all magnitudes are equally probable *a priori*, since we have comparatively little information about this scaling. We similarly ascribe uniform priors over the magnitude of the current scaling parameters ζ_0 and ζ_1 . We condition β_0 and β_1 on additional data, D_β , and so learn them hierarchically (Sec. III.C).

Following the discussion of Ref. 34, we note that to first order the emission current is not a function of ψ , under our approximation that the hydraulic impedence to each emission site is equal and our assumption that the charge to mass ratio is invariant between emission sites. Consequently, we choose a fixed, arbitrary value in conducting the analysis (the porosity of the substrate). We include it as a parameter in our treatment because it is relevant to other potential observables (i.e., the number of emission sites), and it correspondingly appears as an input to our implementation of the model.⁴⁰ The minimum meniscus size we tie to the minimum pore size in the substrate, 1 μ m, with a nominal error of $\pm 10\%$. The magnitude of the maximum meniscus size we motivate loosely from Ref. 34, ascribing a median around 10 μ m with



Parameter	Prior Distribution
Noise parameters θ'	
σ_0	$\ln \sigma_0 \sim \mathcal{N}\left(\ln 1 \text{ }\mu\text{A}, \left(\frac{\ln 10}{1.96}\right)^2\right)$
σ_1	$\ln \sigma_1 \sim \mathcal{U}(-\infty, \infty)$
Scaling parameters, Pm	
ζ_0	$\ln \zeta_0 \sim \mathcal{U}(-\infty, \infty)$
ζ_1	$\ln \zeta_1 \sim \mathcal{U}(-\infty, \infty)$
β_0	nested learning, see Sec. III.C
β_1	nested learning, see Sec. III.C
Surface properties, Sf	
ψ	$\psi = 0.48$
p_1	$p_1 \sim \mathcal{N}\left(1 \text{ }\mu\text{m}, \left(\frac{0.1 \text{ }\mu\text{m}}{1.96}\right)^2\right)$
\bar{p}_2	$\ln \bar{p}_2 \sim \mathcal{N}\left(\ln 10 \text{ }\mu\text{m}, \left(\frac{\ln 2}{1.96}\right)^2\right)$
m	$m \sim \mathcal{U}(-\infty, \infty)$
Fluidic properties, F1	
κ	$\kappa \sim \mathcal{N}\left(0.15 \text{ }\mu\text{m}^{-2}, \left(\frac{0.015 \text{ }\mu\text{m}^{-2}}{1.96}\right)^2\right)$
P_r	$P_r = -\frac{4\gamma}{p_{max}}$
	$p_{max} \sim \mathcal{N}\left(16 \text{ }\mu\text{m}, \left(\frac{1.6 \text{ }\mu\text{m}}{1.96}\right)^2\right)$
γ	$\gamma = \gamma(T)$
ρ	$\rho = \rho(T)$
μ	$\mu = \mu(T)$
	$T \sim \mathcal{N}\left(291 \text{ K}, \left(\frac{2 \text{ K}}{1.96}\right)^2\right)$
Beam property, Be	
ξ	$\xi \sim \mathcal{N}\left(92.8 \frac{\text{C}}{\text{g}}, \left(\frac{4.0 \frac{\text{C}}{\text{g}}}{1.96}\right)^2\right)$
Emitter geometry, $\cup_j \mathbf{Em}_j$	
R_{ej}	nested learning, see Sec. III.D
α_j	nested learning, see Sec. III.D
r_{bj}	nested learning, see Sec. III.D
Extractor geometry, $\cup_j \mathbf{Ex}_j$	
R_{Aj}	$R_{Aj} = 230 \text{ }\mu\text{m}$
t_{ej}	$t_{ej} = -460 \text{ }\mu\text{m}$
Δ_j	nested learning, see Sec. III.E

Table 2. Prior distribution over parameters θ and hyperparameters θ' of the current model



error bounds a factor 2 in either direction. The meniscus exponent we assume can take any value with equal probability.

The substrate permeability we use as quoted from the manufacturer's datasheet and again ascribe a nominal $\pm 10\%$ uncertainty since none is given. The reservoir pressure is driven by capillarity for our passively fed thruster, so it is a function of the surface tension of the propellant (itself uncertain) and the maximum pore size in the reservoir substrate, $p_{max} = 16 \text{ } \mu\text{m} \pm 10\%$.

The surface tension, density, and viscosity of the propellant are functions of the propellant temperature T , which we compute using the Electrospray Propulsion Engineering Toolkit Propellant Database.^{41, 42} We assign a distribution over T based on other experiments showing that the temperature of apparatus within the test facility tends to be about 291 K, but varies within a couple degrees.

To bound the charge to mass ratio, we leverage additional experiments reported in Ref. 35, which directly measured the mass flow rate of propellant leaving the thruster using a microbalance. Specifically, we use the value and error bars computed at 1000 V, which was found to be representative of the spray across different voltages.

We base the distribution over the emitter geometries on measurements taken of the emitter chip, D_{Em} . However, it is infeasible to characterize all 6102 emitters on the chip in this way, so we instead measure a sample of the emitters and infer the distribution over emitter geometry (i.e., manufacturing tolerances) from these measurements, forming a nested inference scheme (Sec. III.D). For the extractor geometry, we treat the aperture radius, R_A , and extractor thickness, t_e , as deterministic and assign them their nominal values (230 μm and 460 μm , respectively). These quantities are both well controlled during the machining process (tolerances $\pm 5 \text{ } \mu\text{m}$ typical), and the electric field strength is not a strong function of either quantity at that scale. The extractor offset Δ , however, can vary across the array over a range of order 100 μm as a result of flexure in the extractor structure,³⁵ so we also infer it hierarchically given measurements D_Δ (Sec. III.E).

C. Onset Criterion Parameters

We have cogent prior information about the onset criterion scaling parameters β_0 and β_1 resulting from simulations done by Coffman *et al.*,⁸ who interrogated the dependency of Eq. (3) within the dimensionless space

$$\bar{E}_0 = E_0 \left(\frac{8\gamma}{\varepsilon_0 p} \right)^{-\frac{1}{2}}, \quad \bar{P}_r = P_r \left(\frac{4\gamma}{p} \right)^{-1}, \quad (22)$$

and corresponding onset criterion

$$\bar{E} \geq \bar{E}_0 = \beta_0 - \beta_1 \bar{P}_r. \quad (23)$$

They simulated the onset of individual, archetypal menisci as a function of \bar{P}_r to determine the resulting \hat{E}_0 , forming a set of $N_{d\beta}$ data

$$D_\beta = \bigcup_{k=1}^{N_{d\beta}} \left\{ \bar{E}_0^{(k)}, \bar{P}_r^{(k)} \right\}. \quad (24)$$

We must incorporate these previous results into the prior of Eq. (21), which is done by applying Bayes' rule hierarchically,

$$P(\beta_0, \beta_1 \mid D_\beta) \propto P(D_\beta \mid \beta_0, \beta_1) P(\beta_0, \beta_1). \quad (25)$$

That is, the prior over the β 's for our inference incorporating the data D is the posterior of an inference over the β 's incorporating data D_β , with $P(D_\beta \mid \beta_0, \beta_1)$ and $P(\beta_0, \beta_1) = P(\beta_0) P(\beta_1)$ being the hierarchical likelihood and prior, respectively.

For this nested inference, we again model a Gaussian likelihood where we assume that the $\bar{P}_r^{(k)}$'s are known exactly but that the output of the simulation is subject to some constant error of unknown magnitude, σ_β ,



relative to the model prediction of Eq. (23). We again learn this error hyperparametrically,

$$P(D_\beta | \beta_0, \beta_1) = \int P(D_\beta | \beta_0, \beta_1, \sigma_\beta) P(\sigma_\beta) d\sigma_\beta, \quad (26)$$

$$P(D_\beta | \beta_0, \beta_1, \sigma_\beta) = \prod_{k=1}^{N_{d,\beta}} \frac{1}{\sigma_\beta \sqrt{2\pi}} \exp \left[-\frac{\left(\bar{E}_0 \left(\bar{P}_r^{(k)}; \beta_0, \beta_1 \right) - \bar{E}_0^{(k)} \right)^2}{2(\sigma_\beta)^2} \right]. \quad (27)$$

Finally, we assign comparatively uninformative priors (Tab. 3) to communicate a lack of any additional prior information. This completes our treatment of the onset parameters.

Parameter	Prior Distribution
σ_β	$\ln \sigma_\beta \sim \mathcal{U}(-\infty, \infty)$
β_0	$\beta_0 \sim \mathcal{U}(-\infty, \infty)$
β_1	$\beta_1 \sim \mathcal{U}(-\infty, \infty)$

Table 3. Hierarchical prior distribution over onset criterion parameters

D. Emitter Geometry

The data D do not carry much information about the geometry of any individual emitter \mathbf{Em}_j , because the current emitted by each is anonymized in taking the sum of Eq. (14). As such, without additional prior information about the \mathbf{Em}_j 's, our analysis would be polluted by uncertainty in the geometry and may fail to conclude anything much useful about the parameters of interest, θ^* . Moreover, individual emitter dynamics are of chief interest in understanding emission nonuniformity and the emergent properties of the array derived therefrom. Therefore, we must ground our understanding in the emitter geometries through additional measurements. Since it is infeasible to measure every emitter in this way, we instead measure a subset of the emitters (Sec. III.D.1) and infer from them properties of the entire population (Sec. III.D.2).

1. Measurement

To measure individual emitter geometries, we apply the characterization methodology of Ref. 13, whereby surface profilometry is used to produce a topographic map of the emitter chip and the macroscopic parameters, \mathbf{Em}_j , are regressed on an emitter by emitter basis by fitting the emitter geometry of Fig. 1 to the data. Emitter chip 202 was inspected with a laser confocal microscope (Olympus OLS 4000 LEXT) after the experiments of Ref. 35. Before imaging, the chip was soaked in a bath of isopropanol and allowed to dry to remove most of the propellant from the chip and improve image quality (by making the emitters less translucent). For the objective selected, the microscope had a lateral resolution of 625 nm (parallel to the emitter basal plane) and a nominal vertical resolution of order 10 nm (orthogonal to the emitter basal plane). A single emitter takes approximately 2 minutes to image on the machine in this way, motivating our adoption of a sampling based strategy.

To select a representative subset of emitters to measure, we first divided the area of the emitter chip (a 70 mm diameter circle) into 2.5 mm \times 2.5 mm regions containing variable numbers of emitters, then dilated the borders of the regions to minimize the total imaging area while still providing a 330 μ m buffer around each emitter. Figure 3 shows the results of this process; the black circles mark the locations of emitters on the chip, (x_j, y_j) , and the red lines mark the boundaries of different imaging regions. We then aggregated these regions into groups of 40 ± 10 emitters each, ensuring each group contains regions distributed from the beginning to the end of the fabrication process (to be able to diagnose systematic changes in emitter geometry over time). Finally, we randomly chose three of these groups to sample. The regions selected for sampling total 143 emitters and are those with bold borders in Fig. 3.

In Fig. 4, we plot the geometries of all 143 emitters measured from these regions as a scatter matrix. This figure mimics the structure of a covariance matrix, such that the diagonal entries are histograms for each parameter and the off-diagonal elements are scatter plots over the pairwise subspaces. In addition to R_c , α , and r_b , we also plot the emitter height h for reference (see Fig. 1).






Figure 3. Division of emitter chip 202 into regions for emitter geometry characterization; the black circles are the emitters, the red squares are the regions, and the bold regions are those that were measured




Figure 4. Measured emitter geometries from emitter chip 202, with colors denoting separate populations (indexed $n = 1, 2, 3$); the diagonal entries are histograms for each parameter, while the off-diagonal entries are scatter plots over pairs of parameters

We find the emitter tip radius is peaked between 20-25 μm with spread order 15 μm , the cone angle spans a range 16-19 deg, the basal radius varies from 130-150 μm , and the emitter height lies within 400-425 μm . R_c and h exhibit long tails and a strong anticorrelation with each other, evidenced in their pairwise scatter plot. This relationship occurs because the cutting tools produce very consistent bodies to the cone— r_b and α are comparatively narrowly distributed—but poorly defined wear processes (e.g., grinding by machining swarf) cause the emitter to become duller and shorter as extra material is removed.^{10,13}

Notably, the emitters exhibit three distinct populations (the peaked structure to the distributions over α and r_b most clearly illustrates this division). We color these populations separately and overlay them to help differentiate, where light blue is the first population, red-orange the second, and yellow-orange the third. We attribute these differences to the fact that the emitter chip was machined using multiple cutting tools. The emitters are machined in a fixed sequence $j = 1, \dots, 6102$ —the same as for the \mathbf{Em}_j 's. The CNC mill is programmed such that a tool begins cutting emitters at $j = 1$ and continues to cut until it is swapped for a new tool at some threshold j_1 , at which point a second tool cuts until some threshold j_2 , and so on, such that the populations of emitters cut by each tool occupy contiguous blocks along the sequence.

The three populations denoted in Fig. 4 are similarly contiguous along j . They do not overlap each other, nor does any population overlap a threshold corresponding to a tool change. The difference between the populations, then, we credit to the finite tolerances with which the tools are fabricated. For example, the diameter of the tool, which must be compensated for to control the emitter basal radius, is quoted only within a tolerance of 13 μm by the manufacturer.

To further illustrate this segmentation of the emitter population, in Fig. 5 we plot the basal radius of each measured emitter as a function of position in the array. As in Fig. 3, each circle represents the lateral




Figure 5. Measured basal radii for emitter chip em202 as a function of position in the array

position of an emitter in the array, with the empty circles being the emitters that were not measured and the others having color scaled to their base radius. Noting that the emitters are machined sequentially spiraling inward—i.e., as successive hexagonal shells from the exterior to interior of the chip—we confirm the distinct

and sequential shift between populations in the figure (the change in color on the colorscale between the outer, middle, and innermost sites). Indeed, a few of the regions imaged in the array (compare with Fig. 3) include emitters from multiple populations.

The only difficulty with our explanation for the different populations is that only two tools were used to cut the chip, and hence one tool has apparently produced two populations. Upon examination of the offending cutting tool under a microscope, we determined that one of the teeth of the end mill had been egregiously chipped, much more so than any tooth on the same or the other tool. We therefore posit that this chip developed at some point while the tool was cutting, the resulting change in its geometry causing subsequent emitters to be cut differently. This post-chip population is rendered in red-orange in Fig. 4; the heightened cone angle and basal radius are consistent with there being a chip in the tool, such that it is removing less material from the substrate (and leaving more behind). Since we cannot say exactly when, sequentially, this change occurred, we must account for its uncertainty in our analysis.

2. Inference

To determine what the emitters we measured imply about the emitters we did not, we must first model the populations themselves. Given the cohesive identity of the three populations (being cut by effectively distinct tools), we suppose that they are independent from one another, and index them sequentially by $n = 1, 2, 3$. We further suppose that each population is described by a distribution of the form

$$\begin{bmatrix} \ln R_{cj} \\ \ln(\tan \alpha_j) \\ \ln r_{bj} \end{bmatrix} \sim \mathcal{N}(\boldsymbol{\mu}_n, \text{diag}(\boldsymbol{\sigma}_n) \mathbf{R}_n \text{diag}(\boldsymbol{\sigma}_n)), \quad (28a)$$

$$\boldsymbol{\mu}_n = \begin{bmatrix} \mu_{R_{cn}} \\ \mu_{\alpha_n} \\ \mu_{r_{bn}} \end{bmatrix}, \quad (28b)$$

$$\boldsymbol{\sigma}_n = \begin{bmatrix} \sigma_{R_{cn}} \\ \sigma_{\alpha_n} \\ \sigma_{r_{bn}} \end{bmatrix}, \quad (28c)$$

$$\mathbf{R}_n = \begin{bmatrix} 1 & \rho_{R_c, \alpha_n} & \rho_{R_c, r_{bn}} \\ \rho_{R_c, \alpha_n} & 1 & \rho_{\alpha, r_{bn}} \\ \rho_{R_c, r_{bn}} & \rho_{\alpha, r_{bn}} & 1 \end{bmatrix}, \quad (28d)$$

where the last emitter in each population is denoted j_n (with $j_0 = 0$) such that $j_{n-1} < j \leq j_n$ for population n . Eqs. (28) indicate that each emitter is treated as a random variable independent and identically distributed according to a multivariate normal distribution with mean vector $\boldsymbol{\mu}_n$, standard deviation vector $\boldsymbol{\sigma}_n$, and correlation matrix \mathbf{R}_n .

Notably, we describe this normal distribution over log space, such that the marginal distributions for R_c , $\tan \alpha$, and r_b are log-normally distributed. Phenomenologically, this expresses that these quantities are strictly positive and helps model the long tail toward larger values of R_c evident in Fig. 4 and previously seen in Ref. 13. Encapsulating α within the tangent merely serves to transform it monotonically from the interval $(0, \frac{\pi}{2})$ to the interval $(0, \infty)$ —consistent with the other two parameters—and is better posed because α only appears in the model as mediated by trigonometric functions in the first place.

The parameters $\boldsymbol{\mu}_n$, $\boldsymbol{\sigma}_n$, and \mathbf{R}_n encode the manufacturing tolerances in fabricating emitters for the different populations, and we infer them from our measurements so that we can feed them forward into our analysis over θ . To do so, we condition the distribution Eq. (28). We denote the set of all indices j corresponding to measured emitters as Ω_m and our measurement of an emitter geometry as $\mathbf{Em}^{(j)}$, such that the set of all emitter geometry measurements, $D_{\mathbf{Em}}$, is

$$D_{\mathbf{Em}} = \bigcup_{j \in \Omega_m} \mathbf{Em}^{(j)} = \bigcup_{j \in \Omega_m} \{R_c^{(j)}, \alpha^{(j)}, r_b^{(j)}\}. \quad (29)$$

As in Sec. III.C, our prior for inference conditioned on D is the posterior from our inference conditioned on $D_{\mathbf{Em}}$, $P(\cup_j \mathbf{Em}_j \mid D_{\mathbf{Em}})$. Approximating there to be zero error in the measurement of the emitter geometries,



we factorize the emitters we measured directly out of the joint distribution as

$$P(\cup_j \mathbf{Em}_j \mid D_{\mathbf{Em}}) = P(\cup_{j \in \Omega_m^C} \mathbf{Em}_j \mid D_{\mathbf{Em}}) \prod_{j \in \Omega_m} \delta(\mathbf{Em}_j - \mathbf{Em}^{(j)}), \quad (30)$$

where Ω_m^C denotes the set of indices not measured directly and δ denotes the delta distribution, indicating we are certain of the geometry of the measured emitters.

Letting

$$\theta_{\mathbf{Em}} = \bigcup_n \theta_{\mathbf{Em}n} = \bigcup_n \{\boldsymbol{\mu}_n, \boldsymbol{\sigma}_n, \mathbf{R}_n\} \quad (31)$$

be the set of parameters describing the populations of emitters, and recalling j_1 is uncertain because we are unsure at which index the tool chipped, we express our knowledge in the emitters we did not measure directly as

$$\begin{aligned} P(\cup_{j \in \Omega_m^C} \mathbf{Em}_j \mid D_{\mathbf{Em}}) &= \int P(\cup_{j \in \Omega_m^C} \mathbf{Em}_j, \theta_{\mathbf{Em}}, j_1 \mid D_{\mathbf{Em}}) d\theta_{\mathbf{Em}} dj_1, \\ &= \int P(\cup_{j \in \Omega_m^C} \mathbf{Em}_j \mid \theta_{\mathbf{Em}}, j_1) P(\theta_{\mathbf{Em}}, j_1 \mid D_{\mathbf{Em}}) d\theta_{\mathbf{Em}} dj_1, \end{aligned} \quad (32a)$$

$$P(\theta_{\mathbf{Em}}, j_1 \mid D_{\mathbf{Em}}) = P(j_1 \mid D_{\mathbf{Em}}) P(\theta_{\mathbf{Em}} \mid D_{\mathbf{Em}}). \quad (32b)$$

The second equality follows by factorization. It is not necessary to retain the conditioning on $D_{\mathbf{Em}}$ in the first term because the distribution over unmeasured geometries is conditionally independent from the values of the measured ones given the parameters describing the population—that is, $D_{\mathbf{Em}}$ only informs $\cup_{j \in \Omega_m^C} \mathbf{Em}_j$ through $\theta_{\mathbf{Em}}$, so conditioning on $\theta_{\mathbf{Em}}$ already provides this information. The third equality follows under an assumption our division of populations as presented in Fig. 4 is correct, which is very highly likely given the data. The distribution $P(\cup_{j \in \Omega_m^C} \mathbf{Em}_j \mid \theta_{\mathbf{Em}}, j_1)$ is exactly the multivariate normal of Eqs. (28), which, letting $\Omega_{mn}^C = \{j \in \Omega_m^C : j_{n-1} < j \leq j_n\}$, is expressed

$$P(\cup_{j \in \Omega_m^C} \mathbf{Em}_j \mid \theta_{\mathbf{Em}}, j_1) = \prod_n \prod_{j \in \Omega_{mn}^C} (2\pi)^{-\frac{3}{2}} \det(\text{diag}(\boldsymbol{\sigma}_n)^2 \mathbf{R}_n)^{-\frac{1}{2}} \exp[\Lambda_{n,j}], \quad (33)$$

$$\Lambda_{n,j} = -\frac{1}{2} \left(\begin{bmatrix} \ln R_{cj} \\ \ln(\tan \alpha_j) \\ \ln r_{bj} \end{bmatrix} - \boldsymbol{\mu}_n \right)^T \left(\text{diag}(\boldsymbol{\sigma}_n)^2 \mathbf{R}_n \right)^{-1} \left(\begin{bmatrix} \ln R_{cj} \\ \ln(\tan \alpha_j) \\ \ln r_{bj} \end{bmatrix} - \boldsymbol{\mu}_n \right) \quad (34)$$

Our prior belief in the emitter geometries is encapsulated, then, by our inference over the population parameters, $\theta_{\mathbf{Em}}$, which we again express hierarchically as

$$P(\theta_{\mathbf{Em}} \mid D_{\mathbf{Em}}) \propto P(D_{\mathbf{Em}} \mid \theta_{\mathbf{Em}}) P(\theta_{\mathbf{Em}}). \quad (35)$$

The likelihood $P(D_{\mathbf{Em}} \mid \theta_{\mathbf{Em}})$ is also described by Eqs. (28). It thus has the same form as Eq. (33), merely making the substitutions $R_{cj} \rightarrow R_c^{(j)}$, $\alpha_j \rightarrow \alpha^{(j)}$, $r_{bj} \rightarrow r_b^{(j)}$ and $\Omega_{mn}^C \rightarrow \Omega_{mn} = \{j \in \Omega_m : j_{n-1} < j \leq j_n\}$. That is, the measured geometries are assumed generated from the same distribution as are the unknown geometries.

We close our specification of the prior over the \mathbf{Em}_j 's by then assigning the hierarchical priors of Tab. 4. We prescribe uniform priors on the $\boldsymbol{\mu}$'s (location parameters) and uniform priors on the magnitudes of the $\boldsymbol{\sigma}$'s (scale parameters). For the correlation matrix, we elect a Lewandowski-Kurowicka-Joe (LKJ) distribution.⁴³ It is not possible to independently prescribe distributions over the individual correlation terms because the correlation matrix is constrained to be positive definite. For example, if $\rho_{R_c, \alpha_n} = \rho_{R_c, r_{b_n}} = 1$, then $\rho_{\alpha, r_{b_n}} = 1$ as well (if R_c and α are perfectly correlated and R_c and r_b are perfectly correlated, then α and r_b must necessarily be perfectly correlated). The LKJ distribution encodes this dependency such that for unity parameter (chosen here) the density is distributed uniformly over all valid correlation matrices.

Finally, we assume it is equally likely for the tool to have become chipped at any point between the last emitter measured to be in population 1 and the first emitter measured to be in population 2. During the fabrication process, the machine will occasionally have to pause cutting emitters to remove extra material



Parameter	Prior Distribution
$\mu_{R_{cn}}$	$\mu_{R_{cn}} \sim \mathcal{U}(-\infty, \infty)$
$\mu_{\alpha n}$	$\mu_{R_{cn}} \sim \mathcal{U}(-\infty, \infty)$
$\mu_{r_{bn}}$	$\mu_{R_{cn}} \sim \mathcal{U}(-\infty, \infty)$
$\sigma_{R_{cn}}$	$\ln \sigma_{R_{cn}} \sim \mathcal{U}(-\infty, \infty)$
$\sigma_{\alpha n}$	$\ln \sigma_{\alpha n} \sim \mathcal{U}(-\infty, \infty)$
$\sigma_{r_{bn}}$	$\ln \sigma_{r_{bn}} \sim \mathcal{U}(-\infty, \infty)$
\mathbf{R}_n	$\mathbf{R}_n \sim \text{LKJ}(1)$
j_1	$j_1 \in \{\max_{\Omega_{m_1}} j, \dots, \min_{\Omega_{m_2}} j\}$ $P(j_1 = j \mid D_{\mathbf{Em}}) \propto \# \text{ of cuts from } [j, j+1)$

Table 4. Prior distribution over parameters defining the emitter geometry populations, $\theta_{\mathbf{Em}_n}$ and j_1

where emitters cannot be placed (e.g., where they would be shadowed by the extractor support frame, see the star pattern of Fig. 5). The tool must make an additional cut to remove each of these would-be emitters, so we leverage a probability proportional to the total number of cuts the tool has to make between two realized emitters—that is, a distribution uniform over all cuts. This more evenly distributes the prior density over all chances the tool would have to break.

E. Extractor Offset

Similar to our treatment of the emitter geometries, we ground our inference over the extractor offsets in measurements. For the experiments of Ref. 35, the extractor offset was measured by varying the focus of a microscope at 5 different emitters across the chip, four at the edge of the chip in each of four directions and one at the center, which we report in Tab. 5.

Location	Extractor offset Δ [μm]
North	558
South	544
East	602
West	609
Center	507

Table 5. Extractor offset measurements for chipset 202

Each of the locations corresponds to a specific emitter at which the extractor offset was measured, $j \in \Omega_{m,ex}$ following earlier notation, and so the measurements constitute the set of data

$$D_{\Delta} = \bigcup_{j \in \Omega_{m,ex}} \Delta^{(j)}, \quad (36)$$

on which we seek to condition our knowledge of the Δ_j ’s as with earlier analyses. In contrast to the \mathbf{Em}_j ’s, which vary stochastically from emitter to emitter, the extractor offset tends to vary smoothly across the array. As an approximation, then, we model the spatial variation in the extractor offset across the array as an interpolation—MATLAB’s linear scattered 2d interpolant—between the extractor offsets at the locations where we took measurements, i.e.,

$$\Delta = \Delta(x, y; \cup_{j \in \Omega_{m,ex}} (x_j, y_j, \Delta_j)), \quad (37)$$

where again (x, y) is the lateral position in the array (see Fig. 3), and $\Omega_{m,ex}$ is the set of all indices where the extractor offset was measured. The terms behind the semicolon indicate that the interpolant is parameterized by a series of ordered sets of test points.



We model the x_j 's and y_j 's as deterministic, but the offset measurements are subject to error because of the finite depth of focus of the microscope. We treat this as Gaussian noise,

$$\Delta^{(j)} = \Delta_j + \eta_{\Delta}^{(j)}, \quad \eta_{\Delta}^{(j)} \sim \mathcal{N}(0, \sigma_{\Delta}^2), \quad (38)$$

where for the microscope used we take $\sigma_{\Delta} = 2.5 \mu\text{m}$. Following earlier treatment of the emitter geometries, we express our prior knowledge in the Δ_j 's conditioned on D_{Δ} as

$$P(\cup_j \Delta_j \mid D_{\Delta}) = P(\cup_{j \in \Omega_{m,ex}^C} \Delta_j \mid \cup_{j \in \Omega_{m,ex}} \Delta_j) P(\cup_{j \in \Omega_{m,ex}} \Delta_j \mid D_{\Delta}), \quad (39)$$

where

$$P(\cup_{j \in \Omega_{m,ex}^C} \Delta_j \mid \cup_{j \in \Omega_{m,ex}} \Delta_j) = \prod_{j \in \Omega_{m,ex}^C} \delta(\Delta_j - \Delta(x_j, y_j; \cup_{j \in \Omega_{m,ex}} (x_j, y_j, \Delta_j))). \quad (40)$$

Hence,

$$P(\cup_{j \in \Omega_{m,ex}} \Delta_j \mid D_{\Delta}) \propto P(D_{\Delta} \mid \cup_{j \in \Omega_{m,ex}} \Delta_j) P(\cup_{j \in \Omega_{m,ex}} \Delta_j), \quad (41)$$

$$P(D_{\Delta} \mid \cup_{j \in \Omega_{m,ex}} \Delta_j) = \prod_{j \in \Omega_{m,ex}} \frac{1}{\sigma_{\Delta} \sqrt{2\pi}} \exp \left[-\frac{(\Delta_j - \Delta^{(j)})^2}{2(\sigma_{\Delta})^2} \right], \quad (42)$$

$$P(\cup_{j \in \Omega_{m,ex}} \Delta_j) = \prod_{j \in \Omega_{m,ex}} \mathcal{U}(-\infty, \infty). \quad (43)$$

That is, the extractor offsets at the sites in the array we did not measure are exactly known given those at the sites we did, under assumption of the interpolant model, Eq. (40); the likelihood of our observations, Eq. (42), is given by our noise model, Eq. (38); and we assign a uniform prior, Eq. (43), over the Δ_j 's at the sites we took measurements.

F. Field Model

With our description of the Δ_j 's given, the functional form for the electric field profile, Eq. (11), remains the sole outstanding component of our inference over the current model parameters. The field profile on an emitter plays a key role in our predictions because it dictates both when a meniscus deforms into a Taylor cone—Eq. (3)—and the amount of current sprayed from that cone when it does, Eq. (4). Furthermore, the spatial dependence of $\hat{E}(s)$ can be crucial in capturing the macroscopic behavior of the emitter as the composition over multiple emission sites, because these sites are themselves distributed in space.^{27, 34, 44, 45} To capture this dependency with comparatively high fidelity, then, we rely on electrostatic simulations of the emitter-extractor geometries.

We model each emitter and extractor using the geometry of Fig. 1, which is 2d axisymmetric about the emitter. While in practice the emitters are not perfectly concentric with their respective apertures, for our system this misalignment is small (order 5-10 μm on the scale of a 230 μm radius aperture) and can be safely neglected (compare also with the 3d simulation of Ref. 34). The surface of the emitter and the emitter basal plane are modeled as an equipotential surface at unity voltage, and the resilient extractor architecture is modeled as a thin (1 μm) grounded sheet on top of a dielectric block the thickness of the extractor. For this block the relative permittivity is $\varepsilon_r = 6$, consistent with the properties of MACOR. Lastly, the domain is discretized with an unstructured mesh and solved using a finite element method (MATLAB's Partial Differential Equations Toolbox).

Figure 6 shows the 1d parameterized field profile $\hat{E}(s)$ extracted from the simulation of a representative emitter-extractor geometry, while Fig. 7 shows the full 2d simulation results of the same, including the meshing (a), potential profile (b), and electric field magnitude (c). The black line in Fig. 6 is the raw result from the simulation, while the dashed red line is the result of fitting a parameterized Bézier curve to the profile, with control points the red asterisks. We return to this topic in Sec. III.G.2. The gray dashed line indicates the position at the edge of the emitter tip, by which point the field has decayed to $\sim 70\%$ of its peak value, tending to restrict emission to the emitter tip.³⁴






Figure 6. Electric field profile extracted from the simulation of Fig. 7, with raw simulation in black and a Bézier compression of that profile in red, including control points

G. Numerical Methods

Fully and exactly characterizing the posterior distribution over the parameters θ and θ' is intractable as a result of its nonlinearity and high dimensional nature—there are 4 uncertain geometric parameters defining each of 6102 emitters in the array. Further, some parameters cannot be easily differentiated: swapping the geometry of two different emitters, for example, would change their individual emitter currents but not the total array current. As such, we must limit the scope of our inference and introduce numerical approximations to make the problem tractable.

1. Pseudomarginalization

Crucially, we recognize that some of the parameters in θ are merely incidental to the analysis. That is, their value is necessary to compute the current model—and hence their uncertainty must be accounted for in performing our inference—but we ultimately do not need to learn their value. This is either because we anticipate our data will not be very informative relative to our prior knowledge or because they are unique to this thruster and not extensible to other systems (e.g., the exact emitter geometries). As such, we partition θ as $\theta = \theta^* \cup \tilde{\theta}$, where θ^* denotes the subset of parameters we will learn. For the analysis presented here, we focus our attention to the parameters $\theta^* = \{\zeta_0, \zeta_1, \bar{p}_2, m\}$. These constitute the key missing links in the model that our prior knowledge does not inform well, so we must infer them that we may leverage the model as a diagnostic and predictive tool. The remaining parameters, $\tilde{\theta} = \{\beta_0, \beta_1, \psi, p_1, \kappa, P_r, \gamma, \rho, \mu, \xi\} \cup \bigcup_j \{R_{cj}, \alpha_j, r_{bj}, R_{Aj}, t_{ej}, \Delta_j\}$, we then marginalize out:

$$P(\theta^*, \theta' | D) \propto \int P(D | \theta^*, \tilde{\theta}, \theta') P(\tilde{\theta}) d\tilde{\theta} P(\theta^*) P(\theta'). \quad (44)$$

Essentially, the marginalization integral of Eq. (44) indicates that our knowledge of the target parameters, θ^* , is an average over all possible values of $\tilde{\theta}$ within our prior knowledge.

Practically, this integral is also intractable, and so in evaluating the likelihood, we substitute a Monte Carlo estimate, forming the pseudomarginalization⁴⁶

$$P(\theta^*, \theta' | D) \approx \frac{1}{N_r} \sum_{l=1}^{N_r} \left[P(D | \theta^*, \tilde{\theta}_l, \theta') \right] P(\theta^*) P(\theta'), \quad \tilde{\theta}_l \sim P(\tilde{\theta}). \quad (45)$$




Figure 7. Representative electrostatic field simulations of an emitter: a) simulation mesh; b) electrostatic potential; c) field strength

where N_r is the number of samples forming the estimate. That is, we take many samples, the $\tilde{\theta}_i$'s, of the nuisance parameters according to their prior distribution, $P(\tilde{\theta})$, and compute the arithmetic mean of the likelihood evaluated for each of the samples. In the limit $N_r \rightarrow \infty$, this estimate converges to the true likelihood.

In this work, we select $N_r = 1000$, limited principally by computational expense. The largest bottleneck forms in performing the electrostatic simulations to compute the position-dependent field profile, $\hat{E}(s)$. A simulation for a single emitter takes order 0.1 CPU seconds to perform, and hence all $N_e \times N_r = 6.1$ million evaluations for a single evaluation of the posterior would take 170 CPU hours. To mitigate this bottleneck, we further approximate the pseudomarginalization by using a fixed set of samples for the $\tilde{\theta}_i$'s. Ideally, one would draw a new set of $\tilde{\theta}_i$'s each time the posterior was evaluated, such that the error inherent in the pseudomarginalization of Eq. (45) tends to average out over many evaluations. Using a fixed set of samples permits bias by preventing this from happening. Critically, however, it allows us to compute the field profiles only once for each of the N_r realizations and then simply store these results to use each time we need to evaluate the posterior, dramatically lowering computational costs. Additionally, the bias is small for a sufficiently large number of samples.

2. Bézier parameterization

While the computational overhead in terms of CPU time is greatly reduced by introducing the fixed set of pseudomarginal samples, the need to store the field profiles remains a memory bottleneck. Since the field profiles as output by the electrostatic model constitute a series of order 1000 discrete s - $\hat{E}(s)$ points, the total memory needed to store every field profile for $N_r = 1000$ is of order 100 GB. Additionally, performing the integrals of Eq. (4) over so finely discretized a representation may unnecessarily slow analysis.

Rather than store the field profiles directly, we instead parameterize each by a sequence of two cubic Bézier curves. Bézier curves are an expressive family of parametric polynomial curves that excel at approximating smooth geometries.⁴⁷ We divide the s - \hat{E} curve into two domains, the emitter cap, $s \in [0, s_c]$, and the emitter body, $s \in [s_c, s_c + s_b]$. Defining a cubic Bézier curve over each domain, parameterizing as $s = s(\nu)$

and $\hat{E}(s) = \hat{E}(s(\nu))$ for $\nu \in [0, 2]$, and anchoring the curves to the 3 points $(0, \hat{E}(0))$, $(s_c, \hat{E}(s_c))$, and $(s_c + s_b, \hat{E}(s_c + s_b))$, we have then

$$s(\nu) = \begin{cases} s_c [3s_1(1-\nu)^2\nu + 3s_2(1-\nu)\nu^2 + \nu^3], & \nu \in [0, 1], \\ s_c + s_b [3s_3\nu^2(\nu-1) + 3s_4\nu(\nu-1)^2 + (\nu-1)^3], & \nu \in [1, 2], \end{cases} \quad (46a)$$

$$\hat{E}(s(\nu)) = \begin{cases} \hat{E}(0) [(1-\nu)^3 + 3\hat{E}_1(1-\nu)^2\nu + 3\hat{E}_2(1-\nu)\nu^2 + \hat{E}_c\nu^3], & \nu \in [0, 1], \\ \hat{E}(0) [\hat{E}_c\nu^3 + 3\hat{E}_3\nu^2(\nu-1) + 3\hat{E}_4\nu(\nu-1)^2 + \hat{E}_b(\nu-1)^3], & \nu \in [1, 2], \end{cases} \quad (46b)$$

$$\hat{E}_c = \frac{\hat{E}(s_c)}{\hat{E}(0)}, \quad (46c)$$

$$\hat{E}_b = \frac{\hat{E}(s_c + s_b)}{\hat{E}(0)}. \quad (46d)$$

Apart from s_c and s_b , which can be computed online, the field profile is defined solely by the set of 11 parameters $\{\hat{E}(0), s_1, s_2, \hat{E}_1, \hat{E}_2, \hat{E}_c, s_3, s_4, \hat{E}_3, \hat{E}_4, \hat{E}_b\}$. The quantities $\hat{E}(0)$, \hat{E}_c , and \hat{E}_b encode the field strength—computed at the emitter tip, edge of the cap, and base of the body, respectively—and act as fixed anchor points that tie the Bézier parameterization to the simulated profile. The remaining parameters, (s_1, \hat{E}_1) , (s_2, \hat{E}_2) , (s_3, \hat{E}_3) , and (s_4, \hat{E}_4) are free anchor points for the curve that are adjusted to match the computed profile. In particular, we fit them to the simulated profile by least-squares regression.

In this way, we have formed a compact representation of the field profile, $\hat{E}(s)$. We use this parameterized form directly in computing our current model predictions, as we find that they are able to reconstruct the simulated profiles with very little relative error (¡0.1% typical on the cap, ¡1% typical on the body).

3. Sampling

It is necessary in forming the pseudomarginalization of Eq. (45) to sample over parameters whose probability distributions, while known, do not fall or factorize into forms that are readily sampled. To do so, we leverage a Metropolis-Hastings sampler, a Markov chain Monte Carlo technique.⁴⁸ This powerful family of techniques has the capability to draw a sequential chain of samples from broad classes of distributions, but has the key drawback that the samples it generates are generally not independent from adjacent samples in the chain, necessitating more samples be drawn to ensure the same effective sample size (i.e., to encode the same information about the distribution).

Relatedly, while introducing the pseudomarginalization has greatly reduced the dimensionality of the posterior over θ^* and θ' , the nonlinearity of the current model in its parameters means that moments over the posterior (e.g., to compute the average value of a parameter or make a probabilistic prediction for array current) lack closed form. Consequently, we also rely on sampling the posterior via Metropolis-Hastings to approximate functionals over the distribution via Monte Carlo.

IV. Inference Results

We now execute the inference we defined in the preceding section, characterizing our state of knowledge in the parameters. We begin by reviewing the results of our nested inference over the onset criterion parameters (Sec. IV.A) and the emitter and extractor geometry (Sec. IV.B). After performing these component inferences, we leverage them to draw samples directly over the parameters of interest, θ^* (Sec. IV.C).

A. Onset Parameters

Consistent with our treatment in Sec. III.C, we drew 20000 samples from $P(\beta_0, \beta_1 \mid D_\beta)$ to infer the onset criterion parameters from the data produced by the simulations of Coffman *et al.*⁸ Given the form of Eq. (26) it was useful to do so by sampling instead from $P(\beta_0, \beta_1, \sigma_\beta \mid D_\beta)$ and then to obtain the target distribution by marginalization. We tabulate the *maximum a posteriori* (MAP), median, and interquartile range over each parameter in Tab. 6. The MAP is the highest posterior probability point (the mode) of the distribution, and represents the single best estimate to the parameters. The median serves as one measure of the average



Parameter	MAP	Median	IQR
β_0	0.519	0.519	0.007
β_1	0.400	0.400	0.013

Table 6. Summary statistics for prior inference over the onset criterion, $P(\beta_0, \beta_1 \mid D_\beta)$

value of the distribution; they are equal to the MAP because there is not skew in the distribution. The interquartile range (IQR) is the difference between the 3rd and 1st quartile (75th and 25th percentile), and is a robust measure of the spread of the distribution—that is, the uncertainty in the parameter.

We find the MAP and median values agree with the least-squares fit employed in previous analysis.³⁴ This agreement follows because we assumed no prior knowledge (uniform priors) and modeled the likelihood as constant Gaussian noise, under which assumptions the MAP is mathematically equivalent to a least squares estimate. The key difference is that by adopting a probabilistic description over the parameters we have a quantifiable measure of their uncertainty. This stems fundamentally from the fact that the model does not perfectly reconstruct the data and that the data are subject to error. There exist, then, multiple plausible values for the parameters, the measure of their plausibility being exactly the posterior.

To illustrate this uncertainty, we form a posterior predictive distribution. That is, we take our knowledge in the model parameters informed by the data and then use the model to predict additional measurements subject to our uncertainty,

$$P(D'_\beta \mid D_\beta) = \int P(D'_\beta, \beta_0, \beta_1, \sigma_\beta \mid D_\beta) d\beta_0 d\beta_1 d\sigma_\beta \quad (47)$$

$$= \int P(D'_\beta \mid \beta_0, \beta_1, \sigma_\beta) P(\beta_0, \beta_1, \sigma_\beta \mid D_\beta) d\beta_0 d\beta_1 d\sigma_\beta. \quad (48)$$

This expresses that our predictions for new observations, D'_β , are generated according to our likelihood model, $P(D'_\beta \mid \beta_0, \beta_1, \sigma_\beta)$, but integrated over our uncertainty in the model parameters. Functionally, we generate samples for D'_β by taking our samples for β_0 and β_1 , predicting the model at a range of measurement locations for each one, and then sampling noise consistent with σ_β to model the experimental error in the measurement. We show the results of applying this process in Fig. 8. The blue are the training data, and the




Figure 8. Posterior predictions for the onset criterion; blue are the training data and black are the predictions from the trained model

solid black line is the median prediction as a function of \bar{P}_r . The shaded region between the dashed black lines contains 90% of the predicted probability, spanning the 5th percentile prediction to the 95 percentile prediction.

These posterior predictions serve as a check that our model is able to reconstruct the training data. Essentially, if we went to predict new measurements and found them inconsistent with our existing ones, we would conclude that we had failed to model some key feature of the problem. As is evident in Fig. 8, this is not the case. The data are linear with some modest deviation, and the model captures this.

B. Emitter and Extractor Geometry

Our inference over the emitter geometries is encapsulated by the population parameters $\theta_{\mathbf{Em}}$, and so we drew 50000 samples over the distribution $P(\theta_{\mathbf{Em}} | D_{\mathbf{Em}})$. Noting that the three different populations are independent, this distribution factorizes as

$$P(\theta_{\mathbf{Em}} | D_{\mathbf{Em}}) = \prod_n P(\theta_{\mathbf{Em}n} | D_{\mathbf{Em}n}), \quad (49)$$

where $D_{\mathbf{Em}n} = \cup_{j \in \Omega_{mn}} \mathbf{Em}^{(j)}$. That is, inferring parameters over all three populations is equivalent to inferring the parameters of each population separately over its respective dataset, which is more efficient.

Figure 9 summarizes our inference over the emitter populations. Rather than plot distributions over




Figure 9. Distribution over median (first row) and IQR (second row) inferred for tip radius, cone angle, and base radius of the three emitter populations (different colors)

the parameters directly, we show their implications for the emitter geometry (once transformed out of log space, etc.). The first row of figures shows the distribution over medians for R_c , α , and r_b . That is, for each sample of parameters drawn from Eq. (49), we compute the corresponding median over the emitter geometries, consistent with Eqs. (28), for each of the geometry parameters. We then render histograms over these samples for each population, where we have colored them consistent with the division of populations in Fig. 4 and normalized so that the histograms integrate to unity. The second row of figures applies the same methodology to plot the IQR for each parameter and population.

These distributions represent our uncertainty in the average emitter geometry (median) and spread in emitter geometry (IQR) for their respective parameters and populations. For α and r_b , we find there to

be comparatively little uncertainty in the distribution. The median cone angle of each population is known within 5 milliradians, and the median base radius of each population has an uncertainty only $\pm 0.5 \mu\text{m}$. Further, these distributions corroborate quantitatively our earlier assertion that the cone angle and basal radius are well controlled by the machining process: the IQRs for α are about 2% of their medians, while the IQRs for r_b are about 1% of their medians—i.e., they are subject to small error on the scale of the parameter. Indeed, we discern, consistent with the clustering of measured geometries of Fig. 4, that the variation between populations (and hence, tools) is substantially larger than the variation within a population.

In contrast, the inference on the tip radii of the populations is muddled. The data only provide enough information to estimate the median tip radius of each population within $\pm 5 \mu\text{m}$, of order 15-25% relative error. We also note that the IQRs for tip radius are of comparable order to the median themselves, and are greatly uncertain; taking the red-orange population as an example, the spread in tip radius could lie plausibly anywhere from 10 to 30 μm . Hence the emitter tip radius is not only highly varied across the array, but it is also difficult to assess exactly how highly varied.

This uncertainty is disconcerting because the model is a strong function of R_c . The tip radius plays the key role of amplifying the electric field on the emitter, much more so than does the spacing between the electrodes or the other geometric parameters.⁴⁹ Since the electric field is so central in determining onset and emission, we anticipate that the large uncertainty in R_c will propagate into our predictions. We return to this point in the discussion (Sec. V).

To generate realizations from the distribution $P(\cup_j \mathbf{Em}_j \mid D_{\mathbf{Em}})$ and form the $\tilde{\theta}_i$'s of our pseudomarginalization, Eq. (45), we take single samples over the population parameters and then draw from the conditional distribution $P(\cup_{j \in \Omega_m^c} \mathbf{Em}_j \mid \theta_{\mathbf{Em}}, j_1)$ —i.e., we sample the mean and covariance for each population and then draw from those distributions for each emitter within each of those populations. This produces a map of what the geometry of each emitter in the array might look like, subject to our uncertainty. We plot the tip radius for each emitter in the array for one such realization in Fig. 10. As in earlier figures, the position of each circle corresponds to the position of an emitter, with its color is mapped to different values for R_c . The larger circles in the array mark the emitters that were measured directly (compare with Fig. 3), while the outlines color code the regions corresponding to the different populations.

Evident in Fig. 10 is the large variance in R_c found by our measurements and subsequently inferred for the populations (Fig. 9). Specifically, for this realization the medians are (23.4, 18.7, 25.2) μm and the IQRs (15.0, 10.9, 12.1) μm for the respective populations (blue, red-orange, yellow-orange). While this means only order half of all the emitters have a tip radius 20 μm or larger, the “fat” tail of the log-normal distribution of Eq. (28) means that a comparatively great number of emitters have larger tip radii, more than 25% being 30 μm or more. This fat tail models the tendency for some emitters to break during the machining process.^{11,13,50} The resulting large population of dull emitters will have weak field amplification, suggesting many emitters may be inactive in the array.

Similarly, in Fig. 11 we plot a realization of the extractor offset for each emitter in the array (the same realization as for Fig. 10). This figure reflects the bulk tendency for the extractor electrode to sag toward its center (evidenced by the radial gradient), but that there is also azimuthal inhomogeneity resulting from warping in each spoke of the superstructure (that each of the wedges has a different offset relative to the emitter chip).

C. Model Parameters

The target parameters θ^* represent the key unknowns in the model and the crux of our inference. We first composed the set of pseudomarginal samples, $\cup_i \tilde{\theta}_i$. For the onset criterion parameters— β_0 and β_1 —and the emitter-extractor geometries—the \mathbf{Em}_j 's and \mathbf{Ex}_j 's—we subsampled from the results of our nested inferences (Secs. III.C and IV.B, respectively). For the remaining parameters— ψ , p_1 , κ , P_r , γ , ρ , μ and ξ —we drew the $N_r = 1000$ realizations directly from the priors of Tab. 2. After precomputing the field profiles for each emitter of each realization consistent with Sec. III.F (totaling $1000 \times 6102 = 6.1$ million simulations), we were then able to evaluate the pseudomarginal posterior of Eq. (45).

Before drawing samples from the posterior distribution, we ran a numerical optimizer to identify the MAP parameter set for θ^* and θ' , which we collect in Tab. 7. Initializing our Markov chain at this point, we then drew 10000 samples from the posterior using a Metropolis-Hastings sampler. We tabulate the median and interquartile range for each parameter alongside the MAP. We find that ζ_0 , ζ_1 , and \bar{p}_2 are comparatively narrowly distributed, with an IQR order 5-10% of their median. Since our priors were largely chosen to be uninformative, this precise inference means our data were highly informative of these parameters—the






Figure 10. Emitter tip radius for each emitter in the array for a single realization over the emitter geometry distribution $P(\cup_j \mathbf{Em}_j \mid D_{\text{Em}})$; the larger circles denote emitters that were measured directly, and the outlines divide the populations

Figure 11. Extractor offset for each emitter in the array for a single realization over the extractor offset distribution $P(\cup_j \Delta_j \mid D_{\Delta})$; the larger circles denote emitters that were measured directly

likelihood dominates the posterior and there is a narrow range of values that plausibly explain the data under this likelihood. In contrast, the distribution over m is broad, a consequence of the model not being a strong function of m in this regime—large changes in m produce small changes in the predicted current, so it is difficult to determine it very precisely. Conversely, knowing it precisely is not critical to analysis, because it does not have much effect on our predictions.

For other parameters where we did not prescribe arbitrarily uninformative priors, comparing our prior knowledge against the posterior knowledge is an essential means to assess that we learned something about them. Indeed, we find that while the MAP value of \bar{p}_2 is close to our *maximum a priori*, the IQR is narrower (to 1.5 μm from 4.8 μm); that is, while our initial guess to \bar{p}_2 was accurate, we are more confident in that fact given the data. This is because \bar{p}_2 plays a key role in establishing the onset voltage of the emitters (larger menisci have less capillary pressure to resist onset. Large changes in \bar{p}_2 thus induce large changes in the current predicted near onset, where our data is most precise—Eq. (18). Consequently, the likelihood is highly sensitive to \bar{p}_2 .

Relatedly, the learning for the noise floor σ_0 , which we had also given a weakly informative prior, was more modest, with comparable uncertainty to our prior distribution and only a modestly shifted median, from 1 μm to 0.71 μm . This indicates that the data are consistent with our prior assumption but cannot more conclusively confirm it. Referencing Tab. 7, we see this is so because the 5% proportional error we inferred for σ_1 is larger than this noise floor over the entire domain: it indicates a noise increasing from 1

Parameter	MAP	Median	IQR
ζ_0	3.25×10^{-3}	3.25×10^{-3}	0.35×10^{-3}
ζ_1	2.64×10^{-3}	2.69×10^{-3}	0.16×10^{-3}
\bar{p}_2	$12.9 \mu\text{m}$	$12.6 \mu\text{m}$	$1.5 \mu\text{m}$
m	1.73	1.48	0.68
σ_0	$0.50 \mu\text{A}$	$0.71 \mu\text{A}$	$0.91 \mu\text{A}$
σ_1	4.95%	5.55%	1.55%

Table 7. Summary statistics for the posterior over the target parameters, θ^* , and hyper parameters, θ'

μA at 650 V up to 280 μA at 2000 V. Hence, it is difficult to estimate the comparatively small σ_0 term when the noise contributed by the σ_1 term is so much larger.

The other key comparison point for our inference over the model parameters is our previous study inferring the current emission model over a single electrospray.³⁴ Motivated by the suspicion that the parameters ζ_0 and ζ_1 may not be universal physical constants, we declined to incorporate these results into the prior. Indeed, we find that the onset parameters learned here disagree substantially with those inferred in Ref. 34, which found $\zeta_0 = -4.5 \times 10^{-3}$ and $\zeta_1 = 4.1 \times 10^{-2}$ (cf. Tab. 7). On balance, we also prescribed 0 prior probability that ζ_0 could be negative in our analysis. Practically, we first attempted to include the possibility for negative ζ_0 , but revised our analysis finding it could not reconstruct the data. In principle this constitutes a modest “bootstrapping” from the data, but realistically our prior was not active in the inference and serves more as a numerical convenience.

We attribute this discrepancy to the fact that the single emitter experiments of Ref. 27, which formed the training data for the study of Ref. 34, exhibited an anomalously smooth onset. That is, rather than for the spray to turn on at some voltage and there be a jump in the current, as the form of Eq. (4) implies and is often the case,^{5, 18, 45, 50} the authors saw a tenuous beam form at comparatively small current, growing rapidly and smoothly with voltage thereafter. The only means by which Eq. (4) can reconstruct this behavior is for ζ_0 to be negative, predicting current to be suppressed even after onset has been achieved. In the study of Ref. 34, this also induced a strong anticorrelation in ζ_0 and ζ_1 , meaning that the model could not differentiate between a delayed onset at higher current and an earlier onset at lower current.

For the thruster we trained the model on here, we discern that such emitter behavior is an implausible explanation for the array current. Rather, the array current is explained as the composition of many emitters which turn on at different voltages and experience a jump discontinuity in current when they do (consistent with positive ζ_1), a point we return to in Sec. V.A. Importantly, this discrepancy suggests there exist some physics which differentiate the two systems that is not captured in the dimensionless linearization of Eq. (4) and encapsulated via the parameters ζ_0 and ζ_1 . Further, we do not observe the same correlation between these two parameters, instead inferring a weak positive correlation of 0.1. This further suggests the physical consequences of ζ_0 and ζ_1 are distinguishable from each other in present analysis.

V. Discussion

Now that we have inferred the missing model parameters θ^* from our data—training the model—we leverage it as a diagnostic and prediction tool to better understand the dynamics underlying the array. First, as in Sec. IV.A, we form a posterior predictive. Here we do so over the array current,

$$P(D' | D) = \int P(D', \theta, \theta' | D) d\theta d\theta' \quad (50)$$

$$= \int P(D' | \theta^*, \tilde{\theta}, \theta') P(\theta^*, \tilde{\theta}, \theta' | D) d\theta^* d\tilde{\theta} d\theta', \quad (51)$$

where again $P(D' | \theta^*, \tilde{\theta}, \theta')$ is our likelihood and $P(\theta^*, \tilde{\theta}, \theta' | D)$ is the posterior.

However, in adopting the marginalized forms of Eqs. (44) and (45) as concessions to facilitate analysis, we also lost access to the information the data held about the marginalized parameters $\tilde{\theta}$. As such, in making



our posterior predictions we approximate the fuller posterior as

$$P(\theta^*, \tilde{\theta}, \theta' | D) \approx P(\theta^*, \theta' | D)P(\tilde{\theta}). \quad (52)$$

That is, we substitute our prior knowledge in $\tilde{\theta}$ for our posterior knowledge. Supposing the data D are only weakly informative of $\tilde{\theta}$ —as we anticipated was the case—this constitutes a decent approximation, because our prior knowledge and posterior knowledge would not differ much. This approximation is also logically equivalent to making predictions for another thruster, for which we had the same prior assumptions about its propellant properties, charge to mass ratio, etc., and for which we had observed the same things in characterizing its geometry.

We generate our samples from the posterior predictive by taking 1000 of our samples over the marginalized posterior, $P(\theta^*, \theta' | D)$, and the 1000 realizations from our prior over $\tilde{\theta}$ and sampling from the likelihood conditional on each pair of these samples. That is, we make 1 million different model predictions (1000 samples of θ^* by 1000 samples of $\tilde{\theta}$) and then sample from the noise model of Eq. (18) for each one to simulate a measurement. Fig. 12 displays these predictions. The experimental data appear in blue, the




Figure 12. Posterior predictions for array current; experimental data³⁵ (blue) against the median (solid black) and a 90% credible interval (shaded between dashed black)

median prediction is the solid black line, and 90% of the posterior predictive probability is included in the shaded region between the dashed black lines. To highlight both low and high current regimes, we plot the same data on linear and logarithmic scales.

We see that the data lie within the uncertainty of the predictions, which confirms that—with training—the model is able to reconstruct the experimental observations. In particular, the model captures that the array current tends to increase at an increasing rate from 500 V to about 1500 V. The chief regime in which the data and model diverge is at the highest voltages, where the model predicts a saturation to linear growth but the data suggest growth in current may still be faster, a point we return to in Sec. V.A. We observe that the data do not appear to be randomly distributed about the median as Eq. (20) might imply. Rather, the median prediction underpredicts the data about 1000 V but overpredicts the data below 1000 V. This is because the noise of Eq. (20) tends to be a small contribution to the overall uncertainty compared to the structural uncertainty in key parameters like the emitter geometry and propellant properties, which we expand on in Sec. V.B. Altogether, the capability to make these probabilistic predictions represents a key capability in trying to design around these sources of uncertainty (Sec. V.C).

A. Emitter Variability

Understanding emergent phenomena in electrospray array thrusters—bulk behavior that arises only as the composition over many individual emitters—is crucial to predicting array performance because key processes like grid shorting can occur at the level of single sites in an array. Ideally, each emitter in the array would behave the same, such that the array current was simply a multiple of a known, single emitter. As a result of fabrication tolerances and other confounding factors, however, the array current—in particular the voltage-dependent array current—is not so trivially determined.

We next examine predictions made for a single realization of the marginalized parameters $\tilde{\theta}$. Isolating a realization this way serves to illuminate the role of emitter variability in the system by considering one example geometry for an array (for which the geometry of each emitter is different). In particular, we consider the sample we drew to form the pseudomarginalization that maximizes the likelihood at the *maximum a posteriori* (MAP) target parameter values. That is, with

$$(\hat{\theta}^*, \hat{\theta}') = \operatorname{argmax}_{\theta^*, \theta'} P(\theta^*, \theta' | D), \quad (53)$$

being the MAP over our inference, then the realization we examine, $\hat{\hat{\theta}}$, is given by

$$\hat{\hat{\theta}} = \operatorname{argmax}_{\hat{\theta}_i} P(D | \hat{\theta}^*, \hat{\theta}', \hat{\theta}_i). \quad (54)$$

Insofar as the MAP parameters of Eq. (53) represent the single best fit of the target parameters and hyper-parameters to the data, this realization represents the best fit to the data over all realizations forming the pseudomarginalization.

Figure 13 is a scatter plot of the emitter current of each site in the array predicted for $\hat{\theta}^*$ and $\hat{\hat{\theta}}$ at a voltage of $V = 1410$ V. We find the emitter current to vary substantially across the array, from the highest output emitters at about $0.7 \mu\text{A}$ to emitters that have not turned on and so emit no current. Indeed, at this voltage, approximately 50% of all of the emitters in the array have yet to activate. The realization $\hat{\hat{\theta}}$ was also used to create Fig. 10, such that the two figures serve as an emitter by emitter comparison of tip radius to emitter current. We observe, then, that the inactive sites have essentially one to one correspondence with the duller (larger R_c) emitters of Fig. 10, reiterating the dominance of the tip radius in amplifying the electric field and hence determining onset.

We discern that there are comparatively few emitters with a predicted current from 0.1 to $0.4 \mu\text{A}$. Rather, emitters exhibit a dichotomy between being on and sourcing substantial current and being off and sourcing none at all. This is because emitters are predicted to exhibit a strong jump in current toward onset. To elucidate this, in Fig. 14, we plot the individual emitter currents as a function of voltage. The partially transparent curves each show the voltage response of a different emitter, where for comparison we also plot the mean, median, and standard deviation over all emitters (solid, dashed, dot-dashed).

Indeed, Fig. 14 shows that each emitter turns on at some onset voltage and exhibits a jump discontinuity in current when it does, evidenced by the many vertical gray lines. These onset events span the entirety of the domain, such that as voltage is increased on the thruster, new emitters are constantly turning on. At lower voltages, only a handful of the very sharpest emitters are predicted to be active, and it is not until above 1400 V that more than half of them are, as evidenced by the abrupt rise in the median emitter current. Even at 2250 V, the highest voltage predicted here, 10% of the emitters remain inactive.

This individual emitter behavior is a consequence of the positive offset, ζ_0 , in the linearized ionic emission scaling law of Eq. (4) and is consistent with a wide body of single emitter experiments demonstrating a jump in current at onset.^{18,19,50–52} Interestingly, we also observe that the magnitude of the corresponding onset current increases with voltage—duller emitters are predicted to turn on at higher voltage but to start at higher current when they do. That blunt emitter geometries can produce current at heightened sensitivity to operating voltage has similarly been reported,^{53,54} and is consistent with these emitters tending to have lower hydraulic impedance. That is, emission tends to be concentrated at the very tip of an emitter because the electric field is strongest here (compare with Figs. 6 and 7). A sharper tip helps amplify the electric field and hence tends to promote earlier onset and enhanced emission, but because the emission region has smaller area, fluid flow is more heavily constricted—e.g., the inverse R_c scaling of Eq. (5). A duller emitter has a wider area where emission sites can develop, with lower hydraulic impedance, and thus is predicted to source more current for the same electric pressure.






Figure 13. Emitter current predicted for each emitter in the array for the MAP parameters, $\hat{\theta}^*$ and maximum likelihood realization, $\hat{\hat{\theta}}$; the larger circles denote emitters whose geometry was measured, and the outlines divide the different populations




Figure 14. Individual emitter current (transparent) curves as a function of voltage predicted for $\hat{\theta}^*$ and $\hat{\theta}$, alongside the mean (solid), median (dashed), and standard deviation (dotted) emitter current over the entire array

The aggregate effect of these individual current-voltage characteristics is to produce the mean current of Fig. 14, which is simply the array current divided by the total number of emitters. This inflected current versus voltage response—where the current grows faster than linearly at lower voltages but tends to grow more slowly at higher voltages—is endemic to porous array architectures,^{5,10,11,55,56} and is explained by dissimilarity in emitter geometry. Because the field amplification provided by the sharpness of the tip varies from emitter to emitter, the voltage at which an emitter first achieves an electric field strong enough to induce emission does as well. Additionally, once an emitter has activated, the current it sprays increases with voltage. As such, an increasing number of emitters each source an increasing quantity of current, producing the inflection. This effect tends to attenuate at higher voltages, however, because most emitters have already activated and growth reduces to the linear scaling of Eq. (4).

The width of this transition region of positive inflection in voltage space is proportional to the variance in the emitter geometry. If the emitters were identical, they would all activate at the same voltage and the array would immediately saturate to linear growth consistent with the dependence of a single emitter (compare with the comparatively narrow onset region of Ref. 5). This suggests the possibility that inspecting the I-V trace of an array system with this mechanism in mind could serve as an efficient proxy for the more laborious and precise geometry characterization conducted here (see also Refs. 13 and 37). By inspecting over how broad the onset region of an array is and at which voltage it was centered, one can infer what the corresponding mean and variance over tip radius is (given also the aperture radius, recession, etc.).

This mechanism by which the I-V curve is produced as the composition over dissimilar emitters has been posited previously, motivated by experiments showing emission nonuniformity;¹⁵ our ability to reconstruct it grounded in measurements of fabrication tolerances greatly strengthens this hypothesis, and it underscores the importance of emitter variability in modifying a thruster’s effective operational envelope. Across the domain of the posterior predictions in Fig. 12, the array current spans 3 orders of magnitude. Ideally, this would imply a deep throttling capability, with a thruster able to operate at different powers to achieve different mission goals (e.g., fine positioning, stationkeeping, orbit raising). Practically however, because this dynamic range is a consequence of an uncertain fabrication process, it poses challenges to system reliability, which we discuss further in the next two sections.

B. Model Sensitivity

The uncertainty in our posterior predictions, Fig. 12, is a combination over many different sources, from variability in the emitter geometry to uncertainty in the scaling coefficients that abstract away more complex physics. We characterize the relative influence of different error sources by performing a sensitivity study varying the different parameters of the model. The model prediction for $\hat{\theta}^*$ and $\hat{\theta}$ serves as a reference to which we will compare our results. That is, we compute predictions probabilistically as we did for Fig. 12, using multiple samples over the parameters. Instead of allowing all parameters to vary at the same time, however, we divide the model parameters θ into different subsets. For each subset, we compute model predictions where those parameter values are drawn from their respective samples, but where all other parameter values are kept constant at the $\hat{\theta}^*$ - $\hat{\theta}$ reference point. For example, we compute a prediction where we consider the emitter geometry to be uncertain, but not the propellant properties, charge to mass ratio, or any other parameter.

We summarize the result of applying this procedure in Fig. 15. As a function of voltage, we plot the




Figure 15. Sensitivity study for various parameter sets $\hat{\theta} \subseteq \theta$, plotting the IQR over predictions varying only that parameter set, $\text{IQR}[I(\hat{\theta})]$, normalized by the $\hat{\theta}^*$ - $\hat{\theta}$ predicted current, $I(\hat{\theta})$

interquartile range for array current predictions over varying some parameter set $\hat{\theta} \subseteq \theta$ —abbreviated as $\text{IQR}[I(\hat{\theta})]$ —normalized by the array current predicted varying no parameters—denoted simply $I(\hat{\theta})$. The different choices for $\hat{\theta}$ are labeled in the figure. The colored lines correspond to mutually exclusive subsets of the marginalized parameters $\hat{\theta}$. The dashed black line represents allowing all of the marginalized parameters to vary; the dotted line varies only the learned parameters θ^* ; and the solid black line represents allowing all parameters to vary (i.e., the posterior predictions of Fig. 12, less the simulated experimental noise). In this way, we can examine the magnitudes of different error sources relative to the reference simulation over different regimes.

We note that the error considering the entirety of θ tends to decrease as voltage is increased, from 330% relative error at 500 V to only 10% at 2250 V. This behavior results as a voltage-dependent balance between two major error sources: the emitter geometry (green in Fig. 15) and the fluidic properties (yellow). At low voltages, uncertainty in the emitter geometry dominates. This is because only a handful of emitters—the very sharpest—are predicted to be active. Referring again to Fig. 9, the high variance in the emitter geometry, compounded by uncertainty in that variance, means that the number of emitters sharp enough to have achieved onset here varies greatly: it could be none at all, only a few, or several—we cannot say

precisely given our measurements.

The effect of geometric uncertainty diminishes with increased voltage because a greater proportion of emitters have activated. That is, even though individual emitter current can be highly variable as a result of the variance in geometry implied by σ_n , the effect of any one emitter on the array current diminishes as the number of active emitters increases. Put differently, the average over many emitters becomes better defined the more emitters one considers (i.e., the law of large numbers). As such, at higher voltages it is the hierarchical uncertainty in the parameters themselves that dominates, because this changes the average over all emitters.

Analogously, imagine that we are trying to estimate the average height of all ships within a harbor relative to a fixed point on shore. The height of each ship varies based on size from dinghies up to cruise liners. Additionally, the average altitude of each ship varies with how high the tide is, since it will raise all the ships. If the harbor tends not to be very busy such that there are only a few ships at port, then whether or not there are two cruise ships or two dinghies will create a large variance in the average altitude, on top of the fact that the tide goes in and out. This is the case at lower voltages where very few emitters are active. If the harbor instead is very busy such that there are many ships at port, whether there are 100 or 101 dinghies will not make as much difference to the average, but whether the tide is in or out will. This is the case at higher voltages where many emitters are active.

Consequently, we find that the uncertainty in our predictions at elevated voltages is dominated by uncertainty in the propellant properties—yellow in Fig. 15—and hence the operating temperature of the thruster, T . This sensitivity is a consequence mostly of variation in the viscosity of the propellant, μ . The hydraulic impedance of the source, Eq. (5), is directly proportional to the dynamic viscosity, so the current is as well by Eq. (4). The viscosity is highly sensitive to temperature: from 289.8 to 292.5 K—which represents the IQR over our prior on temperature— μ changes from 0.0419 to 0.0375 Pa-s, a decrease of 11% and comparable to the variance in our predictions induced by the propellant properties.

The dependence of emission current on temperature has been documented before, including to tie it to changes in hydraulic impedance.^{51,57,58} These studies also demonstrated that the temperature can influence other aspects of the spray, including the charge to mass ratio of the plume.⁵⁹ Even without modeling these additional dependencies, our analysis indicates that an uncertainty of even a couple degrees can have major implications for array performance, constituting the dominant error mode above 1300 V. This sensitivity suggests that measuring and reporting the operating temperature of these sources—or indeed controlling it—should be standard in emitter experiments.

The next most major source of uncertainty is the parameters we learned in our main inference, θ^* , with a magnitude about half that of the propellant properties. This underscores the importance of treating these parameters as probabilistic. Above about 1300 V where the geometry error ceases to dominate, failing to account for uncertainty in these model parameters—e.g., if we had done a least-squares fit to the data or considered only the MAP, would underestimate our uncertainty, as evidenced by the difference between the error curves for θ and $\hat{\theta}$. This effect is of increasing importance as we refine our understanding of the other parameters. If we were to measure the temperature more reliably, say, then uncertainty in these parameters may begin to dominate analysis.

Our sensitivity analysis informs what additional experiments we might do to reduce our uncertainty. If we were intending to operate the thruster above 1300 V, these results suggest that refining our estimate to the propellant temperature could help reduce uncertainty. Below 1300 V though, additional emitter geometry measurements to reduce uncertainty in the underlying geometry distributions could be more effective. Following our earlier discussion, however, the effect of additional geometry measurements is weakened by the underlying reliability of the manufacturing process. In the limit that the manufacturing process were perfectly reliable, producing emitters with no variance in geometry, it would only be necessary to measure a single emitter to infer the properties of the entire population.¹³ Thus, at some level, only improving the underlying manufacturing process would help limit variance. The interplay of these effects is a function of the scale of an electrospray array thruster—thrusters with fewer emitters than the MEAT-1.2 have array currents more sensitive to manufacturing tolerances, because they have a smaller sample to average over. This trade represents a key design point in ensuring reliable systems.

C. Implications for Robust Design

A primary motivation in conducting probabilistic analyses like that in this work is to enable robust design. That is, we have a system—in our case an electrospray array thruster—whose performance is uncertain



stemming from multiple sources of error, like those examined in Sec. V.B. We can then formulate a design problem over this system to optimize objectives related to this uncertainty. The exemplar is to ensure mission requirements; if we had some minimum specifications the system needed for a successful mission—thrust, specific impulse, device lifetime, etc.—we would seek a design that maximizes the probability these specifications are met, subject to our uncertainty.^{60–63}

Similarly, we might aim to optimize some aspect of the thruster’s performance subject to a minimum performance constraint. As an example, we consider a toy optimization problem to pick an operating condition for our array. We suppose that due to mission constraints, the thruster must be able to provide at least some thrust setpoint τ_{req} , where—ignoring inefficiencies in the system such as the divergence of the beam—we idealize the thrust produced by the device τ as

$$\tau = I \sqrt{\frac{2V}{\xi}}. \quad (55)$$

We further suppose that we wish to minimize the necessary operating voltage of the thruster—e.g., to reduce demands on electrical isolation for the spacecraft and minimize the size of the needed power supply. We can define this as an optimization problem of the form

$$V^*(C, \tau_{req}) = \min_V \{V : P(\tau \geq \tau_{req}) \geq C\}, \quad (56)$$

where V^* is the optimum maximum operating voltage for some confidence level $C \in [0, 1]$ with which we demand the thrust requirement be met.

We perform this optimization over our posterior predictions for several values of C and τ_{req} , the results of which we plot in Fig. 16. The solid lines show the thrust reliability, $P(\tau \geq \tau_{req})$, for the different values of




Figure 16. Probability of the thruster exceeding a thrust specification, τ_{req} as a function of voltage, for various τ_{req} ; the dashed lines illustrate the minimum voltage satisfying various confidence thresholds

τ_{req} , while the dashed lines highlight different choices for C and the corresponding V^* . The thrust reliability adopts sigmoid shapes resembling an error function.

The $C = 0.5$ threshold at the center of the plot illustrates the design point that would be chosen if variance in thrust were not accounted for (i.e., if we only designed to the median prediction of Fig. 12). The figure indicates that as the required confidence in meeting the specification is increased, additional margin must be incorporated into the operating voltage, relative to the 50% confidence level. For $\tau_{req} = 500 \mu\text{N}$, for example, a margin of only 30 V (from 1610 V to 1640 V), is needed to hedge to 75% confidence, but a 130 V

margin (from 1610 V to 1740 V) must be incorporated to meet the requirement with 99% confidence. The needed margin also increases with the thrust requirement, because the absolute variance in array current increases with voltage, even though the relative error (i.e., normalized to some reference) decreases per Fig. 15. Qualitatively, these results are intuitive and perhaps even unremarkable; the power in the approach pursued here is that quantitative predictions for margin and reliability can be made, and they flow down explicitly via statistical inference.

Another area of concern in electrospray array thrusters is the decrease in device lifetimes associated with variance in emitter current, a consequence of a heightened probability that any one emitter induces a life-ending short.²³ As an illustration of this effect, we assume that the lifetime of an array of N_e emitters is inversely proportional to the maximum current shed by any emitter—i.e., it sprays the most and so will short out first—such that a proxy for the total throughput of an array is the array current divided by the maximum emitter current, within a scaling constant. We plot the median predicted throughput as judged by this proxy as a function of voltage in Fig. 17. For simplicity, we consider predictions made only varying the




Figure 17. Median predicted throughput over varying emitter geometry, $\dot{\theta} = \cup_j \{R_{cj}, \alpha_j, r_{bj}, \Delta_j\}$

emitter geometry, $\dot{\theta} = \cup_j \{R_{cj}, \alpha_j, r_{bj}, \Delta_j\}$, since this is the primary driver of differences between emitters.

Our study predicts the maximum throughput of the device is achieved at the maximum operating voltage. This follows noting the finding in Sec. V.A that the array current is positively inflected over voltage as a result of an increasing proportion of emitters activating. The total current grows more quickly than does the current of any one emitter—and hence the maximum—so the throughput increases monotonically. Once most emitters have become active, however, the growth in total and maximum current become comparable, and so the predicted throughput plateaus.

For a less idealized analysis, this scaling will be balanced by divergence in the beam. As voltage is increased and emitters form additional emission sites off center, these additional beamlets can increase the spread of the beam, increasing interception to the extractor.^{27,45} Similarly, duller emitters, which preferentially activate at higher voltages, may also have heightened divergence or are otherwise more deeply recessed relative to the extractor (see the anticorrelation of R_c and h in Fig. 4) and so also increase interception. These additional dependencies will tend to decrease expected throughput at higher voltages and likely induce a maximum at lower operating voltages, where emission is not as stressed.

Given finite manufacturing tolerances like those inferred in Sec. IV.B, designing to maximize throughput becomes a problem of selecting a nominal emitter design (i.e., a target base radius, cone angle, emitter height, etc.) that best balances these effects under uncertainty. Similarly, broader goals of ensuring reliability in the thruster are predicated on treating the performance as something that cannot be controlled or predicted

exactly, as something uncertain. By applying probabilistic tools of inference and learning like those employed here, one can rigorously derive uncertainty in performance from underlying error sources, driving development with confidence.

VI. Conclusion

The variability of individual emitters within an electrospray array induces emergent phenomena which can have deleterious effects on thruster performance (lifetime, current yield, etc.). This and other sources of uncertainty in the operation of the device must be accounted for to make design decisions with confidence—to ensure reliability. To that end, in this study we interrogated the effect of this emitter nonuniformity and quantified the effects of different sources of uncertainty by performing an inference problem to train a multi-site emission model—predicting individual emitter behavior—on data taken on the MEAT-1.2—an array of over 6000 porous conical emitters.

We did so by adopting a Bayesian formalism to describe our state of knowledge in the inputs and outputs of the current emission model (the temperature of the propellant, individual emitter geometries, predicted current, etc.) as probability distributions. Most of these parameters were drawn from direct measurements, datasheets, and similar sources, but others represent scaling parameters that cannot be assigned *a priori*. Consequently, we inferred these missing scaling parameters from experiments measuring the emission current of the array as a function of operating voltage. To capture the uncertainty in individual emitter geometries, we measured the geometry of 143 emitters in the array via surface profilometry. Concluding the emitter chip was composed of three distinct populations of emitters (defined by cutting tool changes while machining), we then inferred the distribution in emitter geometry for each of these populations to feed back into our main learning problem. The posterior distribution over the model parameters represented a space of more than 20000 dimensions (due to individual emitter geometry), so to facilitate analysis we adopted an approximately pseudomarginal form for the inference, which restricted inference to only the parameters we did not have strong prior understanding of by integrating out the effect of the other parameters.

We drew samples from the distribution over the parameters learned in this way to characterize our state of knowledge. With regard to the emitter geometry, these results indicated that the emitter base radius and half angle did not vary substantially within each population (interquartile ranges of $< 2 \mu\text{m}$ and $< 8 \text{ mrad}$, respectively). However, we observed that these parameters varied more substantially between populations (order $10 \mu\text{m}$ and 20 mrad) between populations. The tip radius of the emitters was more consistent across populations (approximately $25 \mu\text{m}$ mean), but varied more substantially within each population (interquartile range of order $20 \mu\text{m}$), confirming that this parameter is comparatively poorly controlled during the manufacturing process. With regard to the scaling parameters of the model, we identified that those related to the formation of menisci on the surface of the emitter were similar to values inferred when training the model on data taken on a single emitter—indicating that menisci can be as large as $13 \mu\text{m}$ in diameter and that their relative number density as a function of diameter scales inversely as the $\frac{3}{2}$ power. However, we found that parameters dictating emission current from individual menisci differed from previous studies, and concluded this discrepancy derived from the source in prior studies having a current trace that increases smoothly above their onset voltage while the source in our study has emitters that experience a jump discontinuity in current at onset. We posited this dichotomy in behavior was a nonlinear effect not captured by our linearized current model.

Using these parameter samples to then make probabilistic predictions over emitter performance, we confirmed that the model was able to reconstruct the behavior of the array, a positive but not conclusive demonstration that it is capturing the physics. Our predictions retrieved the common observation that while individual emitter currents are often linear functions of voltage, the emission current over an entire array tends to be a positively inflected curve that eventually saturates to be nearly linear. By inspecting predictions for a single realization of the emitter geometry, we showed this is a consequence of variations in onset voltage of emitters across the array: an increasingly large number of active emitters each yield an increasingly large amount of current, producing the inflection. At sufficiently high voltages such that a majority of emitters have activated, then, this growth saturates to the linear dependence of an individual emitter.

To establish which sources of uncertainty were dominant in our analysis, we made additional predictions where we considered only one subset of the model parameters as uncertain at a time, to form a sensitivity study. Uncertainty in emitter geometry tended to be the dominant source of error in the predictions at



very low voltages, which we attributed to the fact that only very few emitters, if any, are active at low voltages. Consequently, whether or not any individual emitter is active has a larger effect, in a relative sense, on the total prediction over the array. At higher voltages where many more emitters are active, this uncertainty diminishes however, essentially a consequence of the law of large numbers: increasingly many samples over a distribution—the number of active emitters—provide an increasingly precise estimate to its average—the relative variance in the array current. Thus, for increasing voltage the dominant source of error in our analysis was uncertainty in the properties of the propellant. Though our uncertainty in temperature may seem small (order 2 K), the change in viscosity of the propellant over this range is order 10%, which directly propagates into the current measurement since current is predicted to be inversely proportional to the viscosity. As such, we concluded that providing greater control over propellant temperature could substantially reduce uncertainty in the performance of a system, echoing experiments which have sought to characterize this dependency.

Finally, we leveraged these probabilistic predictions to perform a toy robust design optimization, where we considered minimizing the operating voltage of the source subject to the constraint that it must produce some minimum thrust with confidence. We verified that as greater confidence in meeting the requirement was sought, greater voltage margin was necessary (since thrust increases with voltage). Additionally, increasing the requirement on thrust also tended to demand greater voltage margin because uncertainty in the prediction scales proportionally to the average prediction. As a result of these effects, while only a voltage margin of only 30 V (relative to the average prediction) was necessary to meet a thrust requirement of 100 μ N with 75% confidence, a margin of 50 V was necessary to meet a requirement of 1 mN with the same confidence, and a margin of 200 V to meet it with 99% confidence.

These three main discussion points—individual emitter behavior, model sensitivity, and robust design—constitute a primary motivation for the present study. Deconvolving the behavior of individual emitters from the performance of an entire array is experimentally challenging, and so modeling it like we have here provides a means to interpret what might be going on at the single source level, directly informed from array dynamics. Similarly, examining which sources of uncertainty are dominant in our analysis motivates improvements in manufacturing methodologies (e.g., to tighten fabrication tolerances) and operations (e.g., to more thoroughly control propellant temperature) that we might make to minimize uncertainty and increase reliability. Finally, quantifying uncertainty in our predictions is the only means by which we can make design decisions with confidence, determining what changes may be necessary to meet requirements more robustly. In aggregate, then, the present study represents a step forward in ensuring performant electrospray array thruster systems.

Data Availability

The data that support the findings of this study and the codes used to generate and analyze them are available in an external repository, Ref. 40.

Acknowledgments

This work was supported by a NASA Space Technology Graduate Research Opportunity (grant number 80NSSC21K1247). This research was supported in part through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor. This work was performed in part at the University of Michigan Lurie Nanofabrication Facility.

References

- ¹Lozano, 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, pp. 842–849.
- ²Zeleny, J., “The Electrical Discharge from Liquid Points, and a Hydrostatic Method of Measuring the Electric Intensity at Their Surfaces,” *Phys. Rev.*, Vol. 3, Feb 1914, pp. 69–91.
- ³Zeleny, J., “Instability of Electrified Liquid Surfaces,” *Phys. Rev.*, Vol. 10, Jul 1917, pp. 1–6.
- ⁴Taylor, G. I., “Disintegration of water drops in an electric field,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, Vol. 280, No. 1382, 1964, pp. 383–397.
- ⁵Krejci, D., Mier-Hicks, F., Thomas, R., Haag, T., and Lozano, P., “Emission Characteristics of Passively Fed Electrospray Microthrusters with Propellant Reservoirs,” *Journal of Spacecraft and Rockets*, Vol. 54, No. 2, 2017, pp. 447–458.



- ⁶Ziener, 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," *35th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2017, p. 578.
- ⁷Mair, G., "Emission from liquid metal ion sources," *Nuclear Instruments and Methods*, Vol. 172, No. 3, 1980, pp. 567–576.
- ⁸Coffman, 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," *Phys. Rev. E*, Vol. 99, Jun 2019, pp. 063108.
- ⁹Goebel, D. M., Katz, I., and Mikellides, I. G., *Fundamentals of electric propulsion*, John Wiley & Sons, 2023.
- ¹⁰Natisin, 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 fully conventionally machined, high-performance porous-media electrospray thruster," *Journal of Micromechanics and Microengineering*, Vol. 30, No. 11, sep 2020, pp. 115021.
- ¹¹Ma, C., Bull, T., and Ryan, C. N., "Plume Composition Measurements of a High-Emission-Density Electrospray Thruster," *Journal of Propulsion and Power*, Vol. 37, No. 6, 2021, pp. 816–831.
- ¹²Cisquella-Serra, A., Galobardes-Esteban, M., and Gamero-Castaño, M., "Scalable Microfabrication of Multi-Emitter Arrays in Silicon for a Compact Microfluidic Electrospray Propulsion System," *ACS Applied Materials & Interfaces*, Vol. 14, No. 38, 2022, pp. 43527–43537, PMID: 36112012.
- ¹³Whittaker, C. B. and Jorns, B. A., "Statistical Characterization of Emitter Fabrication over Electrospray Array Thruster," *Journal of Propulsion and Power*, Vol. 40, No. 6, 2024, pp. 896–904.
- ¹⁴Wright, P. L., Uchizono, N. M., Collins, A. L., and Wirz, R. E., "Characterization of a Porous Tungsten Electrospray Emitter," *37th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2022, p. 232.
- ¹⁵Guerra-Garcia, C., Krejci, D., and Lozano, P., "Spatial uniformity of the current emitted by an array of passively fed electrospray porous emitters," *Journal of Physics D: Applied Physics*, Vol. 49, No. 11, feb 2016, pp. 115503.
- ¹⁶Fedkiw, T., Wood, Z. D., and Demmons, N. R., *Improved Computed Tomography Current Mapping of Electrospray Thrusters*, American Institute of Aeronautics and Astronautics, 2021, p. 3389.
- ¹⁷Natisin, M. R., Zamora, H. L., Holley, Z. A., Ivan Arnold, N., McGehee, W. A., Holmes, M. R., and Eckhardt, D., "Efficiency Mechanisms in Porous-Media Electrospray Thrusters," *Journal of Propulsion and Power*, Vol. 37, No. 5, 2021, pp. 650–659.
- ¹⁸Perez-Martinez, C. S. and Lozano, P. C., "Ion field-evaporation from ionic liquids infusing carbon xerogel microtips," *Applied Physics Letters*, Vol. 107, No. 4, 07 2015, pp. 043501.
- ¹⁹Coffman, C., Martínez-Sánchez, M., Higuera, F. J., 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, 12 2016, pp. 231602.
- ²⁰Schroeder, M., Gallud, X., Petro, E., Jia-Richards, O., and Lozano, P. C., "Inferring electrospray emission characteristics from molecular dynamics and simulated retarding potential analysis," *Journal of Applied Physics*, Vol. 133, No. 17, 05 2023, pp. 173303.
- ²¹Thuppul, A., Wright, P., and Wirz, R. E., *Lifetime Considerations and Estimation for Electrospray Thrusters*, American Institute of Aeronautics and Astronautics, 2018, p. 4652.
- ²²Brikner, N. A., *On the Identification and Mitigation of Life-limiting Mechanisms of Ionic Liquid Ion Sources Envisaged for Propulsion of Microspacecraft*, Ph.D. thesis, Massachusetts Institute of Technology, 2015.
- ²³Jorns, B. A., Gorodetsky, A., Lasky, I., Kimber, A., Dahl, P., St. Peter, B., and Dressler, R., "Uncertainty Quantification of Electrospray Thruster Array Lifetime," *36th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, September 2019, p. 317.
- ²⁴Wright, P. and Wirz, R. E., *Transient Flow in Porous Electrospray Emitters*, American Institute of Aeronautics and Astronautics, 2021, p. 3437.
- ²⁵Villegas-Prados, D., Kazadi, J., Fajardo, P., and Wijnen, M., "Emission characterization of externally wetted electrospray thrusters using computed tomography," *38th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, September 2024, p. 205.
- ²⁶Jia-Richards, O. and Lozano, P. C., "Ensemble Kalman Update for Inference of Spatial Uniformity of Emission across an Electrospray Array," *2022 IEEE Aerospace Conference (AERO)*, 2022, pp. 1–8.
- ²⁷Dressler, R. A., St. Peter, B., Chiu, Y.-H., and Fedkiw, T., "Multiple Emission Sites on Porous Glass Electrospray Propulsion Emitters Using Dielectric Propellants," *Journal of Propulsion and Power*, Vol. 38, No. 5, 2022, pp. 809–821.
- ²⁸Petro, E. M., Gallud, X., Hampl, S. K., Schroeder, M., Geiger, C., and Lozano, P. C., "Multiscale modeling of electrospray ion emission," *Journal of Applied Physics*, Vol. 131, No. 19, 05 2022, pp. 193301.
- ²⁹Asher, J., Huang, Z., Cui, C., and Wang, J., "Multi-scale modeling of ionic electrospray emission," *Journal of Applied Physics*, Vol. 131, No. 1, 01 2022, pp. 014902.
- ³⁰Gallud, X. and Lozano, P. C., "The emission properties, structure and stability of ionic liquid menisci undergoing electrically assisted ion evaporation," *Journal of Fluid Mechanics*, Vol. 933, 2022, pp. A43.
- ³¹COFFMAN, C. S. and Lozano, P. C., *On the Manufacturing and Emission Characteristics of Dielectric Electrospray Sources*, American Institute of Aeronautics and Astronautics, 2013, p. 4035.
- ³²Courtney, D. G. and Shea, H., "Influences of porous reservoir Laplace pressure on emissions from passively fed ionic liquid electrospray sources," *Applied Physics Letters*, Vol. 107, No. 10, 09 2015, pp. 103504.
- ³³Courtney, D. G., Li, H. Q., and Lozano, P., "Emission measurements from planar arrays of porous ionic liquid ion sources," *Journal of Physics D: Applied Physics*, Vol. 45, No. 48, nov 2012, pp. 485203.
- ³⁴Whittaker, C. B. and Jorns, B. A., "Modeling Multi-Site Emission in Porous Electrosprays Resulting from Variable Electric Field and Meniscus Size," *Journal of Applied Physics*, Vol. 134, No. 8, 08 2023, pp. 083301.
- ³⁵Whittaker, C. B., Jorns, B. A., Arestie, S. M., and Marrese-Reading, C. M., "Characterization of a High Power Electrospray Array Thruster," *38th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2024, p. 730.



- ³⁶Higuera, F. J., “Model of the meniscus of an ionic-liquid ion source,” *Phys. Rev. E*, Vol. 77, Feb 2008, pp. 026308.
- ³⁷Whittaker, C. B., *Designing Porous Electrospray Array Thrusters Under Uncertainty*, Ph.D. thesis, University of Michigan, 2025.
- ³⁸Courtney, D. G., *Ionic Liquid Ion Source Emitter Arrays Fabricated on Bulk Porous Substrates for Spacecraft Propulsion*, Ph.D. thesis, Massachusetts Institute of Technology, 2011.
- ³⁹Whittaker, C. B. and Jorns, B. A., “A Self-Healing Extractor for Electrospray Arrays,” *37th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2022, p. 206.
- ⁴⁰Whittaker, C. B. and Jorns, B. A., “Emitter Model Inference from Electrospray Array Thruster Tests Dataset,” Deep Blue Data, 2025, <https://doi.org/10.7302/92gy-3n37>.
- ⁴¹Peter, B. S., Dressler, R. A., Chiu, Y.-h., and Fedkiw, T., “Electrospray Propulsion Engineering Toolkit (ESPET),” *Aerospace*, Vol. 7, No. 7, 2020.
- ⁴²SSI, Inc., “ESPET Web App,” <http://espet.spectral.com/>, 2024, Accessed: 17 Oct 2024.
- ⁴³Lewandowski, D., Kurowicka, D., and Joe, H., “Generating random correlation matrices based on vines and extended onion method,” *Journal of Multivariate Analysis*, Vol. 100, No. 9, 2009, pp. 1989–2001.
- ⁴⁴Wright, P. L. and Wirz, R. E., “Multiplexed electrospray emission on a porous wedge,” *Physics of Fluids*, Vol. 33, No. 1, 01 2021, pp. 012003.
- ⁴⁵Schroeder, M., Cidoncha, X. G., Bruno, A., Jia-Richards, O., and Lozano, P., *Angular Properties of Ionic Liquid Electrospray Emitters*, American Institute of Aeronautics and Astronautics, 2023, p. 1408.
- ⁴⁶Andrieu, C. and Roberts, G. O., “The Pseudo-Marginal Approach for Efficient Monte Carlo Computations,” *The Annals of Statistics*, Vol. 37, No. 2, 2009, pp. 697–725.
- ⁴⁷Hughes, J. F., Van Dam, A., McGuire, M., Sklar, D. F., Foley, J. D., Feiner, S. K., and Akeley, K., *Compute Graphics: Principles and Practice*, chap. Splines and Subdivision Curves, Addison-Wesley, Upper Saddle River, NJ, 3rd ed., 2014.
- ⁴⁸Mira, A., “On Metropolis-Hastings algorithms with delayed rejection,” *Metron*, Vol. 59, 04 2001.
- ⁴⁹Eckels, J. D., Whittaker, C. B., Jorns, B. A., Gorodetsky, A. A., St.Peter, B., and Dressler, R. A., “Simulation-based surrogate methodology of electric field for electrospray emitter geometry design and uncertainty quantification,” *37th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2022, p. 271.
- ⁵⁰McGehee, W., Arnold, N., Natisin, M., Holley, Z., and Eckhardt, D., “The Single Emitter Air Force Electrospray Thruster (SEAFET) Program: Goals and Initial Results,” *37th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2022, p. 600.
- ⁵¹Lozano, P. and Martínez-Sánchez, M., “Ionic liquid ion sources: characterization of externally wetted emitters,” *Journal of Colloid and Interface Science*, Vol. 282, No. 2, 2005, pp. 415–421.
- ⁵²Miller, C. and Lozano, P. C., *Measurement of the Fragmentation Rates of Solvated Ions in Ion Electrospray Thrusters*, American Institute of Aeronautics and Astronautics, 2016, p. 4551.
- ⁵³Castro, S. and Fernández de la Mora, J., “Effect of tip curvature on ionic emissions from Taylor cones of ionic liquids from externally wetted tungsten tips,” *Journal of Applied Physics*, Vol. 105, No. 3, 02 2009, pp. 034903.
- ⁵⁴Quraishi, A., Dworski, S., Ryan, C. N., Ferreri, A., Vincent, G., Croos, A., Wild, A., Garbayo, A., Vozarova, M., and Neubauer, E., *Designing and Commercialization of Porous Emitter Electrospray Thruster for Space Applications*, American Institute of Aeronautics and Astronautics, 2023, p. 0262.
- ⁵⁵Petro, E., Bruno, A., Lozano, P., Perna, L. E., and Freeman, D., *Characterization of the TILE Electrospray Emitters*, American Institute of Aeronautics and Astronautics, 2020, p. 3612.
- ⁵⁶Sun, W., Guo, Y., Wu, Z., Du, Z., Wang, N., and Sun, Z., “Design and Performance of an Electrowetting Ionic Liquid Electrospray Thruster Prototype,” *Journal of Propulsion and Power*, Vol. 40, No. 3, 2024, pp. 380–387.
- ⁵⁷Romero-Sanz, I., de Carcer, I. A., and de la Mora, J. F., “Ionic Propulsion Based on Heated Taylor Cones of Ionic Liquids,” *Journal of Propulsion and Power*, Vol. 21, No. 2, 2005, pp. 239–242.
- ⁵⁸Fujiwara, Y., “Effects of the temperature of a protic ionic liquid on ion beam production by vacuum electrospray,” *Journal of Vacuum Science & Technology B*, Vol. 41, No. 2, 03 2023, pp. 024202.
- ⁵⁹Bruno, A. R., Corrado, M. N., and Lozano, P. C., “Temperature-Dependent Performance of Electrospray Thrusters with ASCENT Monopropellant,” *38th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2024, p. 782.
- ⁶⁰Marrese-Reading, C., Arestie, S., Ziemer, J. K., Wirz, R., Collins, A., Uchizono, N., no, M. G.-C., and Demmons, N., “Electrospray Thruster Lifetime Modeling with Uncertainty Quantification and Experimental Test Validation,” *37th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, 2022, p. 234.
- ⁶¹Ziemer, J. K., Marrese-Reading, C., Arestie, S., Demmons, N. R., Wirz, R. E., Collins, A., and Gamero, M., *Progress on Developing LISA Microthruster Technology*, American Institute of Aeronautics and Astronautics, 2020, p. 3609.
- ⁶²Gorodetsky, A., Whittaker, C. B., Szulman, A., and Jorns, B., *Robust Design of an Electrospray Emitters*, American Institute of Aeronautics and Astronautics, 2021, p. 3422.
- ⁶³Ben-Tal, A., El Ghaoui, L., and Nemirovski, A., *Robust Optimization*, Princeton University Press, Princeton, NJ, second edition ed., 2009.

