Reproducing a physics-informed simplified model of a Hall thruster: a well-posed formulation, a reference solution, and what a small residual loss does not show

TechUPGRADE Digital Twin Platform

TechUPGRADE technical report TU-TR-2026-001 (tenant `hall`)

2026-10-31

Preprint: 6/8 requirements met, 1 partial

Abstract

We reimplement the physics-informed neural network (PINN) of Burrows et al. [1] for the steady, one-dimensional channel of a xenon Hall thruster, and add what that work did not have: a conventional reference solution of the same equations. Solved as a boundary-value problem with the discharge current as an unknown, the equations as printed, with the paper's boundary conditions, have no exact solution: the anode conditions are overdetermined by one, and the electron temperature, carried by an electron flux that flows towards the anode, is fixed at the wrong end. We show that the paper's fixed normalisation lets a network reach a small loss (1.1×10^{-6}) while its profiles do not satisfy the equations (discharge current 0.13 A against 6.91 A). With four changes and standard values (5 mg/s, charge-exchange collisions), the problem is well posed: at 100 V the reference gives 6.91 A, 691 W, 44.0 mN and 82 % propellant utilisation, and the paper's network, trained on residuals scaled by their own terms, reproduces it to within 1.3 % in every field in about 38 s on four CPU cores. The whole study runs as a tenant of the TechUPGRADE platform, where every result in this report can be recomputed.

Keywords: physics-informed neural networks, Hall thruster, quasineutral plasma, boundary-value problem, reproducibility

1 Introduction

Physics-informed neural networks [5] minimise the residuals of differential equations at collocation points and are attractive as fast, mesh-free surrogates. Burrows et al. [1] apply one to a simplified, steady, one-dimensional model of a Hall thruster derived from Roy and Pandey [2, 3], report a residual loss of order 10^{-8} after 4 min 32 s of training, and a discharge current of 6.5859 A at 100 V. The paper's results are compared qualitatively with the literature; no solution of the same equations by another method is given. This report reproduces the approach on the TechUPGRADE digital twin platform and adds such a solution, so the network can be judged against an answer rather than against its own loss.

2 Contributions

- An analysis of the paper's boundary conditions: as printed, the system has no exact solution (the anode is overdetermined by one condition, and the electron temperature is fixed at the wrong end).
- A demonstration that the paper's fixed normalisation lets a network reach a small loss while its fields do not solve the equations, and a term-relative scaling that removes this.
- A well-posed variant of the model with two independent references (collocation, and finite differences solved by Newton's method) that agree to 0.01 %, against which the PINN is measured over five seeds.
- A one-parameter comparison with the SPT-100's published operating point, a sensitivity study of the inputs the paper leaves open, and a parametric PINN tested at voltages it was not trained on.

3 The source model

The channel is 4 cm long with cross-section $A = 0.0045$ m^2; the radial magnetic field follows a Gaussian with a 200 G maximum (Eq. 5 of [1]). The unknowns are the ion density n_i, ion velocity V_i, electron temperature T_e (eV) and potential ϕ, with the electric field $E = -\phi'$ and the current density J_T as a trainable constant. Quasineutrality gives $n_e = n_i$; current conservation $V_e = V_i - J_T/(en_i)$; mass conservation $n_n = (\Gamma_m - n_i V_i)/V_n$ with a constant neutral velocity V_n. The four ODEs are

$$(n_i V_i)' = S = k_i(T_e)\, n_n n_e, \tag{1}$$

$$V_i V_i' = \tfrac{e}{m_i} E + \tfrac{m_e}{m_i}\nu_{ei}(V_e - V_i) - \nu_{in}(V_i - V_n) - \tfrac{S}{n_e}(V_i - V_n), \tag{2}$$

$$0 = -\tfrac{e}{m_e} E - \Omega\omega_{ce} V_e - \nu_{ei}(V_e - V_i) - \nu_{en}(V_e - V_n) - \tfrac{S}{n_e}(V_e - V_n), \tag{3}$$

$$\left(\tfrac{5}{2} n_e V_e T_e\right)' - n_e V_e \phi' = \sigma_I S\left(\tfrac{3}{2} T_e + \alpha E_I\right) + 3\tfrac{m_e}{m_i} n_e \nu_{ei}(T_i - T_e) + 3\tfrac{m_e}{m_n} n_e \nu_{en}(T_n - T_e), \tag{4}$$

with $\Omega = \omega_{ce}/(\nu_{ei} + \nu_{en} + \alpha_B \omega_{ce})$, k_i the cubic fit of Eq. 4, $\nu_{en} = 2.5 \times 10^{-13} n_n$ and the Coulomb frequency of Eq. 2 with $\ln\Lambda = 15$. The paper normalises with $\Gamma_{\text{ref}} = 10^{23}$ m^{-2}s^{-1} and $T_{\text{ref}} = 12.1$ eV, which give $V_{\text{ref}} = 2982$ m/s, $n_{\text{ref}} = 3.35 \times 10^{19}$ m^{-3} and $z_{\text{ref}} = 1.693$ mm.

Values not given, and printing errors. The Bohm coefficient α_B (we take $1/16$), T_i (0.3 eV) and the unit of z in Eq. 5 (read as z/L, which places the maximum at 3.2 cm, "upstream of the exit plane") are not stated. The mass flow is not stated either; it follows from $n_n(0) = n_{\text{ref}}$ and $V_n = 0.4 V_{\text{ref}}$ as 39 mg/s, about eight times that of an SPT-100. Eq. 17 drops the factor $(V_e - V_n)$ of the ionisation drag that Eqs. 6 and 40 contain; Eq. 47 drops the ionisation energy of Eq. 18; Eqs. 9 and 18 add the ionisation energy as a gain ($\sigma_I = +1$) where it is a loss ($\sigma_I = -1$); the trained "current density" is reported in amperes; and the inlet densities in the figures ($n_n = 2.0 \times 10^{19}$, $n_i = 0.35 \times 10^{18}$ m^{-3}) differ from the stated n_{ref} and $0.014\, n_{\text{ref}}$. We keep the factors and terms, and list every choice on the tenant's page.

4 Analysis of the boundary conditions

The anode is overdetermined. The paper imposes $n_i(0) = n_{i0}$, $V_i(0) = 0$, $T_e(0) = 3$ eV, $\phi(0) = 100$ V and $\phi(L) = 0$ for four first-order ODEs and one unknown J_T. With $V_i(0) = 0$ and n_{i0} finite, Eq. (1) forces $V_i'(0) = k_i n_{n0}$, and Eq. (2) then requires its right-hand side to vanish at $z = 0$. But Eq. (3) gives $E(0) = -\frac{m_e}{e}[\Omega\omega_{ce} + \nu_{ei} + \ldots]V_e(0) > 0$ for every $J_T > 0$ (since $V_e(0) = -J_T/(en_{i0}) < 0$), and every term of the right-hand side of Eq. (2) is then positive. No current satisfies the anode conditions exactly. The paper's "physics layer" enforces $V_i'(0)$ and $\phi'(0)$ so that the ion equations hold at $z = 0$; the electron momentum equation must then carry the error there. In our run of the paper's case its residual at $z = 0$ is 0.43 (Fig. 6).

The electron temperature is fixed at the wrong end. Without heat conduction Eq. (4) is first order in T_e and is carried by the electron flux $\Gamma_e = n_e V_e < 0$, which flows from the cathode to the anode. Its characteristic enters at the exit. Marched from the anode with $T_e(0) = 3$ eV, the temperature collapses for every current we tried.

A well-posed variant. We use four changes: ions enter at the neutrals' speed, $V_i(0) = V_n$; the electrons' temperature is given at the exit, $T_e(L) = 3$ eV; the ionisation energy is a loss; and $\nu_{in} = n_n \sigma_{cx} |V_i - V_n|$ with $\sigma_{cx} = 5.5 \times 10^{-19}$ m^2 (the paper's constant 1.8×10^6 s^{-1} is discussed in Section 6.5), with an SPT-100 class flow of 5 mg/s. These give four first-order ODEs, one parameter and five conditions.

5 Methods

Reference. The variant is solved by collocation (`scipy.integrate.solve_bvp` [6]) in logarithmic variables with J_T as a parameter, to a tolerance of 10^{-6}; one solve takes about 0.04 s. Putting finite differences of the solution back into the residuals gives a relative imbalance below 2×10^{-2} in every equation, limited by the differencing.

A second reference. The same equations, written as trapezoidal differences in every cell of a uniform grid, with the five boundary conditions and J_T as unknowns, solved together by Newton's method (`scipy.optimize.root`) from a generic starting guess, at $N = 100$, 200 and 400 cells. Neither method uses the other's solution.

PINN. As in [1]: one input $s = z/L$, four hidden layers of 32 tanh units with Glorot-uniform initialisation [7], four outputs, $J_T = \exp(\theta)$ trainable, Adam [8] (10^{-3}) followed by L-BFGS, in JAX [10] with optax, in double precision. The paper's physics layer is reproduced for its case ($n_i(0)$, $V_i(0) = 0$, $V_i'(0) = k_i n_{n0}$, $T_e(0)$, $\phi(0)$, $\phi(L) = 0$ and $\phi'(0)$ from Eq. (2) at $z = 0$); for the variant, the layer imposes its own conditions ($n_i = n_{i0} e^{sN_n}$, $V_i = V_n e^{sN_V}$, $T_e = T_e(L) e^{(1-s)N_T}$, $\phi = \phi_0(1-s) + \phi_0 s(1-s)N_\phi$). The loss is the mean of the four squared residuals at 2,000 fixed points, 6,000 + 6,000 iterations by default. Two normalisations are compared: the paper's fixed scales, and each residual divided by the sum of the magnitudes of its own terms at that point.

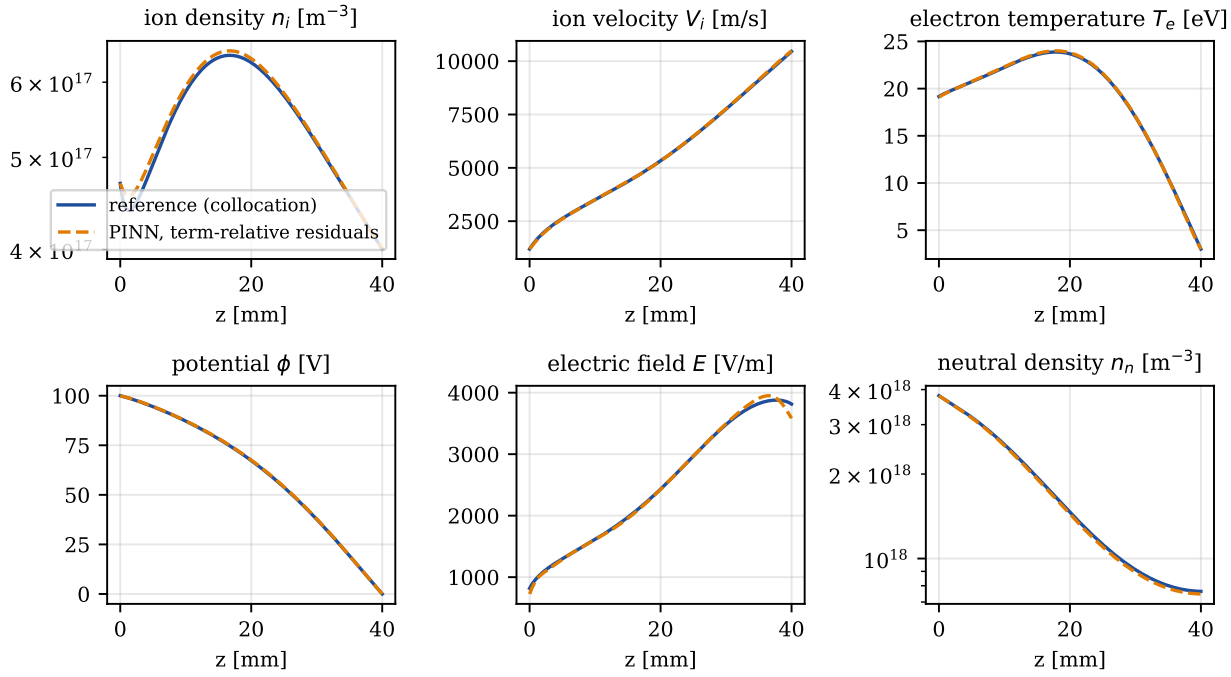


Figure 1: The well-posed variant at 100 V, 5 mg/s: the reference (solid) and the PINN with term-relative residuals, seed 0 (dashed).

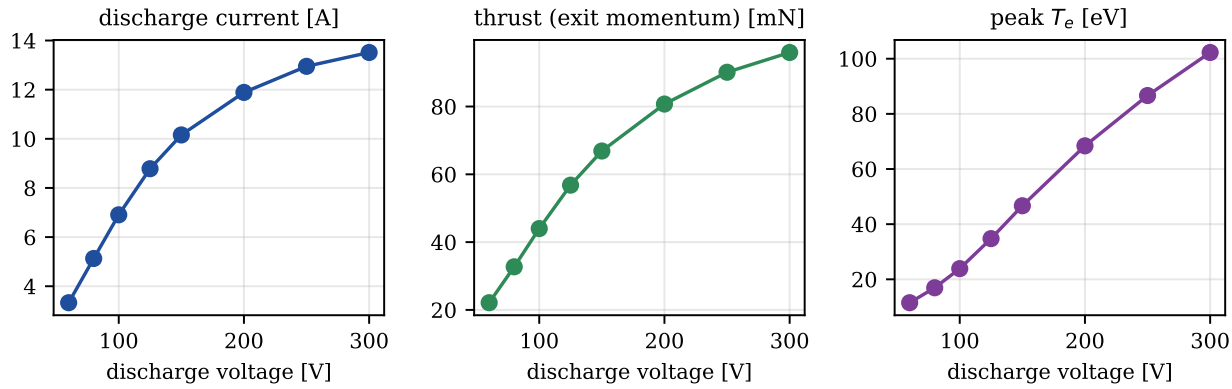


Figure 2: Reference solutions over the discharge voltage, 5 mg/s.

6 Results

6.1 The reference solution

At 100 V and 5 mg/s the current is 6.91 A ($J_T = 1535$ A m^{-2}), the power 691 W, the thrust from the exit momentum 44.0 mN and 82 % of the propellant leaves ionised. Ions leave at 10.5 km/s (the ideal for 100 V is 12.1 km/s); T_e peaks at 23.9 eV at $z = 17.8$ mm; n_i peaks at 6.40×10^{17} m^{-3} and n_n falls from 3.80×10^{18} to 7.64×10^{17} m^{-3} (Fig. 1, and as a moving picture in Fig. 3). The current grows from 3.33 A at 60 V to 13.51 A at 300 V; without wall losses the peak T_e rises to 102 eV at 300 V (Fig. 2), beyond what thrusters show, as the paper also leaves the walls out.

6.2 The PINN against the reference

Table 1 lists the runs. With term-relative residuals the network matches the reference in every field to within 1.3 % (relative L_2) and the current to within 0.7 %, in 38 s on four cores of an ARM (GB10) host. At the paper's budget (10,000 points, 10,000 + 30,000 iterations) the run takes 821 s; the paper reports 4 min 32 s on unstated hardware.

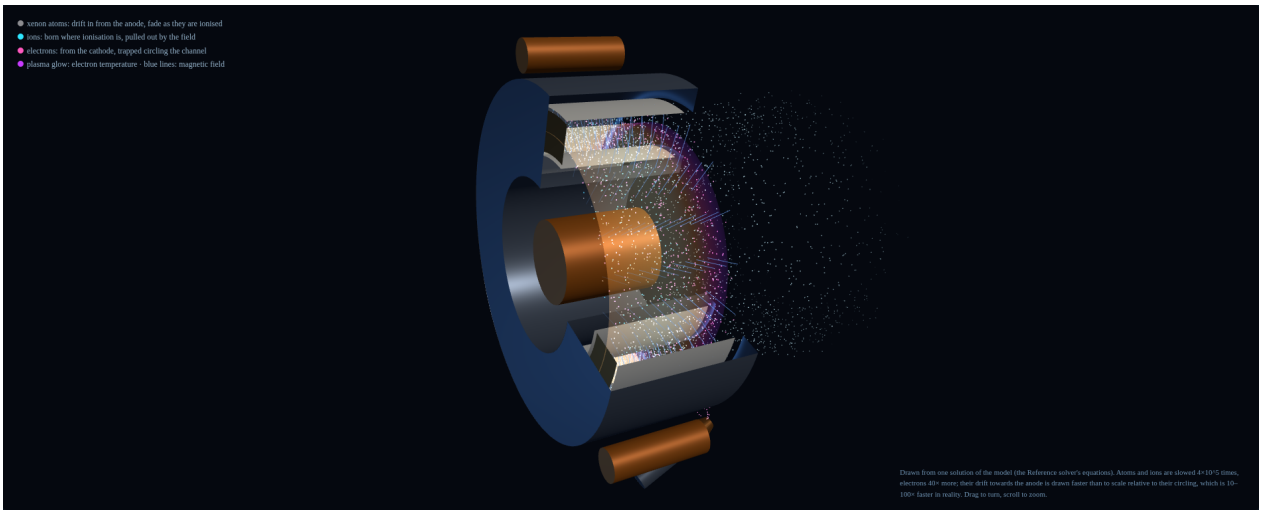


Figure 3: The tenant's Thruster view of the same solution: neutrals (grey) fade as they are ionised, ions (cyan) are born where $S(z)$ is and accelerate into the plume, electrons (pink) come from the cathode and drift round the channel, the plasma glows with $T_e(z)$ and the radial lines brighten with $B(z)$. Motion is slowed for display.

Table 1: PINN runs on the well-posed variant (reference $I_d = 6.906$ A). Field errors are relative L_2 differences from the reference.

run	I_d [A]	final loss	n_i [%]	V_i [%]	T_e [%]	ϕ [%]	time [s]
seed 0	6.956	1.61×10^{-4}	1.3	0.3	0.4	0.3	37
seed 1	6.921	9.71×10^{-5}	0.5	0.2	0.2	0.1	40
seed 2	6.925	1.30×10^{-4}	0.8	0.2	0.2	0.2	38
paper's budget	6.908	1.41×10^{-5}	0.1	0.0	0.0	0.0	821
paper's scales	0.135	1.13×10^{-6}	90.8	28.9	64.0	30.3	23

6.3 A small loss is not a solution

With the paper's fixed scales the plasma's own terms are about 10^{-3} of the reference scales (the ionisation source peaks at $\sim 5 \times 10^{-4}\, n_{\mathrm{ref}} V_{\mathrm{ref}} / z_{\mathrm{ref}}$). Every residual is then small whether or not the equation holds. On the variant, a network trained this way reached a loss of 1.1×10^{-6} with $I_d = 0.13$ A and field errors up to 90.8 % (Fig. 4). Integrating the true ODEs from its own anode values gives a different solution, so its profiles do not satisfy the equations. With term-relative residuals the same network and budget converge to the reference (Fig. 5).

6.4 The paper's case

With the paper's boundary conditions and fixed scales, the network reaches a loss of 1.5×10^{-3} with $I_d = 0.01$ A: an almost empty channel. The ion equations hold exactly at the anode, as the physics layer intends, and the electron momentum residual there is 0.43. With term-relative residuals the loss stays at 2.5×10^{-1} (Fig. 6). This agrees with the analysis of Section 4: there is no exact solution to converge to.

6.5 The constant ion–neutral frequency

With $\nu_{in} = 1.8 \times 10^{6}$ s^{-1} (the paper's value from [2]) the drag on the ions exceeds the field's pull: at 100 V and 5 mg/s they leave at 2.4 km/s instead of 10.5 km/s, and the current is 18.96 A (Fig. 7). With the paper's implied flow of 39 mg/s it is 159 A.

6.6 Two references

Table 2 compares the two references at three voltages. The finite-difference solution converges with the grid to the collocation solution: at $N = 400$ the currents agree to 0.01 % and every field to 0.09 % (largest relative difference).

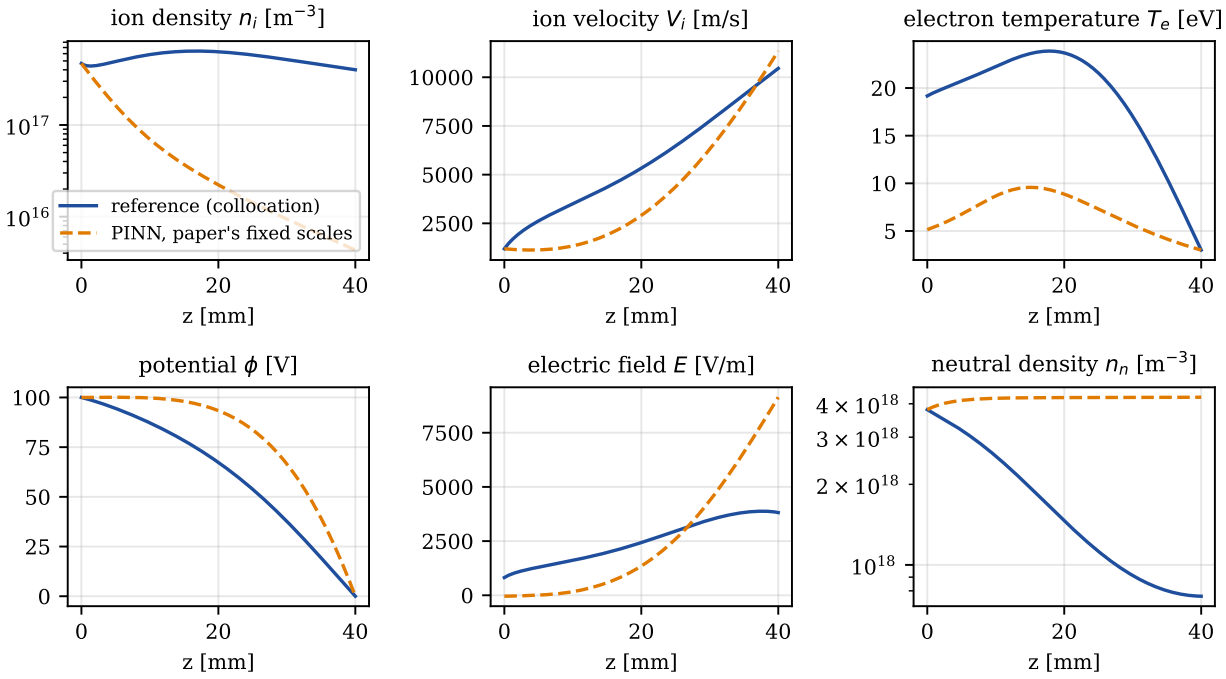


Figure 4: The same network trained on the paper's fixed scales: loss 1.1×10^{-6}, but not a solution of the equations.

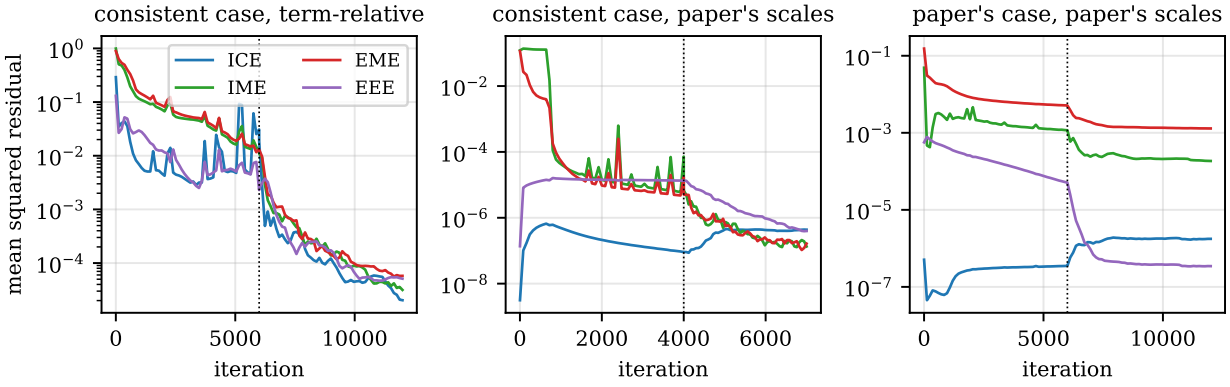


Figure 5: Mean squared residual per equation (ICE, IME, EME, EEE: Eqs. 1–4); the dotted line marks the switch from Adam to L-BFGS.

Table 2: Discharge current from the collocation reference and from finite differences with Newton's method.

voltage [V]	collocation [A]	FD, $N = 100$	FD, $N = 200$	FD, $N = 400$
100	6.906	6.895	6.903	6.905
200	11.888	10.227	11.883	11.887
300	13.513	11.077	13.505	13.511

6.7 Statistics over five seeds

With term-relative residuals, five seeds give $I_d = 6.939 \pm 0.023$ A (mean $\pm$ 95 % confidence interval; reference 6.906 A) and relative field errors of 0.9±0.5 % (n_i), 0.3±0.1 % (V_i), 0.3±0.2 % (T_e) and 0.2±0.1 % (ϕ), each a mean with its 95 % interval; Fig. 8 shows their distribution. On the paper's fixed scales, five seeds give $I_d = 0.13 \pm 0.09$ A, every one of them far from the reference.

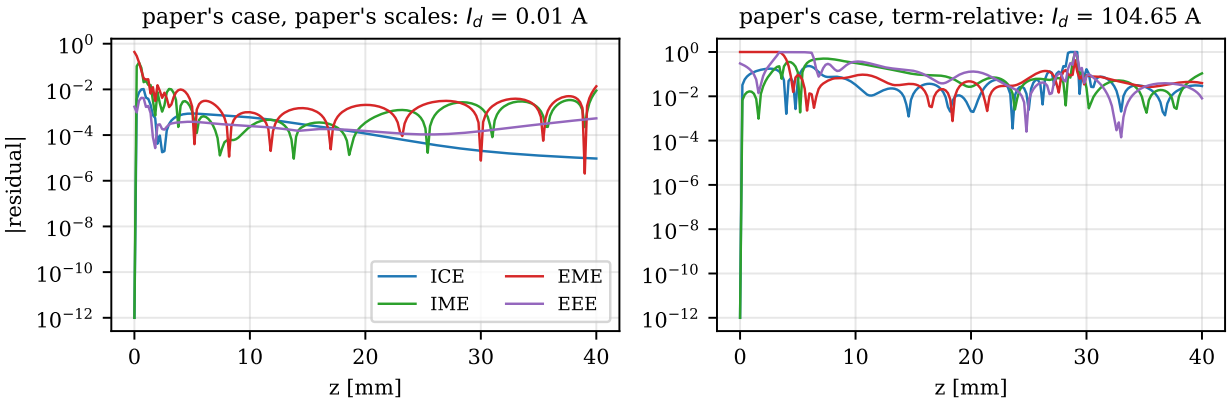


Figure 6: The paper's case: residuals along the channel at the end of training.

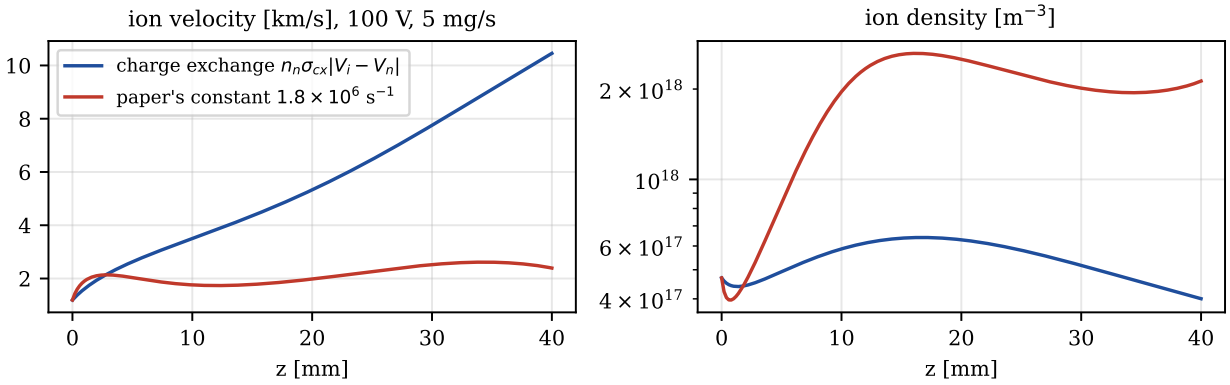


Figure 7: Charge-exchange collisions against the paper's constant frequency, 100 V, 5 mg/s.

6.8 Sensitivity

Fig. 9 varies the four inputs the paper does not give, one at a time at 100 V. The discharge current goes from 4.02 to 10.24 A (+155 %) as mass flow goes from 3 to 7 mg/s, from 2.98 to 14.23 A (+378 %) as α_B goes from 0.0312 to 0.125, from 9.58 to 5.26 A (-45 %) as B_{max} goes from 150 to 250 G, and from 6.79 to 7.13 A (+5 %) as $T_e(L)$ goes from 2 to 5 eV. The Bohm coefficient and the magnetic field, which set the electrons' cross-field transport, move the current most; the exit electron temperature moves it least.

6.9 Comparison with a real thruster

The SPT-100's published operating point is 300 V, 4.5 A, 80–83 mN and a specific impulse of 1,600 s [9], which implies a mass flow of 5.19 mg/s. With $\alpha_B = 1/16$ the model gives 14.2 A and 100 mN there: the anomalous transport is far too large. Fitting α_B alone to the published current gives $\alpha_B = 0.0139$ ($\approx 1/72$), and with it the model predicts 76.7 mN, -8 to -4 % from the published thrust. This is one fitted parameter against one other measured number, and our channel is the paper's simplified one, not the SPT-100's; the model also omits the walls, so its peak electron temperature (49 eV) is higher than such thrusters show.

6.10 Generalisation across voltages

One network with the voltage as a second input, trained on 80–150 V, was evaluated at voltages it had not seen (Fig. 10, five seeds). Inside the training range the ion velocity is within 1.7 % and the density within 8.2 % of the reference; outside it, at 170 and 200 V, the errors grow to 8.8 % and 21.1 %.

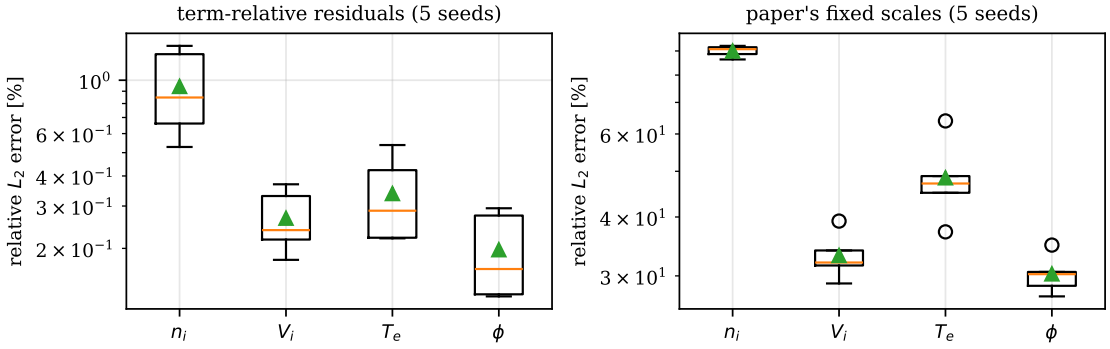


Figure 8: Field errors of five seeds: term-relative residuals (left) and the paper's fixed scales (right).

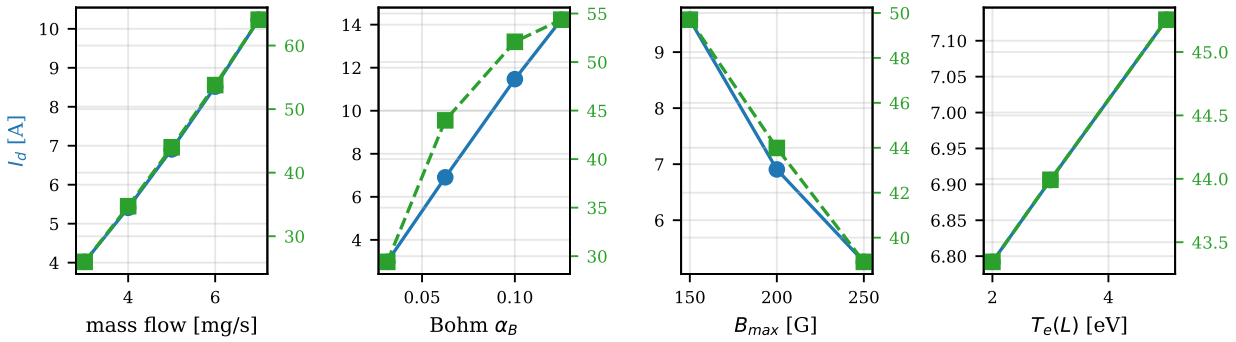


Figure 9: Reference solutions as each open input is varied (100 V, the others at their defaults).

6.11 Comparison with the paper

Table 3: The paper's reported values and ours. Theirs are read from its text and figures.

quantity	Burrows et al. [1]	this report (well-posed variant)
discharge current at 100 V	6.5859 A (trained)	6.906 A (reference), 6.96 A (PINN, seed 0)
mass flow	not stated (39 mg/s implied)	5 mg/s (assumed)
peak ion density	1.25×10^{18} m^{-3}	6.40×10^{17} m^{-3}
neutral density, inlet → exit	$2.0 \times 10^{19} \rightarrow 5 \times 10^{17}$ m^{-3}	$3.80 \times 10^{18} \rightarrow 7.64 \times 10^{17}$ m^{-3}
ion velocity	maximum at the exit	10.5 km/s at the exit
peak electron temperature	about 3/4 into the channel	23.9 eV at 17.8 mm (of 40)
residual loss	$\sim 10^{-8}$ (fixed scales)	1.6×10^{-4} (term-relative)
exact solution exists	not examined	no (paper's case); yes (variant)
training	10,000 pts, 40,000 it., 4 min 32 s	2,000 pts, 12,000 it., 38 s; at their budget 821 s
check against another method	none	collocation reference

The qualitative picture agrees: neutrals are depleted where ions are produced, ions accelerate towards the exit, and T_e peaks downstream. The quantitative values cannot be compared one to one, since the paper's case has no exact solution and its missing inputs had to be chosen.

7 Discussion and limitations

The model is the paper's simplified one: no wall losses, no anode sheath, no ion backflow, a constant neutral velocity and a prescribed magnetic field. These are why T_e grows too much with voltage. Real anode regions have a sonic point and ions moving backwards [4], and that is the physics an exact anode condition would need. Our values for α_B, the flow and σ_{cx} are standard choices, not fitted to a thruster. The paper's code is not available, so we worked from its text and kept to its architecture, physics layer and optimisers. The implementation was written with an AI coding assistant and has not been reviewed by an independent plasma physicist.

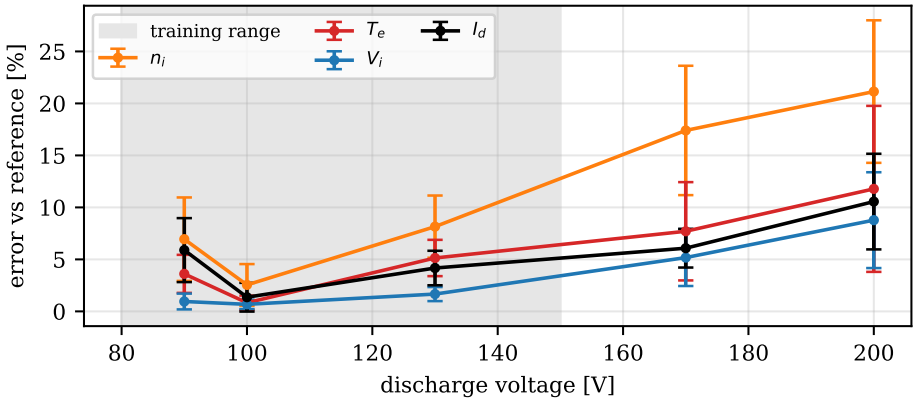


Figure 10: A parametric PINN trained on 80–150 V (shaded), against the reference at five voltages; mean and 95 % interval of five seeds.

8 Conclusion

The PINN approach of [1] works on this problem once the problem is well posed and the residuals are scaled by their own terms; it then reproduces a conventional solution to about a percent within a minute. As printed, the problem has no exact solution, and the paper's normalisation lets a non-solution reach a small loss, so a residual loss alone is not evidence of convergence. We recommend that PINN studies report a check against an independent solution of the same equations.

Publication requirements

Every TechUPGRADE technical report is assessed against eight requirements for a publishable paper (reproducibility, baselines, validation, statistics, sensitivity, generalisation, contribution, independent review). This report is a **preprint**: 6 of 8 requirements are met and 1 partly; the open items are listed.

requirement	status	evidence
R1 Reproducibility	met	Code in sim-manager/hall (model.py, reference.py, fd.py, pinn.py); parameters in model.py (DEFAULTS, CASES, ASSUMPTIONS); make_results.py and make_results_more.py regenerate every number, make_figures.py and gen_tex.py every figure and the text; the tenant hall reruns every case.
R2 Baseline comparison	met	Two independent references solve the same equations: collocation (scipy solve_bvp) and trapezoidal finite differences solved by Newton's method from a generic guess; they agree to 0.01 % in current and 0.09 % in every field at 100, 200 and 300 V.
R3 Validation	partial	Compared with the SPT-100's published operating point (Sankovic et al., IEPC-93-094): with the Bohm coefficient fitted to the published 4.5 A, the model predicts 76.7 mN against the published 80-83 mN. *Open:* measured axial profiles (plasma potential, electron temperature, ion velocity) or performance over several operating points, not one point after a fit.
R4 Statistical analysis	met	Five seeds per case: I_d 6.939 +/- 0.023 A (95 % CI, reference 6.906 A), field errors with 95 % intervals and their distributions; five seeds on the paper's scales (0.13 +/- 0.09 A); five seeds of the parametric network.
R5 Sensitivity analysis	met	Mass flow, Bohm coefficient, peak magnetic field and exit electron temperature varied one at a time; the discharge voltage (60-300 V); the ion-neutral collision model; the residual normalisation.
R6 Generalisation	met	A network with the voltage as an input, trained on 80-150 V, evaluated against the reference at voltages it did not see inside (90, 100, 130 V) and outside (170, 200 V) the range, five seeds.
R7 Scientific contribution	met	A Contributions section lists what is new: the boundary-condition analysis, the false convergence on fixed scales, two independent references, the SPT-100 comparison, sensitivity and generalisation.
R8 Independent review	missing	– *Open:* review by an electric-propulsion / plasma specialist independent of the work.

Reproducibility

Tenant `hall` of the TechUPGRADE platform, Simulations → Hall thruster: the Reference solver and PINN tabs run every case in this report. The code is in `sim-manager/hall` (equations once for NumPy and JAX, the reference, the PINN), and this report is in `docs/reports/hall-iepc-2024-814` (`make_results.py`, `make_results_more.py`, `make_figures.py`, `gen_tex.py`). The results were computed on 2026-10-31 01:30 UTC, and the experiments added for the publication requirements on 2026-11-01 19:29 UTC.

References

[1] S. Burrows, E. Webb, J. Agarwal, K. Naganathan, C. McDevitt, S. Roy, "Physics-Informed Simplified Model of a Hall Thruster," IEPC-2024-814, 38th International Electric Propulsion Conference, Toulouse, France, 2024.

[2] S. Roy, B. P. Pandey, "Plasma-wall interaction inside a Hall thruster," J. Plasma Phys. 68, 305–319 (2002).

[3] S. Roy, B. P. Pandey, "Numerical investigation of a Hall thruster plasma," Phys. Plasmas 9, 4052–4060 (2002).

[4] E. Ahedo, P. Martínez-Cerezo, M. Martínez-Sánchez, "One-dimensional model of the plasma flow in a Hall thruster," Phys. Plasmas 8, 3058–3068 (2001).

[5] M. Raissi, P. Perdikaris, G. E. Karniadakis, "Physics-informed neural networks," J. Comput. Phys. 378, 686–707 (2019).

[6] J. Kierzenka, L. F. Shampine, "A BVP solver based on residual control and the MATLAB PSE," ACM Trans. Math. Softw. 27, 299–316 (2001).

[7] X. Glorot, Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," AISTATS (2010).

[8] D. P. Kingma, J. Ba, "Adam: a method for stochastic optimization," arXiv:1412.6980 (2014).

[9] J. M. Sankovic, J. A. Hamley, T. W. Haag, "Performance Evaluation of the Russian SPT-100 Thruster at NASA LeRC," IEPC-93-094, 23rd International Electric Propulsion Conference (1993).

[10] J. Bradbury et al., "JAX: composable transformations of Python+NumPy programs," http://github.com/google/jax (2018).