Reproducing PINN–PONN plume reconstruction: what the probe data determine, what the prior decides, and a closed-form baseline

TechUPGRADE Digital Twin Platform

TechUPGRADE technical report TU-TR-2026-002 (tenant `ponn`)

2026-10-31

Preprint: 6/8 requirements met, 1 partial

Abstract

We reimplement PINN–PONN [1], a physics-informed network that turns two maps of Langmuir-probe ion-ram current, the mass flow and the thrust into two-dimensional fields of heavy-particle density and axial and radial velocity in an electric-propulsion plume, and test it on an independent synthetic argon plume (the paper's ground truth is not public) under the paper's conditions: 30 % bias and 30 % noise on a 2.5 cm grid. We find that these data do not determine the fields: in a convective plume the ram currents depend only on the mass flux, so plumes with different radial speed profiles and equal thrust give identical probe maps, mass flow and thrust, while their axial velocity differs by 9.8–16.2 %. Over five noise realisations of the probe data, each reconstructed with five seeds, PINN–PONN's selected axial velocity has an error of 17.4 ± 22.6 % (mean ± 95 % confidence interval; paper: 4.3 % on its reference case), although in every realisation the network produced a field within 2.9–6.5 %: the selection by mass flow and thrust picked a poor field in 2 of five. Two baselines without a network use the same data: a closed form, in which the mass flow removes the bias plane by plane (10.5 ± 0.6 %), and a regularised inversion (20.6 ± 1.9 %). On four plumes not used to choose the network's prior, the selected axial-velocity error ranges from 3.5 to 27.0 %. The study runs as a tenant of the TechUPGRADE platform.

Keywords: physics-informed neural networks, electric propulsion, plume diagnostics, Langmuir probe, inverse problems, identifiability

1 Introduction

Two-dimensional maps of density and velocity in a thruster plume are valuable and expensive. Maraqten et al. [1] propose to obtain them from cheap data: the ion-ram current densities $j_z = e n v_z$ and $j_r = e n v_r$ collected by two planar probes facing the axial and radial directions, combined with the facility's mass flow $\dot{m}$ and thrust T and with the convective conservation laws, through a physics-informed neural network (PINN) [3]. On a synthetic MPD plume with 30 % bias and noise they report axial velocity within 4.3 % mean absolute percentage error (MAPE), radial velocity 13.6 % and heavy-particle density 14.7 %, and state that global information renders the locally under-determined problem well posed. We reproduce the method and ask what the data actually determine.

2 Contributions

- An identifiability result: ram currents, mass flow, thrust and the convective conservation laws do not determine the radial speed profile of a convective plume; plumes with identical data differ in velocity by 10–16 %.
- A reimplementation of PINN–PONN with statistics over noise realisations and seeds: the network can produce accurate fields, but selecting them by mass flow and thrust is unreliable.
- Two baselines without a network on the same data – a closed form in which the bias cancels exactly, and a regularised inversion – and tests on plumes not used for any choice of the method.

3 The method

A residual network maps (z, r) to (v_{tot}, ϕ, n_h) through inductive biases: the speed as a correction to a characteristic exhaust speed with a geometric factor that falls off with radius, the angle as a correction to a geometric guess, and the

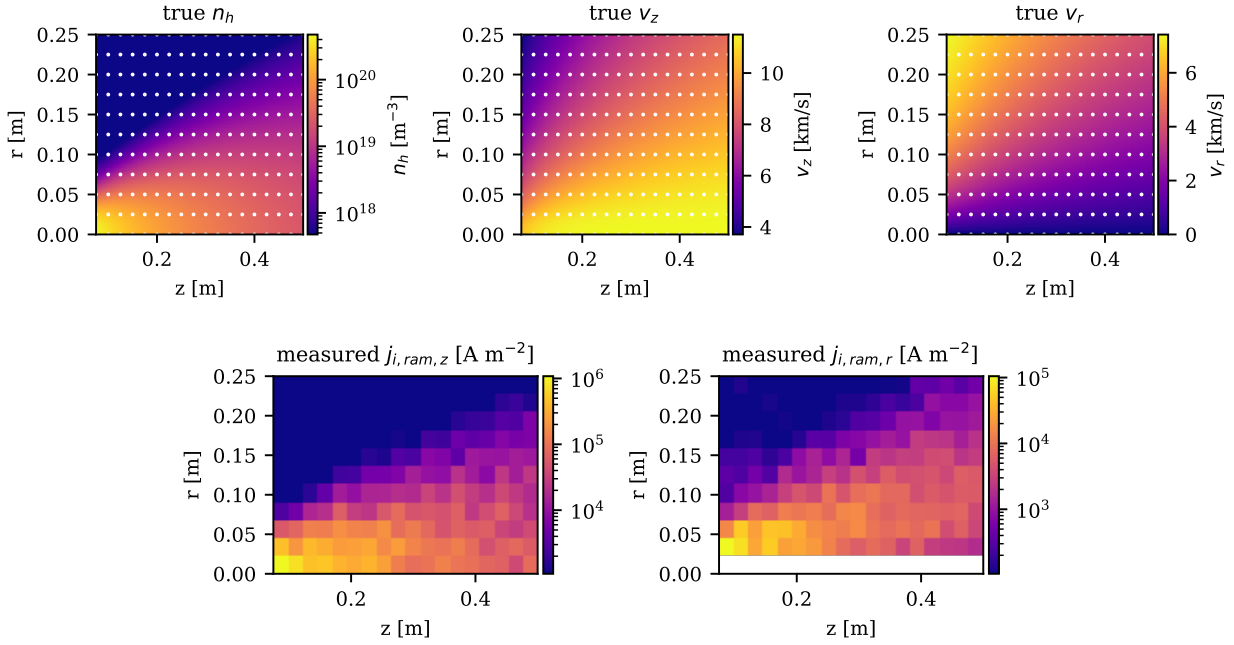


Figure 1: Top: the true fields, with the probe grid (dots). Bottom: the probe data, with 30 % bias and 30 % noise.

density as a bounded correction (two orders of magnitude) to a Gaussian radial profile. Seven losses – axial ram current, flow angle, mass flow and thrust (Riemann sums over each plane, the paper's Eq. 4), and convective mass, momentum and energy residuals – are trained one at a time, chosen at random each step, with AdamW [4] inside a look-ahead optimiser [5]; eight seeds of 20,000 epochs give 160,000 candidates, of which the one minimising $1.8L_{\dot{m}} + L_T$ is kept. The formulas of the inductive biases and of the momentum and energy residuals are given in a thesis we did not have; our versions are listed on the tenant's page.

4 Our synthetic plume and probes

Argon at 0.8 g/s flows along straight rays from a virtual source 5 cm behind the exit plane, with a mass flux per steradian made of a Gaussian core and a weak halo; mass is conserved exactly on every sphere. Along each ray the speed grows over 4 cm (the near field) to a terminal speed of 11.5 km/s on the axis, falling by $\kappa = 30$ % towards the edge. Beyond the acceleration zone the flow is ballistic and satisfies the convective conservation laws. The domain is $z \in [0.075, 0.5]$ m, $r \in [0, 0.25]$ m, as in the paper's figures; the plume's halo carries 4.0 % of the mass out through $r = 0.25$ m by $z = 0.5$ m. The probes sample j_z, j_r on a 2.5 cm grid with a uniform bias of 30 % and multiplicative log-normal noise of 30 %; the thrust is taken at $z = 0.15$ m where the acceleration has ended (Fig. 1).

5 What the data determine

The bias. A uniform bias b multiplies both currents. Their ratio, the flow angle $\phi = \arctan(j_r/j_z)$, does not see it, and $\dot{m} = 2\pi m_h \int n v_z r\, dr$ is the integral of j_z/e: dividing by it removes b exactly, plane by plane. Robustness to bias needs no network.

The speed. In a convective plume the ram current is the mass flux, $n\mathbf{v}$. With rays from a source, $n\mathbf{v} = F(\theta)\,\hat{\mathbf{e}}/(m_h R^2)$ for any speed profile: the probe maps do not contain the speed at all. The mass flow is the same integral. The convective momentum and energy equations, given the flux, only say that the speed is constant along each ray. The thrust, $\int n v_z^2$, is one number per plane, and in the ballistic region it is the same number on every plane. So the radial profile of the speed, and with it the split of $n\mathbf{v}$ into density and velocity, is not determined by the data. We checked this numerically: plumes with $\kappa = 0$ and 0.6, their terminal speed rescaled to give the same thrust, produce probe maps, mass flow and thrust equal to ours to 5×10^{-16}, while their density differs by 11.2 % (13.1 %) and their axial velocity by 9.8 % (16.2 %) (Fig. 2). Any method given only these data returns the same answer for all three plumes; the radial structure it returns comes from its assumptions – for PINN–PONN, the geometric factor of its speed prior.

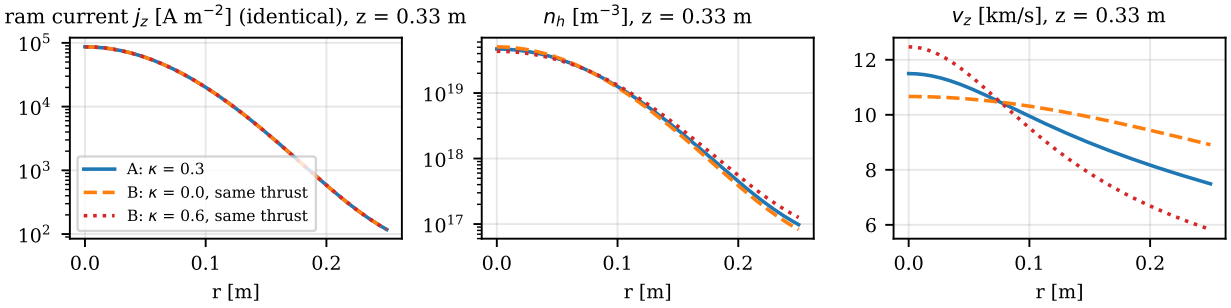


Figure 2: Three plumes with the same probe data, mass flow and thrust, and different density and velocity.

6 Results

6.1 The reference case

Table 1 lists the selected reconstructions. With our prior's speed drop $k_{geo} = 0.3$, equal to the true κ, PINN–PONN selects, on this one realisation of the probe noise, seed 2000, epoch 19,950, and reaches 3.1 % for v_z, 18.0 % for v_r and 34.0 % for n_h (Fig. 3). The axial velocity matches the paper's 4.3 %. The density error is larger than the paper's 14.7 %: it sits near the thruster, where the density is overestimated, and in the low-density halo at large angles, which the Gaussian density guess with a bounded correction cannot follow (Fig. 4). The four seeds took 92 s each on four CPU cores; the paper reports 15 min per seed on an A100.

Table 1: MAPE [%] (and R^2) of the selected fields against the true plume; v_r where $|v_r| \geq 500$ m/s, as in the paper.

reconstruction	v_z	n_h	v_r
Maraqten et al. (their plume, 8 seeds)	4.3 (0.98)	14.7 (0.97)	13.6 (0.88)
PINN–PONN, $k_{geo} = 0.3$ (reference)	3.1 (0.97)	34.0 (0.95)	18.0 (0.77)
PINN–PONN, $k_{geo} = 0$	10.8 (0.69)	33.4 (0.98)	17.7 (0.92)
PINN–PONN, $k_{geo} = 0.6$	26.8 (-0.73)	40.5 (0.91)	35.1 (-0.22)
PINN–PONN, ideal data	2.4 (0.98)	35.7 (0.89)	14.4 (0.85)
PINN–PONN, 50 % bias and noise	6.0 (0.89)	43.4 (0.96)	22.2 (0.79)
PINN–PONN, probes in $z \in [0.30, 0.40]$	3.2 (0.97)	36.8 (0.99)	27.5 (0.79)
closed form, one speed per plane	11.5 (0.63)	18.2 (0.93)	22.4 (0.59)
closed form, one axial speed per plane	24.0 (-0.46)	20.4 (0.88)	38.4 (-1.88)

6.2 The selection

Every checkpoint of the four seeds is shown against the metric in Fig. 5. As the paper reports, a small metric does not guarantee a small error: the best checkpoint by true error (seed 2000, epoch 19,800: v_z 3.9 %, n_h 33.5 %, v_r 18.4 %) is not the one selected.

6.3 The prior, the data quality, and extrapolation

With the prior's speed drop set to 0 or 0.6, the same data give v_z errors of 10.8 % and 26.8 % (Fig. 6): the accuracy follows the prior, as the analysis predicts. Without bias and noise the errors are 2.4/35.7/14.4 % (v_z/n_h/v_r); with 50 % bias and noise 6.0/43.4/22.2 %. With probe data only in $z \in [0.30, 0.40]$ m (the paper's extrapolation test) they are 3.2/36.8/27.5 %.

6.4 Statistics over noise realisations

The reference case above is one realisation of the probe noise. Table 2 and Fig. 7 repeat it for five realisations, each reconstructed with five seeds of 20,000 epochs (the baselines for ten). PINN–PONN's selected fields vary widely, from 3.1 to 45.0 % in v_z; the best checkpoint of each run by true error (which an experiment cannot know) stays within 2.9–6.5 %. The generator works; the selection by $1.8L_{\dot{m}} + L_T$ does not reliably find its good fields, as the paper itself cautions.

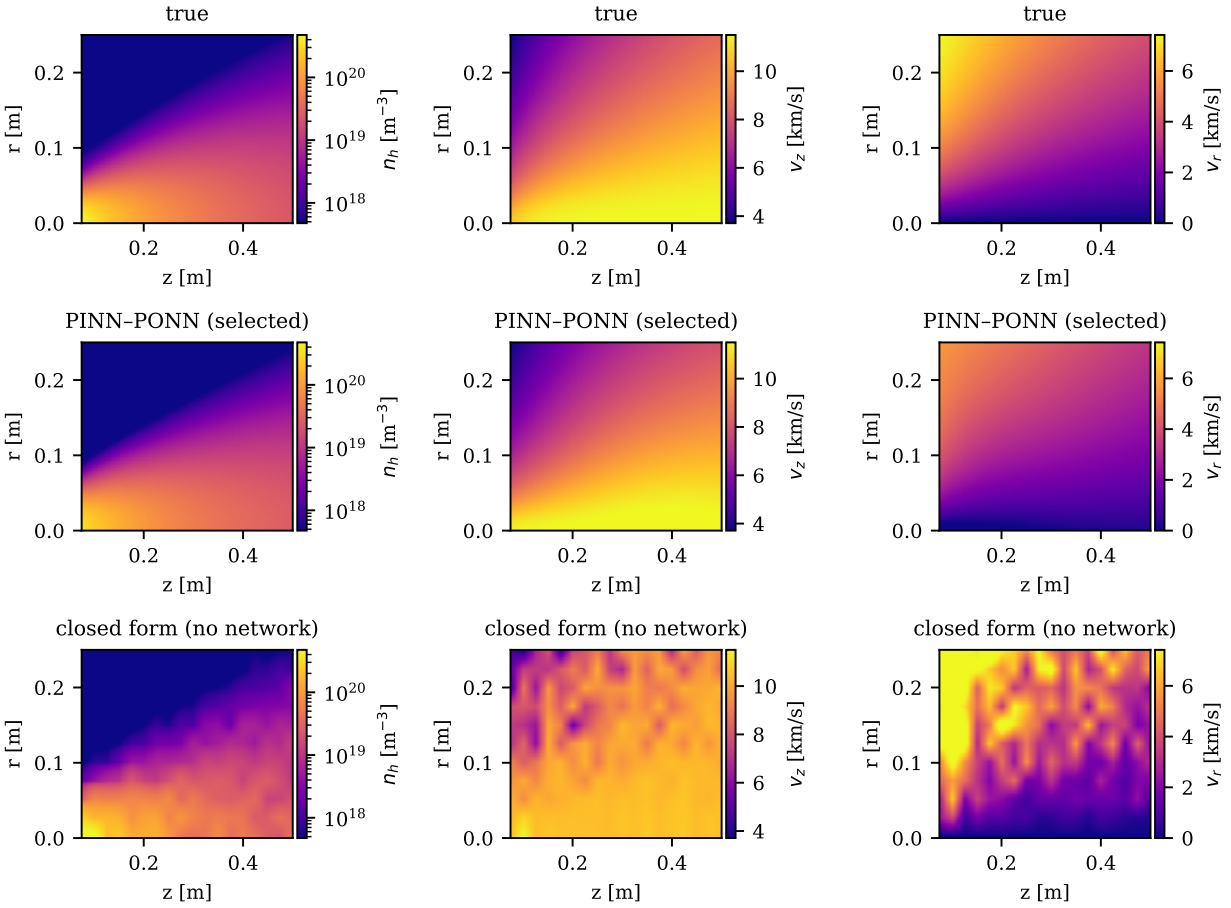


Figure 3: The true fields (top), PINN–PONN's selected reconstruction (middle) and the closed form (bottom), reference data.

Table 2: MAPE [%] over noise realisations: mean ± 95 % confidence interval (five realisations for PINN–PONN, ten for the baselines).

method	v_z	n_h	v_r
PINN–PONN (selected)	17.4 ± 22.6	44.4 ± 13.0	31.6 ± 20.5
closed form	10.5 ± 0.6	17.6 ± 0.7	23.3 ± 0.9
regularised inversion	20.6 ± 1.9	35.6 ± 5.2	26.4 ± 1.9

6.5 A second baseline: regularised inversion

A classical inverse-problem solution of the same data: $\log n_h$, $\log v$ and ϕ on the probe grid and one probe bias, fitted by nonlinear least squares to the axial ram current (in log form), the current ratio, the mass flow of every plane and the thrust past $z = 0.15$ m, with the speed held constant along the flow there (the convective laws) and second-difference smoothing. It estimates the bias as 27.8 % (true 30 %) and gives 20.6 ± 1.9 % for v_z: its speed profile, like every method's here, comes from its regularisation.

6.6 Plumes not used for tuning

The prior's speed drop ($k_{geo} = 0.3$) and the selection weight (1.8) were set on the reference plume. Four other plumes – a gentler and a steeper speed drop ($\kappa = 0.15, 0.45$), a lower mass flow (0.5 g/s) and a narrower core – give the results of Fig. 8: PINN–PONN's selected v_z error is 10.0 % ($\kappa = 0.15$), 3.5 % ($\kappa = 0.45$), 7.5 % (0.5 g/s), 27.0 % (narrow core); the closed form gives 6.6, 19.0, 11.5, 12.6 %.

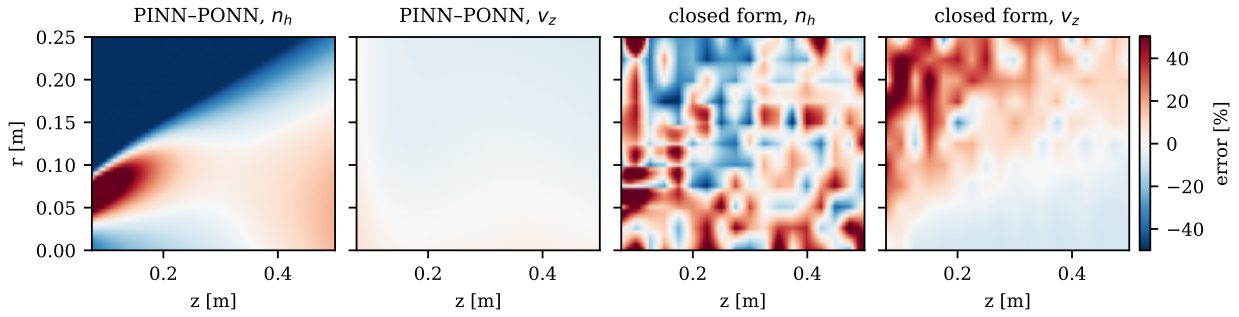


Figure 4: Relative error of density and axial velocity, PINN–PONN and closed form.

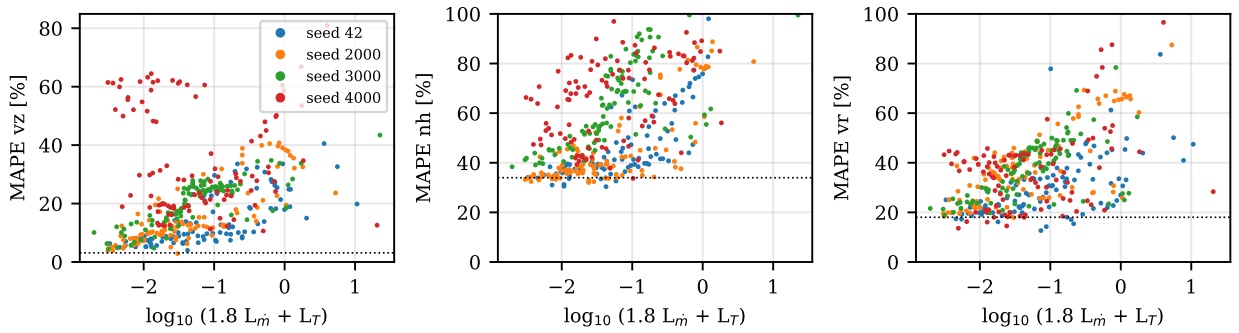


Figure 5: Selection metric against true error for every checkpoint; the dotted line is the selected reconstruction.

6.7 A closed form

Using the same information without a network – the mass flow to remove the bias in each probe plane, the current ratio for the angle, and the thrust for one total speed per plane – gives 11.5/18.2/22.4 % (PINN–PONN: 3.1/34.0/18.0 %). On this reference realisation, the closed form is more accurate for the density; and PINN–PONN is more accurate for the axial velocity and radial velocity. Over noise realisations (Section 6.4) the closed form is the more accurate and the far more stable. Its axial velocity cannot follow the radial speed profile, since it has one speed per plane; PINN–PONN's can, to the extent that its prior has the right profile. Its errors are exactly the same with and without bias (10.7 % and 10.7 % for v_z at 0 and 50 % bias, no noise).

7 Discussion and limitations

Our plume is not the paper's: it has no shocks, its halo is broader, and its speed profile is our choice, so the numbers cannot be compared one to one; we use our plume to test the mechanism. The analysis of Section 5 is independent of the plume: whenever the flow is convective, ram currents, mass flow and thrust leave the radial speed profile free, and a reconstruction can be only as good as its prior on that profile. The paper's inductive biases and its selection weight were set on the same reference case its results are reported for. Our implementation of the biases and of the momentum and energy residuals follows the paper's description, not its code, which is not public. The work was done with an AI coding assistant and has not been reviewed by an independent plasma physicist.

8 Conclusion

On our plume, PINN–PONN can produce accurate fields – within a few percent in axial velocity in every noise realisation – but the reconstruction it selects is accurate only some of the time (17.4 ± 22.6 % over five realisations), and a closed form using the same data is both more accurate on average and far more stable (10.5 ± 0.6 %). The data it uses determine the mass flux, the flow angle and each plane's mean speed, but not how the speed varies with radius: the network supplies that from its prior. A measurement that sees the speed directly (a retarding-potential analyser, Doppler-shifted emission, or probes at more than one orientation per point), and a selection that does not rest on mass flow and thrust alone, would be needed for the method to be reliable; the paper names both as future work.

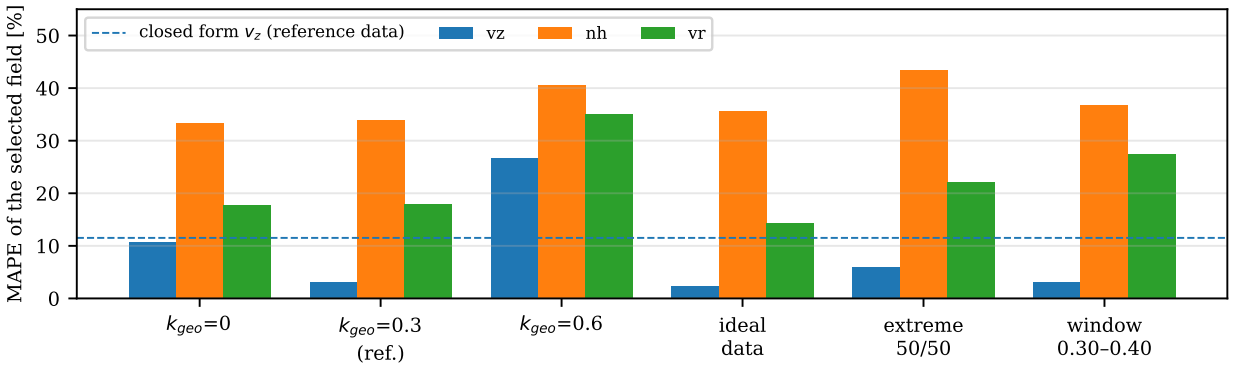


Figure 6: MAPE of the selected fields for three priors, ideal and extreme data, and probes in a window.

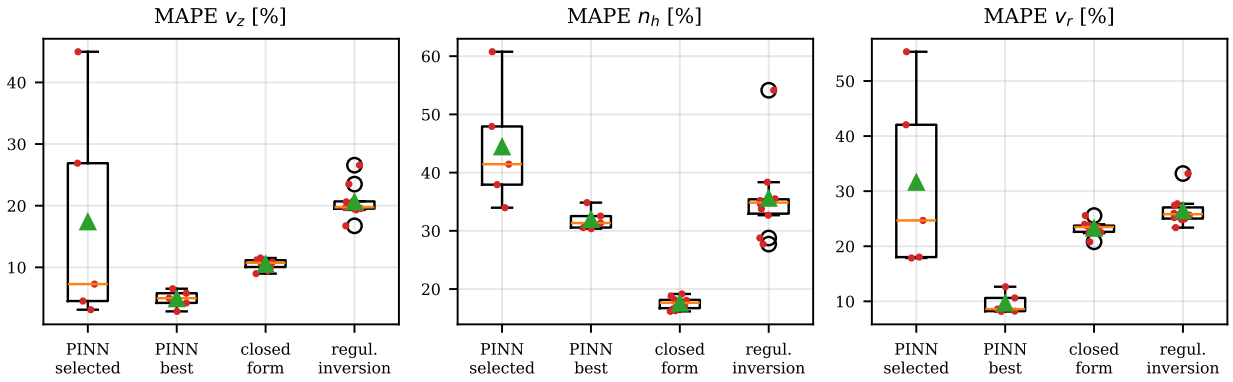


Figure 7: Errors over noise realisations: PINN–PONN's selected field, its best checkpoint by true error, and the two baselines.

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/ponn (plume.py, pinn.py, tikhonov.py); plume, probe and network parameters in plume.py and pinn.py (with 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 ponn reruns every case.
R2 Baseline comparison	met	Two independent baselines without a network on the same data: a closed form (bias removed by the mass flow, angle from the current ratio, one speed per plane from the thrust) and a regularised (Tikhonov) inversion by nonlinear least squares; both compared with PINN-PONN over ten noise realisations.
R3 Validation	partial	Every field is compared with an exactly known synthetic plume; it is our own plume, since the paper's SAMSA data are not public. *Open:* measured plume maps (Langmuir probe, LIF) or the authors' synthetic data.
R4 Statistical analysis	met	Five noise realisations of the probe data, each with five seeds of 20,000 epochs (the baselines over ten); means with 95 % confidence intervals and the distributions of the selected and best fields.
R5 Sensitivity analysis	met	The prior's speed drop (0, 0.3, 0.6), the probe bias and noise (0, 30, 50 %), the plume's speed profile (twin plumes) and the probe window.
R6 Generalisation	met	Four plumes not used to choose the prior or the selection weight (speed drop 0.15 and 0.45, 0.5 g/s, a narrower core), and extrapolation beyond the probe window.
R7 Scientific contribution	met	A Contributions section lists what is new: the identifiability result, the statistics showing the selection is unreliable, and two baselines without a network.
R8 Independent review	missing	– *Open:* review by a plasma-diagnostics / electric-propulsion specialist independent of the work.

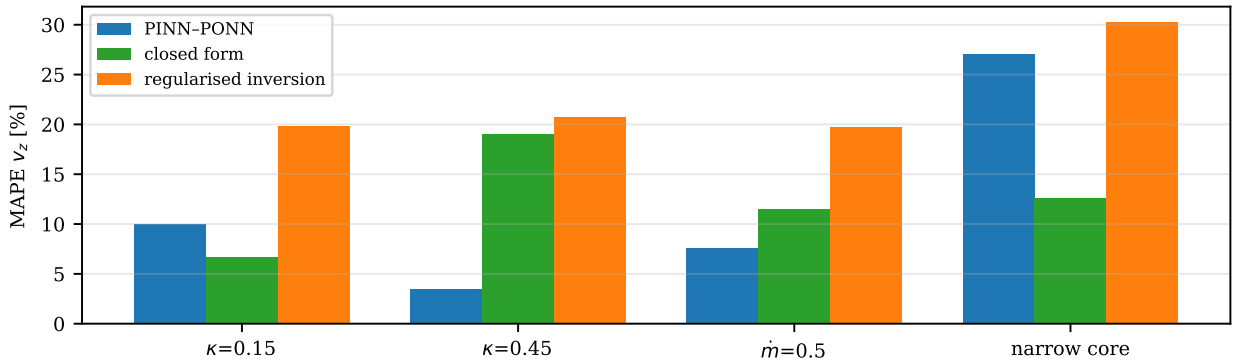


Figure 8: Axial-velocity error on four plumes not used to choose the prior or the selection weight (one realisation each).

Reproducibility

Tenant `ponn` of the TechUPGRADE platform, Simulations → Plume reconstruction: every case here can be rerun (Probes & data, Reconstruct, Same data, different plumes). Code: `sim-manager/ponn`; this report: `docs/reports/ponn-iepc-2025-287` (`make_results.py`, `make_results_more.py`, `make_figures.py`, `gen_tex.py`). Results computed on 2026-10-31 08:37 UTC, and the experiments added for the publication requirements on 2026-11-01 19:18 UTC.

References

[1] N. Maraqten, O. Nimer, J. Grill, M. Ehresmann, G. Herdrich, U. Soergel, "PINN–PONN: A Novel Physics-Informed Deep Learning Method for Plasma Flow Reconstruction," IEPC-2025-287, 39th International Electric Propulsion Conference, London (2025).

[2] R. B. Lobbia, B. E. Beal, "Recommended Practice for Use of Langmuir Probes in Electric Propulsion Testing," J. Propuls. Power 33, 566–581 (2017).

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

[4] I. Loshchilov, F. Hutter, "Decoupled Weight Decay Regularization," arXiv:1711.05101 (2019).

[5] M. R. Zhang, J. Lucas, G. Hinton, J. Ba, "Lookahead Optimizer: k steps forward, 1 step back," arXiv:1907.08610 (2019).

[6] S. Wang, X. Yu, P. Perdikaris, "When and why PINNs fail to train: A neural tangent kernel perspective," J. Comput. Phys. 449, 110768 (2022).