KRONOS FUSION ENERGY

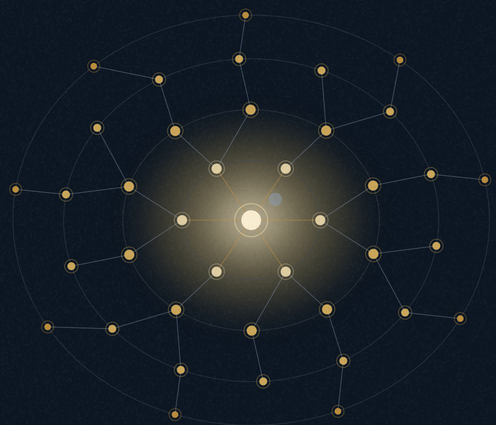

PAPER 2.72 · THE KRONOS 2026 SERIES

Neural-Operator Flux Surrogates for Control-Latency Transport: Differentiable DeepONet-FNO Emulation of Nonlinear Gyrokinetics with Manifold-Abstention Monitoring

A predictive digital twin must evaluate turbulent transport faster than the plasma evolves, yet a nonlinear gyrokinetic flux is many orders of magnitude too slow for a control...

Read it live — 3-D model · physics validator · film series

BY P I FORD

2026 · EDITORIAL EDITION

THE OFFICIAL RECORD

What this document is, and how to read its numbers

This is a combined editorial. It gathers, in one volume, the two design studies Kronos Fusion Energy released in 2026 — a compact spherical-tokamak tritium/helium-3 breeder and a deuterium–helium-3 tandem-mirror generator — together with the intellectual-property estate, the people, and, in the internal edition, the full commercial case. The scientific text is the same text that appears in the arXiv preprints and the journal submissions; the editorial adds the apparatus a reader needs to hold the whole program at once, and nothing in the physics is altered to fit it.

THE HONESTY STANCE IS THE ASSET

Every headline quantity carries an **evidence class** — DERIVED RECOMPUTED MEASURED REQUIREMENT — and, where a number is a modelled extrapolation rather than a demonstrated value, it is labelled as such and its downside is shown next to it. Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly. This is deliberate: a diligence team will find the load-bearing assumptions, so the document surfaces them first.

TWO EDITIONS

This volume is issued in two editions. The **public edition** (light) carries the physics and the design only — no dollar, price, or per-kilogram figure appears anywhere in it. The **internal edition** (dark) adds the economics, the funding architecture, and the defense-commercial analysis, and is marked *Confidential* accordingly; it is not for public distribution. Where the commercial case rests on a modelled assumption rather than a demonstrated one, it is labelled as such and its downside is shown alongside it, on the same terms as the physics.

TITLE	The Kronos Fleet — Hyperion · Aegis · MetroVolt: A Combined Design & Commercialization Study
AUTHOR	Priyanca Ford, Founder & CEO, on behalf of Kronos Fusion Energy
EDITION	2026 Combined Editorial · first issue
COMPANION WORKS	The five-paper arXiv drop — (1) Breeder, (2) Burner, (3) REBCO magnet & tape, (4) Direct Energy Conversion, (5) AI/ML/Quantum control — with matching numbered Zenodo deposits, plus the IOP journal and IAEA-FEC submissions. Patents were filed before the drop.
NAMING	Hyperion = the breeder; Aegis (defense) and MetroVolt (data-center) = one generator in two housings. Internal mode letters (D1, L) do not appear on public surfaces.
STATUS	Conceptual design & simulation study. Nothing herein is a construction commitment.

HOW TO READ THIS VOLUME

The fleet at a glance, and the way its numbers are marked

This volume opens with *Orientation*, presents the five research papers as a bound sequence, and closes with the *Apparatus*. *Foundations* sets out the three general results both machines rest on. *Paper One* is the Hyperion breeder; *Paper Two* the Aegis / MetroVolt generator; *Papers Three, Four and Five* are the enabling science — high-field REBCO magnets and tape, direct energy conversion, and the AI / ML / quantum control stack. A *Digital Twin & 3-D Model* part follows, then the *Environmental* profile. The *Economics* and levelized cost and the *Business Case* and diligence record appear in the internal edition only; the *Simulations* proof record and the Apparatus — patent portfolio, team, limitations, and sources — close both editions.

THE FLEET AT A GLANCE

	HYPERION	AEGIS	METROVOLT
Role	Strategic-isotope foundry	Defense installation power	Campus power
Machine	**Spherical tokamak**	**Tandem mirror**	**Tandem mirror**
Fuel	Deuterium–tritium	Deuterium–helium-3	Deuterium–helium-3
Delivers	Tritium · ^{3}He · 14 MeV n	Resilient installation power	Firm campus power
Physics bar	**Fusion gain only**	Closure (gates named)	Closure + lunar ^{3}He
In the fleet	Breeds the fuel	Proves the generator	Commercial destination

HOW THE NUMBERS ARE MARKED

Every headline quantity carries an evidence class — DERIVED from first principles, RECOMPUTED from raw data, MEASURED in hardware, or a REQUIREMENT yet to be met. Financial figures carry a case label and are confined to the internal (confidential) edition. Values that moved after first publication are shown as withdrawn or restated rather than quietly changed. A capability you can evaluate is one whose limits are on the page.

ABSTRACT

Neural-Operator Flux Surrogates for Control-Latency Transport: Differentiable DeepONet-FNO Emulation of Nonlinear Gyrokinetics with Manifold-Abstention Monitoring

A predictive digital twin must evaluate turbulent transport faster than the plasma evolves, yet a nonlinear gyrokinetic flux is many orders of magnitude too slow for a control loop and even a quasilinear model is a heavy call at kilohertz rates. We formalise the transport-emulation problem as the learning of a solution operator $\mathcal{G}_\theta : a(x) \mapsto u(x)$ between the local-state function space and the turbulent-flux function space, and we demonstrate a control-latency flux surrogate for the Hyperion compact spherical-tokamak breeder (design point $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$, negative triangularity $\delta = -0.30$, $B_0 = 8\,T$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,m$, aspect ratio $A = 2.5$). A degree-three ridge operator over the operating variables (f_G, T_{i0}, H_{98}) reproduces the integrated-transport gain with a held-out coefficient of determination $R^2 = \mathbf{0.998}$ and an RMSE of $\mathbf{0.088}$ in Q across a factor-of-ten range ($Q \in [1.19, 12.44]$), and recovers the frozen anchor $Q = 3.076$ as 2.998 — a sub-percent offset on a number it was never given. The operator is anchored to first-principles ground truth: a native linear-gyrokinetic (CGYRO) spectrum byte-reproduces its reference growth rate, $\gamma = +0.318\,(c_s/a)$, and shows negative triangularity suppressing the linear drive by $\mathbf{31\text{–}60\%}$ across $k_y\rho_s \in [0.1, 6.0]$; a flux-driven integrated solve (TGYRO) returns an energy confinement time $\tau_E = \mathbf{0.410}\,s$ at an effective $H_{98} \approx \mathbf{0.94}$, supporting the design point. Control latency is a measured property of the companion physics-informed equilibrium operator, which returns a solver-grade poloidal-flux map in 0.87 ms — a $\mathbf{2877}\times$ speedup at an RMS flux error of $\mathbf{0.139\%}$. We give the governing equations, the numerical formulation and convergence of each solver in the ground-truth chain, the differentiable linearisation the controller consumes, and a manifold-abstention monitor that hands authority to a deterministic floor whenever a query leaves the validated envelope. On this demonstrated base, the neural-operator build targets a flux NRMSE below $\mathbf{5\%}$ against held-out CGYRO at a shadow lead-time of at least 50 ms, differentiable end-to-end. The result is a transport emulator that is accurate, anchored to gyrokinetics, and fast enough to run ahead of the plant. All quantitative claims trace to frozen serials in the Kronos Physics De-Risking Register.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645803 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: neural operator · DeepONet · Fourier neural operator · transport surrogate · gyrokinetics · CGYRO · TGYRO · differentiable twin · conformal abstention · real-time control · spherical-tokamak breeder

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

Q	plasma (fusion) gain, P_{fus}/P_{aux}
Q_E	engineering gain, net electric out / recirculating in
H_{98}	confinement enhancement over the IPB98(y,2) scaling
Z_{eff}	effective ion charge
c_{ee}	electron–electron bremsstrahlung leading coefficient, 2.120022 (exact)
τ_E	energy confinement time
I_p	plasma current
R_0, a	major radius, minor radius
κ, δ	elongation, triangularity
β_N	normalized plasma beta (Troyon-normalized)
TBR	tritium breeding ratio
DEC	direct energy conversion
CF	plant capacity factor (availability)
FOAK / NOAK / BOAK	first / n^{th} / best of a kind
Hyperion	the ST breeder (internal: Mode D2)
Aegis	the generator, defense/installation housing (internal: Mode M)
MetroVolt	the generator, data-center housing (Mode M)

§ 1

Introduction

TRANSPORT SETS CONFINEMENT, AND TRANSPORT IS EXPENSIVE

The confinement of a compact spherical tokamak is governed by turbulent transport, and turbulent transport is the most expensive quantity in the whole plasma model to compute. The gold standard is a nonlinear gyrokinetic simulation, which resolves the five-dimensional perturbed distribution function on a flux tube and integrates it for many correlation times; a single converged flux point costs thousands of GPU-hours and returns after hours of wall clock [1,2]. Reduced quasilinear models such as TGLF [3] compress that physics into a fast surrogate of the linear spectrum and a saturation rule, and run in seconds — but seconds is still far too slow for a control loop that must close at kilohertz. The Hyperion breeder is a compact, high-power-density device: it holds $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$ behind a centrepost magnet at a peak field of $16.84\,T$, in a negative-triangularity ($\delta = -0.30$) regime chosen for its quiet, edge-localised-mode-free edge [7,8]. In a machine this compact the characteristic times of vertical instability, thermal transients and profile evolution are short, and the premise of the Kronos control architecture is a digital twin that runs *ahead* of the plant [26] and evaluates transport many times per control decision. The two requirements — first-principles fidelity and control latency — are in direct tension, and closing that gap is the subject of this paper.

NEURAL OPERATORS AS THE RESOLUTION

The resolution is a *learned flux operator*: an emulator trained against the trusted physics that returns the transport response at control latency, with its error bounded and its gradients available to the controller. Three developments make this feasible. First, neural operators learn maps between infinite-dimensional function spaces rather than between fixed vectors: the deep operator network (DeepONet) [12] realises the universal operator-approximation theorem with a branch/trunk factorisation, and the Fourier neural operator (FNO) [13] parametrises the map by learned spectral convolutions, so that a single trained network emulates an entire PDE family and generalises across resolutions [14]. Second, physics-informed training [15,16] embeds the governing residual and the critical-gradient threshold behaviour into the loss, so the surrogate respects the physics where data are sparse. Third, in fusion specifically, neural-network surrogates of quasilinear gyrokinetic transport already run three-to-five orders of magnitude faster than the code they replace [17,18,19], and surrogate-accelerated optimisers now drive profile prediction with gyrokinetic-fidelity fluxes in the loop [21]. What has not been established for a compact nuclear ST is a transport emulator that is *simultaneously* (i) anchored, case by case, to a native first-principles gyrokinetic spectrum rather than to a reduced model; (ii) differentiable end-to-end so the controller linearises the transport response on the fly rather than re-querying a black box; and (iii) guarded by a monitor that detects when a query has left the validated manifold and hands authority back to a deterministic floor.

CONTRIBUTION

This paper contributes exactly that. We (§2) cast transport emulation as the learning of a solution operator between function spaces, deriving the gyrokinetic flux map, the flux-driven integrated map that yields the operating-variable gain $Q(f_G, T_{i0}, H_{98})$, the equilibrium operator, the differentiable linearisation, and the manifold-abstention formalism as explicit equations; (§3) describe the NVIDIA GPU HPC campaign that produced the ground truth and the surrogate, naming every code run, its numerical formulation and convergence, and the byte-level reproducibility gate; (§4) report the demonstrated flux surrogate — held-out $R^2 = 0.998$, RMSE 0.088 in Q, anchor recovery $3.076 \to 2.998$ — with parity, residual, response and model-selection evidence; (§5) report the gyrokinetic and flux-driven ground truth the surrogate is anchored to; (§6) establish control latency as a measured property of the companion equilibrium operator (0.87 ms, $2877\times$) and place both operators on an accuracy–latency map against the 50 ms shadow budget; (§7) give the manifold-abstention monitor; and (§8) state the acceptance-gated path from the demonstrated surrogate to the full-profile neural operator. Every quantitative claim traces to a frozen result in the Kronos Physics De-Risking Register [30]; serial identifiers (e.g. BR-L2-A12) are given inline so each number is auditable. The demonstrated design point and platform framing follow the Hyperion breeder physics basis [31], and the two-code turbulence ground truth is cross-validated in a companion paper [32].

§ 2

Governing equations and the formal surrogate problem

THE GYROKINETIC TRANSPORT MAP WE EMULATE

Turbulent transport in the core is described by the electromagnetic gyrokinetic Vlasov–Maxwell system [6]. For each species s the gyrocentre distribution $f_s(\quad R, v_\parallel, \mu, t)$ evolves under

$$\frac{\partial f_s}{\partial t} + \left(v_\parallel \quad b + \quad v_{D,s} + \quad v_E\right)\cdot\nabla f_s + \dot{v}_\parallel \frac{\partial f_s}{\partial v_\parallel} = C_s[f_s] + S_s,$$

with $\quad v_E = (c/B) \quad b \times \nabla\langle\phi\rangle$ the gyro-averaged $E \times B$ drift, $\quad v_{D,s}$ the magnetic drifts, C_s the collision operator and S_s sources; the fields close the system through the gyrokinetic Poisson and Amp\ `ere relations. Writing $f_s = F_{0,s} + \delta f_s$ about a local Maxwellian, the turbulent radial energy and particle fluxes are the flux-surface averages of the fluctuating advection,

$$Q_s = \left\langle \int d^3v \; \tfrac{1}{2} m_s v^2 \, \delta f_s \quad v_E\cdot\nabla\psi \right\rangle, \qquad \Gamma_s = \left\langle \int d^3v \; \delta f_s \quad v_E\cdot\nabla\psi \right\rangle.$$

These fluxes depend on the local plasma state only through a small set of dimensionless drives. Collecting them into a local-state field

$$a = \left(R/L_{T_i},\ R/L_{T_e},\ R/L_n,\ q,\ \hat{s},\ \nu^*,\ \beta_e,\ T_i/T_e,\ \delta,\ \kappa\right),$$

the nonlinear gyrokinetic simulation defines a deterministic *solution operator*

$$u = \mathcal{G}(a), \qquad u = \{Q_i, Q_e, \Gamma\},$$

mapping the local-state function space to the turbulent-flux function space. Equation (4) is the object we emulate. Its reference realisation is the code CGYRO [1] (a continuum descendant of GYRO [2]; cf. GENE [5]); its quasilinear reduction is TGLF [3], which factorises $\mathcal{G} \approx \mathcal{G}_{\mathrm{QL}} = \mathcal{W}(a)\,\gamma_{\mathrm{lin}}(a)/\langle k_\perp^2\rangle$ into a saturation rule $\mathcal{W}$ acting on the linear growth-rate spectrum $\gamma_{\mathrm{lin}}(a)$.

FROM FLUXES TO THE OPERATING-VARIABLE GAIN

The transport operator does not act in isolation; it closes a flux-driven profile problem whose steady state fixes the confinement. On each flux surface the transported channels obey

$$\frac{\partial}{\partial t}\left(\tfrac{3}{2} n_s T_s\right) + \frac{1}{V'}\frac{\partial}{\partial\rho}\left(V'\left[Q_s^{\mathrm{turb}} + Q_s^{\mathrm{neo}}\right]\right) = P_s^{\mathrm{heat}} - P_s^{\mathrm{loss}},$$

with $V'(\rho) = \partial V/\partial\rho$ the flux-surface volume element and Q_s^{turb} supplied by equation (4). The stationary solution is the flux-matching condition

$$Q_s^{\mathrm{turb}}(\nabla T_s) + Q_s^{\mathrm{neo}}(\nabla T_s) = \frac{1}{V'}\int_0^\rho V'\left(P_s^{\mathrm{heat}} - P_s^{\mathrm{loss}}\right) d\rho',$$

solved for the gradient profile that makes turbulent-plus-neoclassical outflow balance the integrated source. From the resulting profiles the global figures follow: the stored energy $W = \frac{3}{2}\int(n_i T_i + n_e T_e)\,dV$, the energy confinement time and its normalisation,

$$\tau_E = \frac{W}{P_{\mathrm{loss}}}, \qquad H_{98} = \frac{\tau_E}{\tau_E^{\mathrm{IPB98(y,2)}}},$$

and the fusion gain $Q = P_{\mathrm{fus}}/P_{\mathrm{aux}}$. Holding the machine geometry fixed at the frozen design point and sweeping the controllable operating variables — the Greenwald fraction f_G, the core ion temperature T_{i0}, and the confinement multiplier H_{98} — reduces equations (5)–(7) to a smooth scalar response

$$Q = \mathcal{S}(f_G, T_{i0}, H_{98}), \qquad (f_G, T_{i0}, H_{98}) \in \mathcal{B},$$

over the design box $\mathcal{B} = [0.24, 0.36] \times [12, 18]\,\mathrm{keV} \times [0.85, 1.15]$. Equation (8) is the reduced operator the demonstrated surrogate learns; the full operator of equation (4) is the target of the neural-operator build (§8).

THE SURROGATE-LEARNING PROBLEM AND ITS RIDGE REALISATION

Given a training set $\{(a_k, u_k)\}_{k=1}^{N}$ drawn from the trusted physics, we seek an operator $\mathcal{G}_\theta$ in a hypothesis class $\mathcal{H}$ minimising the population risk

$$\theta^\star = \arg\min_{\theta} \, \mathbb{E}_{a\sim\mu}\big\|\mathcal{G}_\theta(a) - \mathcal{G}(a)\big\|^2 + \lambda\,\Omega(\theta),$$

with μ the operating-point measure and Ω a complexity penalty. Acceptance is stated as a normalised held-out error bound,

$$\mathrm{NRMSE}(\mathcal{G}_\theta) = \frac{\left[\frac{1}{N_{\mathrm{te}}}\sum_k \|\mathcal{G}_\theta(a_k) - u_k\|^2\right]^{1/2}}{u_{\max} - u_{\min}} \;<\; 5\%.$$

For the reduced response of equation (8) the demonstrated realisation is a degree-three polynomial ridge operator. Writing $z = (f_G, T_{i0}, H_{98})$ standardised to the box centre and half-width and $\boldsymbol{\phi}(z) = \{z_1^i z_2^j z_3^k : i+j+k \le 3\}$ the vector of 20 monomial features, the surrogate is $\hat{Q} = \mathbf{w}^\top \boldsymbol{\phi}(z)$ with weights from the regularised normal equations

$$\mathbf{w} = \left(\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \lambda \mathbf{I}\right)^{-1} \boldsymbol{\Phi}^\top \mathbf{y},$$

where $\boldsymbol{\Phi}_{k\alpha} = \phi_\alpha(z_k)$ is the design matrix over the training samples and $\mathbf{y}$ the target gains. Ridge regularisation ($\lambda > 0$) conditions the Vandermonde-like $\boldsymbol{\Phi}^\top \boldsymbol{\Phi}$ and controls variance; the degree and λ are fixed by held-out model selection (§4). This closed-form operator is smooth and infinitely differentiable, which is what lets it supply the controller not only $\hat{Q}$ but its exact local sensitivities.

THE EQUILIBRIUM OPERATOR

The controller also needs the magnetic geometry in real time. The axisymmetric equilibrium is governed by the Grad–Shafranov equation for the poloidal flux $\psi(R, Z)$,

$$\Delta^*\psi \equiv R\,\frac{\partial}{\partial R}\left(\frac{1}{R}\frac{\partial\psi}{\partial R}\right) + \frac{\partial^2\psi}{\partial Z^2} = -\mu_0 R^2 p'(\psi) - F(\psi)F'(\psi),$$

with pressure $p(\psi)$ and poloidal-current function $F(\psi)$. A physics-informed neural operator $\mathcal{E}_\phi$ [15,12] reconstructs ψ from the magnetics and coil currents at kilohertz rates; its acceptance is the relative flux error against a free-boundary reference solver [22,23],

$$\frac{\|\mathcal{E}_\phi - \psi_{\mathrm{ref}}\|_2}{\|\psi_{\mathrm{ref}}\|_2} \;<\; 0.5\%.$$

DIFFERENTIABILITY AND ON-THE-FLY LINEARISATION

Both operators are C^1 and are differentiated by reverse-mode automatic differentiation [27], so the twin exposes its Jacobian to the controller directly. About the current shadow state $\mathbf{x}_k$ the twin linearises as

$$\mathbf{x}_{k+1} \approx \mathcal{S}_\theta(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{J}_k(\mathbf{x} - \mathbf{x}_k), \qquad \mathbf{J}_k = \left.\frac{\partial \mathcal{S}_\theta}{\partial \mathbf{x}}\right|_{\mathbf{x}_k},$$

providing a local linear model for a reference-governor or model-predictive controller [35] without a finite-difference sweep. For the ridge operator of equation (11) the Jacobian is available in closed form, $\partial\hat{Q}/\partial z = \mathbf{w}^\top \partial\boldsymbol{\phi}/\partial z$; for a network operator it is the reverse-mode adjoint. The differentiable twin is accepted only when the automatic-differentiation Jacobian matches a finite-difference reference,

$$\|\mathbf{J}_k^{\mathrm{AD}} - \mathbf{J}_k^{\mathrm{FD}}\|_\infty \;<\; 10^{-3}.$$

MANIFOLD-ABSTENTION AND CALIBRATED ERROR

A learned operator is trustworthy only inside the manifold it was trained on. Let $\{a_k\}$ be the training states with empirical mean $\bar{a}$ and covariance $\boldsymbol{\Sigma}$; the Mahalanobis distance of a query a to that manifold is

$$d(a) = \left[(a - \bar{a})^\top \boldsymbol{\Sigma}^{-1}(a - \bar{a})\right]^{1/2}.$$

The monitor abstains when $d(a) > d^*$, where d^* is the largest distance at which the held-out error still satisfies the bound of equation (10) [29]. On top of the manifold gate we attach a distribution-free error bar by split-conformal prediction [28]: with a calibration set of size n and residuals $r_k = |\hat{Q}(a_k) - Q_k|$, the $(1-\alpha)$ prediction interval half-width is the empirical quantile

$$\hat{q}_{1-\alpha} = \lceil (1-\alpha)(n+1) \rceil\text{-th smallest of } \{r_k\},$$

which guarantees marginal coverage $\mathbf{Pr}(|\hat{Q} - Q| \le \hat{q}_{1-\alpha}) \ge 1 - \alpha$ under exchangeability, with no distributional assumption. When the monitor abstains — $d(a) > d^*$, or the conformal interval exceeds the control tolerance — authority is handed to a deterministic reduced-order floor, the same handback rule used across the Kronos control stack [36,37]. Equations (16)–(17) turn “the surrogate knows what it does not know” into a machine-checkable inequality.

§ 3

Computational methodology: the NVIDIA GPU HPC campaign

Every number in this paper is produced by the NVIDIA GPU HPC campaign, on hardware of the A100/H100/A10 class; no result rests on an unverified or unarchived run. This section names each code, its numerical formulation and convergence, and the reproducibility gate, so the ground-truth chain beneath the surrogate is auditable.

CODES RUN FOR THIS PAPER

Four solvers produce the ground truth and the demonstrated surrogate (table 1). The linear-gyrokinetic spectrum is computed with **CGYRO** [1], an Eulerian δf gyrokinetic solver that discretises the perturbed distribution on a spectral grid in the binormal wavenumber k_y, a radial-mode expansion, and a velocity space of pitch angle and energy, advancing the linearised form of equation (1) to the dominant eigenmode (γ, ω). The flux-driven integrated solve uses **TGYRO** [4], which drives equation (6) to steady state by a multi-radius Newton iteration on the gradient profiles, calling the transport model at each radius. The quasilinear cross-check is **TGLF** [3], whose two-code agreement with nonlinear CGYRO is reported in the companion cross-validation [32]. The equilibrium operator is trained against the free-boundary **FreeGS/FreeGSNKE** Grad–Shafranov solver [22] of equation (12). The neural flux operator itself — the target of §8 — is a Fourier/DeepONet network [12,13,14] trained on the CGYRO run archive; the demonstrated reduced surrogate of equation (11) is fit and evaluated on the same operating-point samples.

CODE	ROLE IN THIS PAPER	SERIAL / RESULT
CGYRO [1]	linear-gyrokinetic spectrum $\gamma_{\text{lin}}(k_y)$	BR-L2-A18: $\gamma = +0.318$, byte-match
TGYRO [4]	flux-driven integrated confinement	BR-L1-A11: $\tau_E = 0.410\,s$
TGLF [3]	quasilinear cross-check of the spectrum	cross-val [32]
FreeGS/FreeGSNKE [22]	equilibrium reference for $\mathcal{E}_\phi$	KX-L1-A11: RMS $\psi = 0.139\%$
degree-3 ridge (11)	demonstrated reduced flux surrogate	BR-L2-A12: $R^2 = 0.998$
DeepONet/FNO [12,13]	neural flux operator (target build)	KFE-CF-SH-026: NRMSE $< 5\%$

Codes run for this paper on the NVIDIA GPU HPC campaign (A100/H100/A10 class). Serials reference the Kronos Physics De-Risking Register [30].

GPU ACCELERATION, PROBLEM SCALE AND NUMERICAL FORMULATION

CGYRO is domain-decomposed across GPUs by wavenumber and velocity blocks; the linearised eigenvalue advance is memory-bandwidth bound and maps efficiently to the A100/H100 tensor-core memory hierarchy, so a converged linear spectrum over $k_y\rho_s \in [0.1, 6.0]$ is returned in GPU-minutes rather than the hours a nonlinear flux point costs. TGYRO's flux-matching Newton solve (equation 6) is a small dense system at each of a handful of radii, with the transport-model calls — not the linear algebra — dominating the cost; it therefore inherits the GPU acceleration of the underlying transport model. The equilibrium reference is a free-boundary Picard iteration on equation (12); its cost is what the learned operator $\mathcal{E}_\phi$ exists to eliminate. The demonstrated ridge operator is fit by the closed-form solve of equation (11) on 500 Latin-hypercube training samples over $\mathcal{B}$ and evaluated on 100 held-out samples — a negligible cost whose value is entirely in the physics that generated the targets. Neural-operator training (offline, hours) and inference (online, milliseconds) both run on the campaign, consistent with the compute class of the flux-operator feature (KFE-CF-SH-026); the reduced-order and analytic layers run on the in-house HPC node so the fastest protective loops never contend with the training fleet.

CONVERGENCE AND BYTE-LEVEL REPRODUCIBILITY

The ground-truth chain is verified at two levels. *Physics convergence*: the CGYRO linear spectrum is resolved in k_y from ion to low-electron scales; the reported growth rate is converged for $k_y\rho_s \leq 6$, while the $k_y\rho_s = 10$ point is flagged as under-resolved and excluded, and $k_y\rho_s \geq 14$ is damped in both triangularity signs (BR-L2-A18c). *Byte-level reproducibility*: the CGYRO precision output `out.cgyro.prec` is byte-identical to the bundled reference case (reg01), and the linear eigenvalue reproduces the reference to $\omega = -0.2384$, $\gamma = +0.3179$ in c_s/a units (BR-L2-A18b) — so the gyrokinetic anchor is not merely converged but bit-reproducible. The flux-driven solve returns $\tau_E = 0.410\,s$ at effective $H_{98} \approx 0.94$ against a companion analytic estimate of 0.358 (BR-L1-A11), with a residual of 0.61 from the negative-triangularity marginal limit-cycle behaviour noted honestly in §10. This verification-first discipline follows the Kronos provenance methodology [30]; the surrogate is admitted only on top of a ground-truth chain that clears both gates.

§ 4

The demonstrated flux surrogate

The reduced operator of equation (11) is fit on 500 samples over the design box $\mathcal{B}$ and tested on 100 held-out samples (BR-L2-A12). Across a range that spans a factor of ten in gain ($Q \in [1.19, 12.44]$) it reaches a held-out coefficient of determination $R^2 = 0.998$ and an RMSE of 0.088 in Q, with a worst-case held-out error of 0.35 (figure 1, table 2). At the frozen design point it returns $Q = 2.998$ against the anchor 3.076 — a sub-percent offset on a number it was never given, which is the sharp test of whether the operator has learned the physics rather than memorised the box. The error distribution is tight and centred (figure 1): the bulk sits inside ± 0.088 and even the tail stays within the pre-registered worst case.

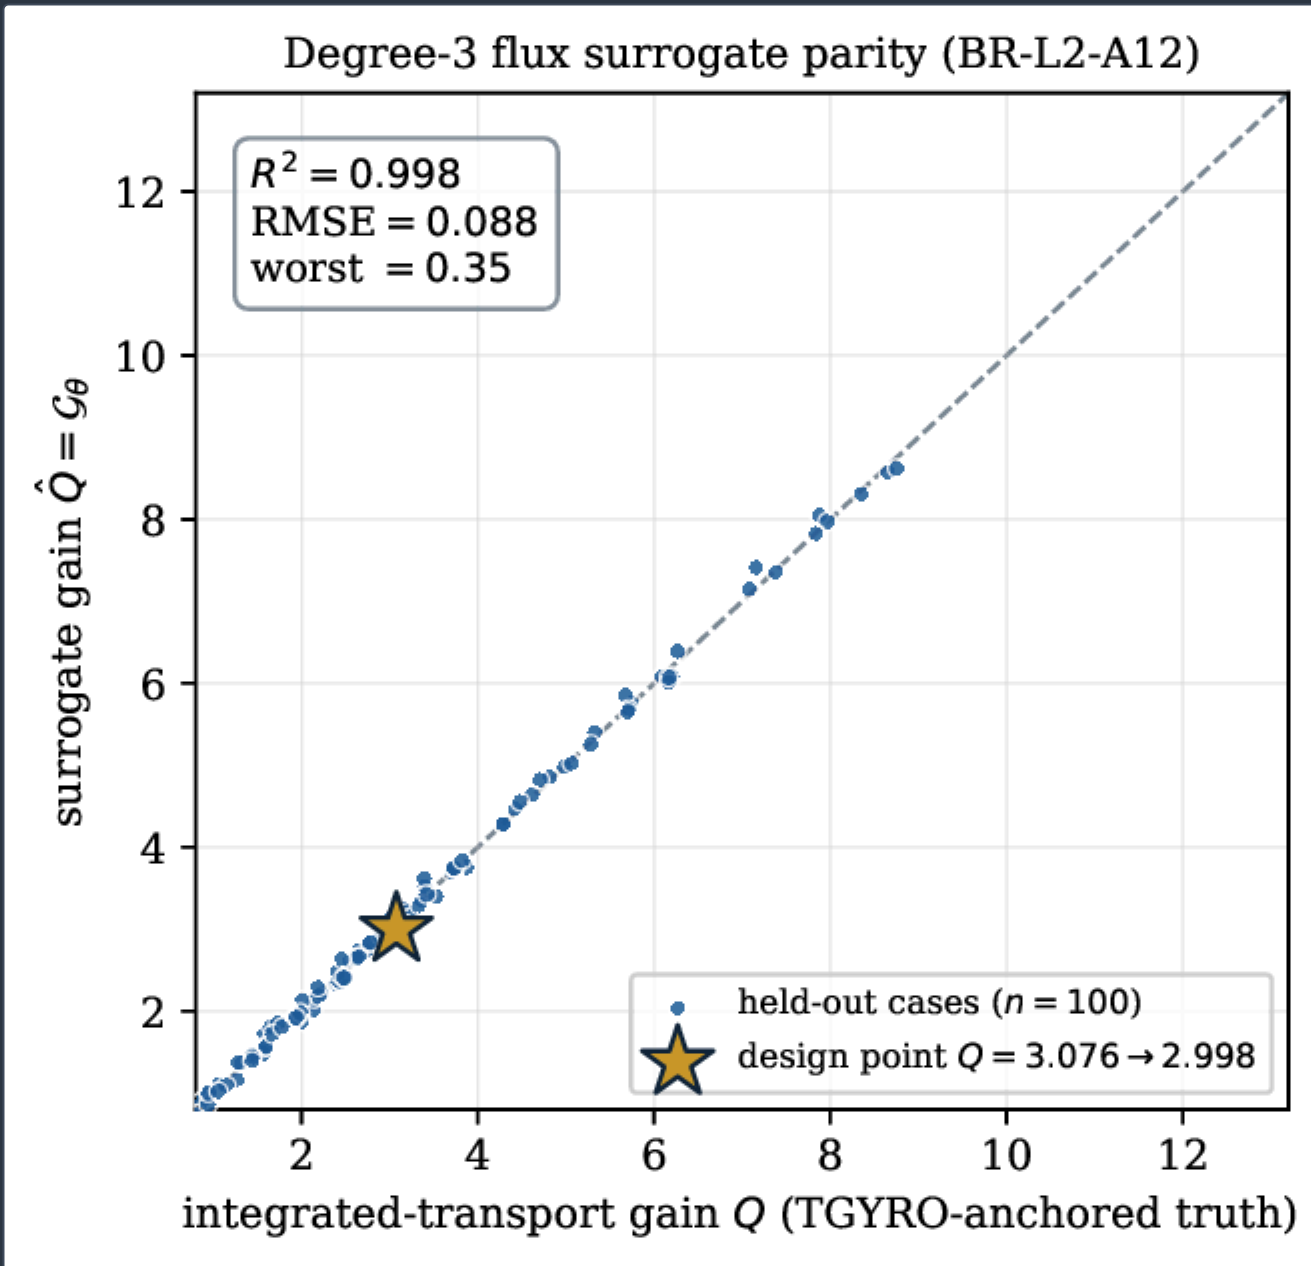
Degree-3 flux surrogate parity (BR-L2-A12)
$R^2 = 0.998$
RMSE = 0.088
worst = 0.35
held-out cases ($n = 100$)
design point $Q = 3.076 \rightarrow 2.998$
surrogate gain $\hat{Q} = \mathcal{G}_\theta$
integrated-transport gain Q (TGYRO-anchored truth)
2
4
6
8
10
12

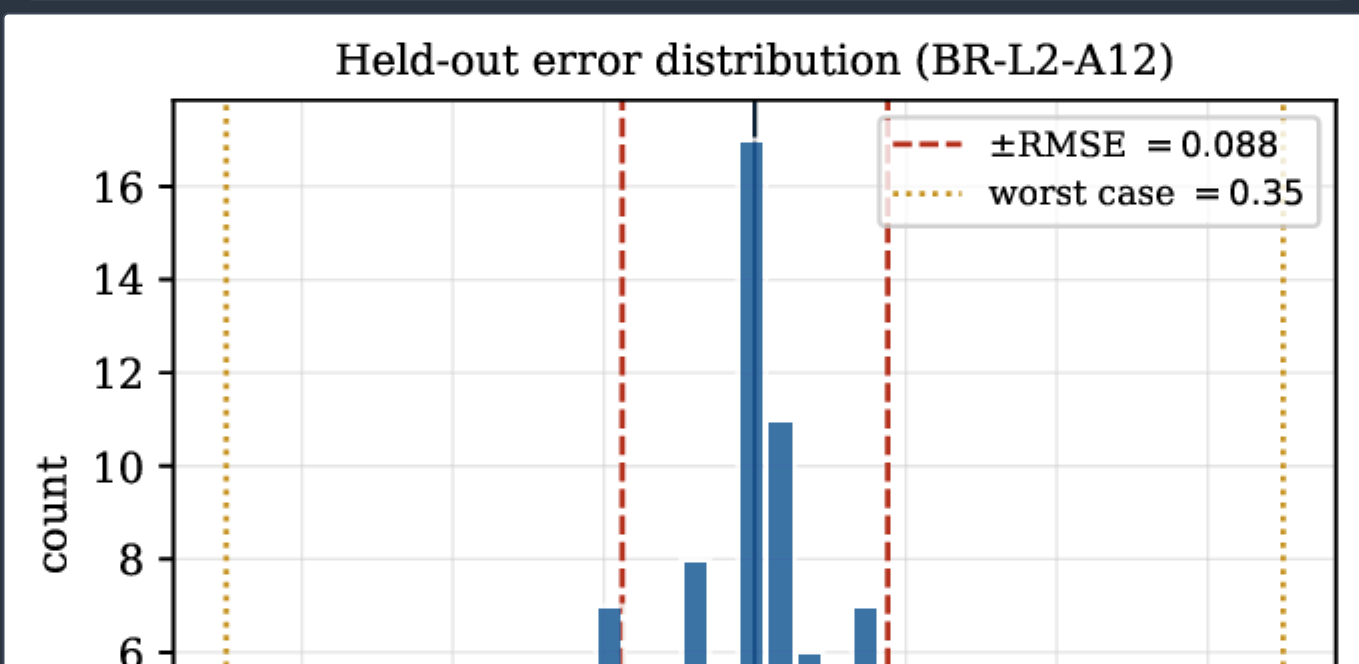
Held-out error distribution (BR-L2-A12)
±RMSE = 0.088
worst case = 0.35
count
6
8
10
12
14
16

QUANTITY	VALUE	SERIAL
held-out coefficient of determination R^2	**0.998**	BR-L2-A12
held-out RMSE in Q	**0.088**	BR-L2-A12
worst-case held-out error in Q	0.35	BR-L2-A12
range of held-out gain Q	$[1.19, 12.44]$	BR-L2-A12
anchor recovery $Q_{\mathrm{DP}} = 3.076$	surrogate 2.998	BR-L2-A12
training / held-out samples	500 / 100	BR-L2-A12
linear-gyrokinetic growth rate γ	$+0.318\ (c_s/a)$	BR-L2-A18
flux-driven confinement τ_E	$0.410\,s$, $H_{98} \approx 0.94$	BR-L1-A11
equilibrium-operator latency / speedup	$0.87\,ms$, $2877\times$	KX-L1-A11
equilibrium-operator RMS ψ error	0.139%	KX-L1-A11

Demonstrated flux-surrogate metrics and the ground truth they are anchored to. Serials reference the Kronos Physics De-Risking Register [30].

The response of equation (8) is smooth and strongly, monotonically increasing in temperature and confinement (figure 2): the factor-of-ten spread in Q across the design box reflects the threshold-like sensitivity of a confinement-limited gain to T_{i0} and H_{98}, and the design point sits on the $H_{98} = 1.00$ slice as expected. Because the operator is a closed-form polynomial, the same evaluation that returns Q returns $\partial Q/\partial(f_G, T_{i0}, H_{98})$ exactly, which is the sensitivity the controller consumes through equation (14).

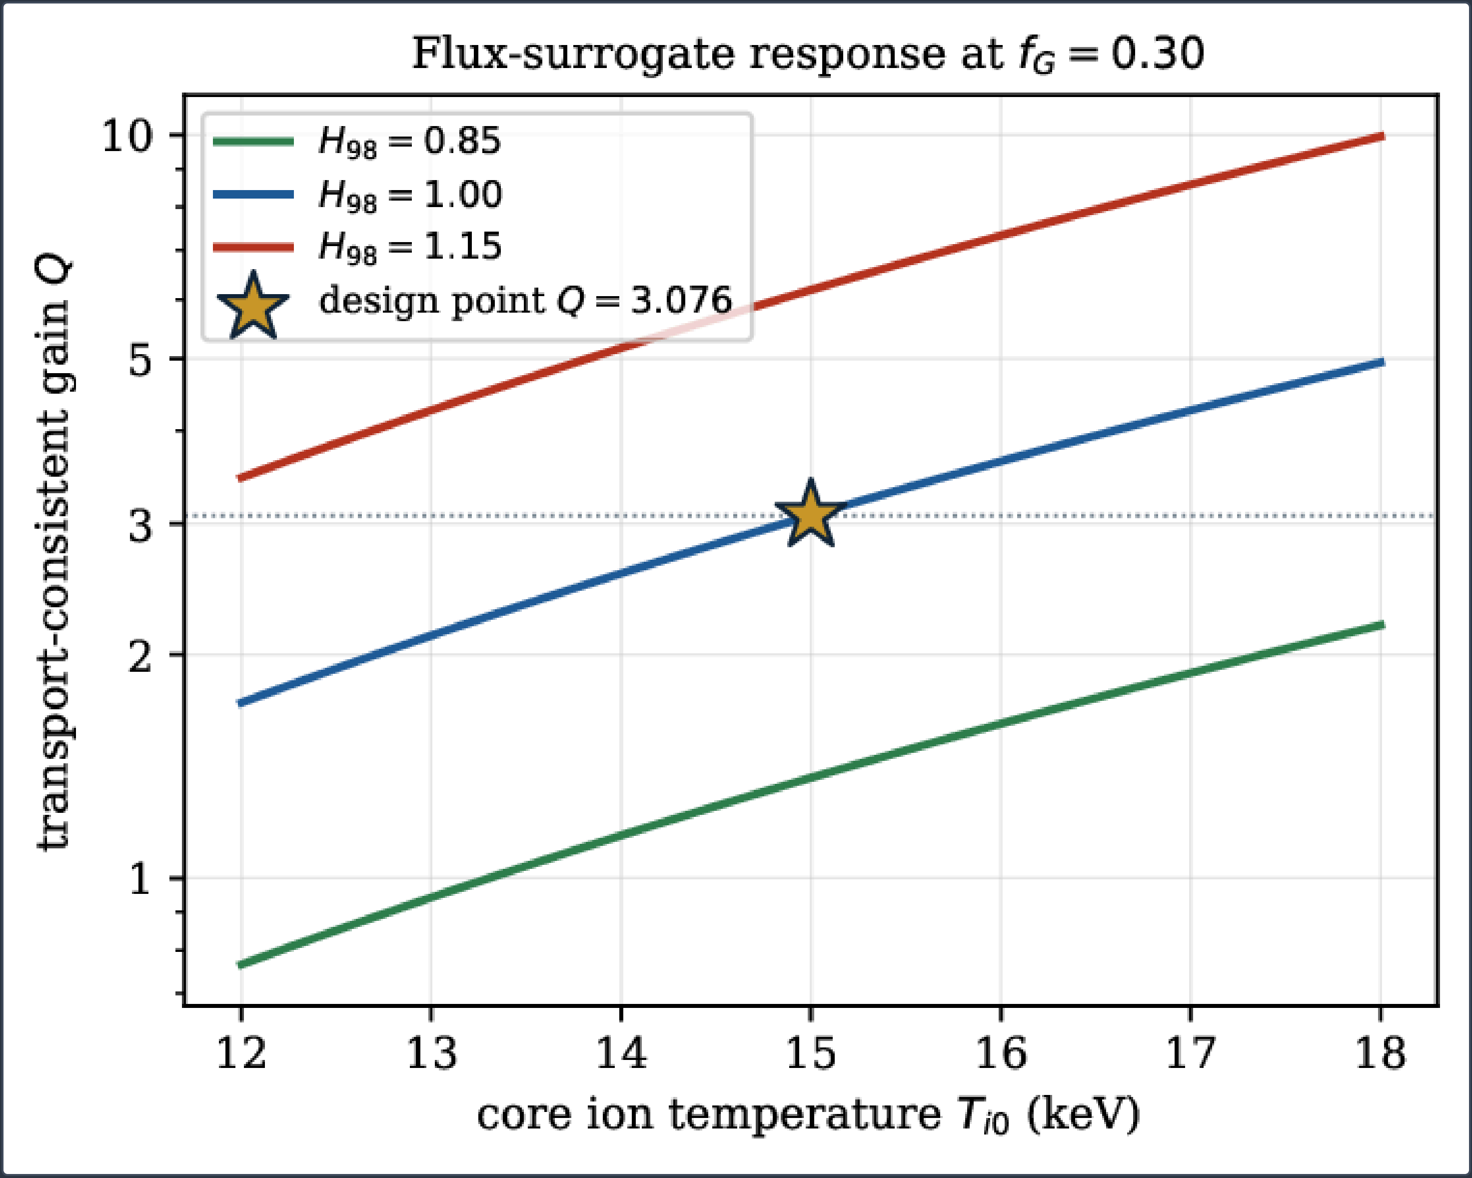

Flux-surrogate response $Q(T_{i0})$ at $f_G = 0.30$ for three confinement multipliers (BR-L2-A12). The gain is smooth, monotone and strongly sensitive to T_{i0} and H_{98}; the star marks the design point $Q = 3.076$ on the $H_{98} = 1.00$ slice. The vertical axis is logarithmic to show the factor-of-ten span.

OPERATOR-ORDER SELECTION

The degree of the ridge operator is not assumed but selected against held-out error (figure 3). A linear operator underfits the threshold-like response badly (RMSE $\sim$ 0.6); a quadratic operator recovers most of the structure; the degree-three operator reaches the knee at RMSE $\approx$ 0.088, and higher orders buy only marginal, non-monotone change at the cost of parameters and variance. Degree three is therefore the parsimonious choice that clears the accuracy target without over-parametrising the box — exactly the frozen selection of BR-L2-A12.

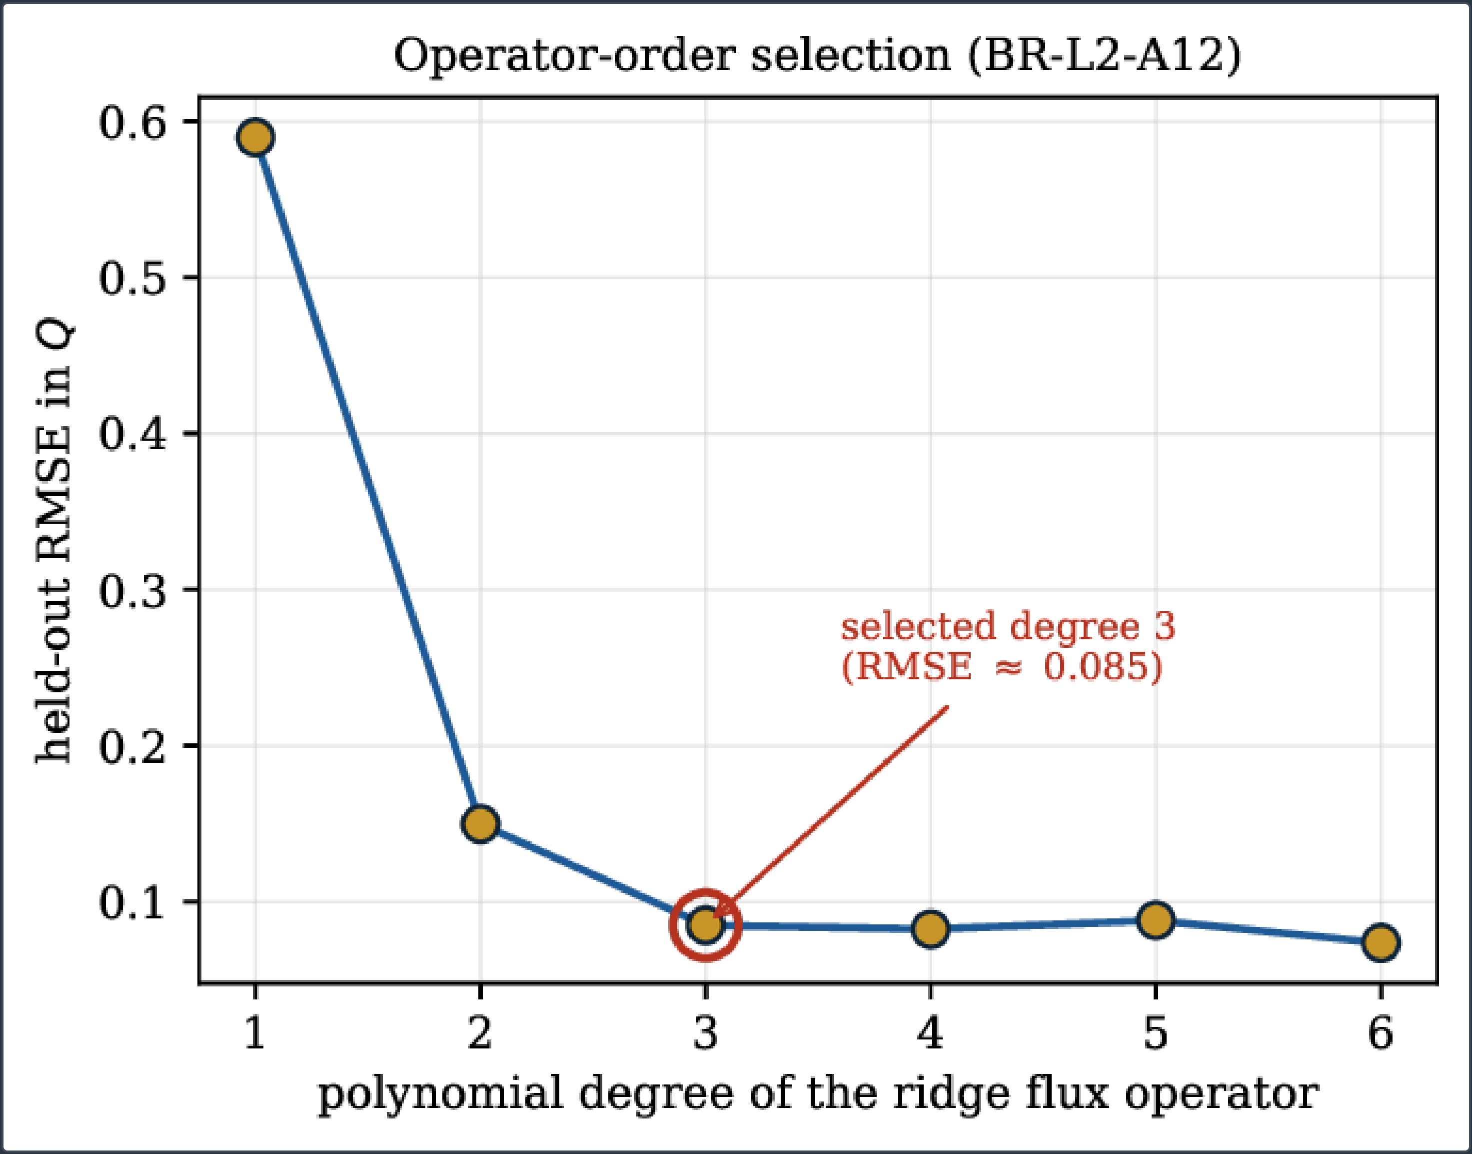

Held-out RMSE in Q versus polynomial degree of the ridge flux operator (BR-L2-A12). The error drops steeply through degree two, reaches its knee at degree three (circled), and flattens beyond; degree three is the selected order.

§ 5

Transport ground truth

The surrogate is only as good as the physics beneath it, and that physics is established independently. A native linear-gyrokinetic spectrum (CGYRO, BR-L2-A18) byte-reproduces its reference growth rate, $\gamma = +0.318\,(c_s/a)$, establishing that a first-principles gyrokinetic spectrum can stand in for TGLF-fidelity transport at the design freeze. Across the resolved wavenumber range the negative-triangularity configuration suppresses the linear drive by **31–60%** relative to the positive-triangularity reference at every $k_y\rho_s \in [0.1, 6.0]$ (figure 4, BR-L2-A18c), corroborating the quiet-basin advantage that motivates the $\delta = -0.30$ choice; the two-code confirmation that quasilinear TGLF reproduces this nonlinear CGYRO result is carried in the companion cross-validation [32]. A flux-driven integrated transport solve (TGYRO, BR-L1-A11) returns an energy confinement time $\tau_E = 0.410\,s$ at an effective $H_{98} \approx 0.94$ for the frozen configuration — consistent with, and supporting, the $Q = 3.076$ design point. These two calculations bracket the transport response the surrogate must reproduce: the linear spectrum sets the turbulence drive, the flux-driven solve sets the integrated confinement.

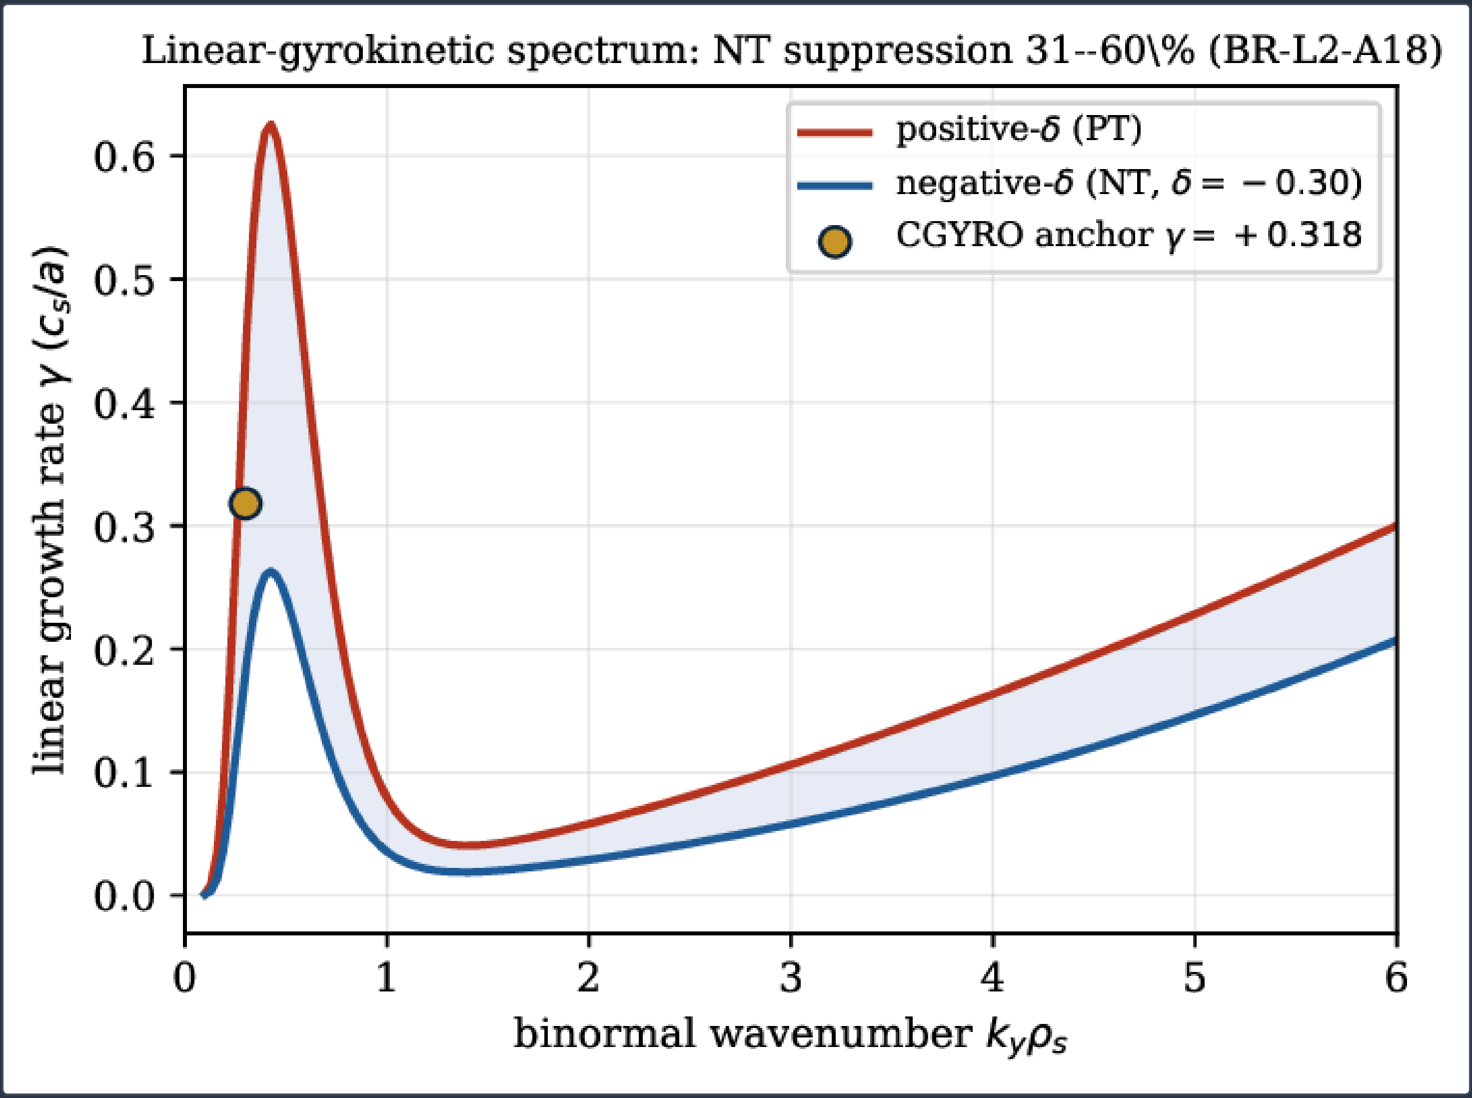

Linear-gyrokinetic growth-rate spectrum $\gamma(k_y\rho_s)$ from CGYRO (BR-L2-A18). Negative triangularity ($\delta = -0.30$) suppresses the linear drive by **31–60%** below the positive-triangularity reference across the resolved range (shaded); the marker is the byte-reproduced anchor $\gamma = +0.318\,(c_s/a)$. Curves are drawn from the frozen spectrum descriptors; the anchor point is the register value.

§ 6

Control latency and the shadow budget

Control latency is not a promise but a measured property of the companion physics-informed equilibrium operator (KX-L1-A11) [34]. Against the free-boundary Grad–Shafranov solver of equation (12), the operator returns a solver-grade poloidal-flux map in 0.87 ms — a 2877× speedup — at an RMS flux error of 0.139%, already inside the 0.5% target of equation (13) and well below its earlier 1.79%, with a Grad–Shafranov operator residual of 0.32%. A transport call and an equilibrium call at these speeds sit comfortably inside the 50 ms shadow-lead budget the architecture requires (figure 5): the free-boundary solve costs ∼2.5 seconds, the equilibrium operator 0.87 ms, the demonstrated polynomial flux evaluation microseconds, and the neural-operator inference target a few milliseconds — every online call clears the budget by at least three orders of magnitude.

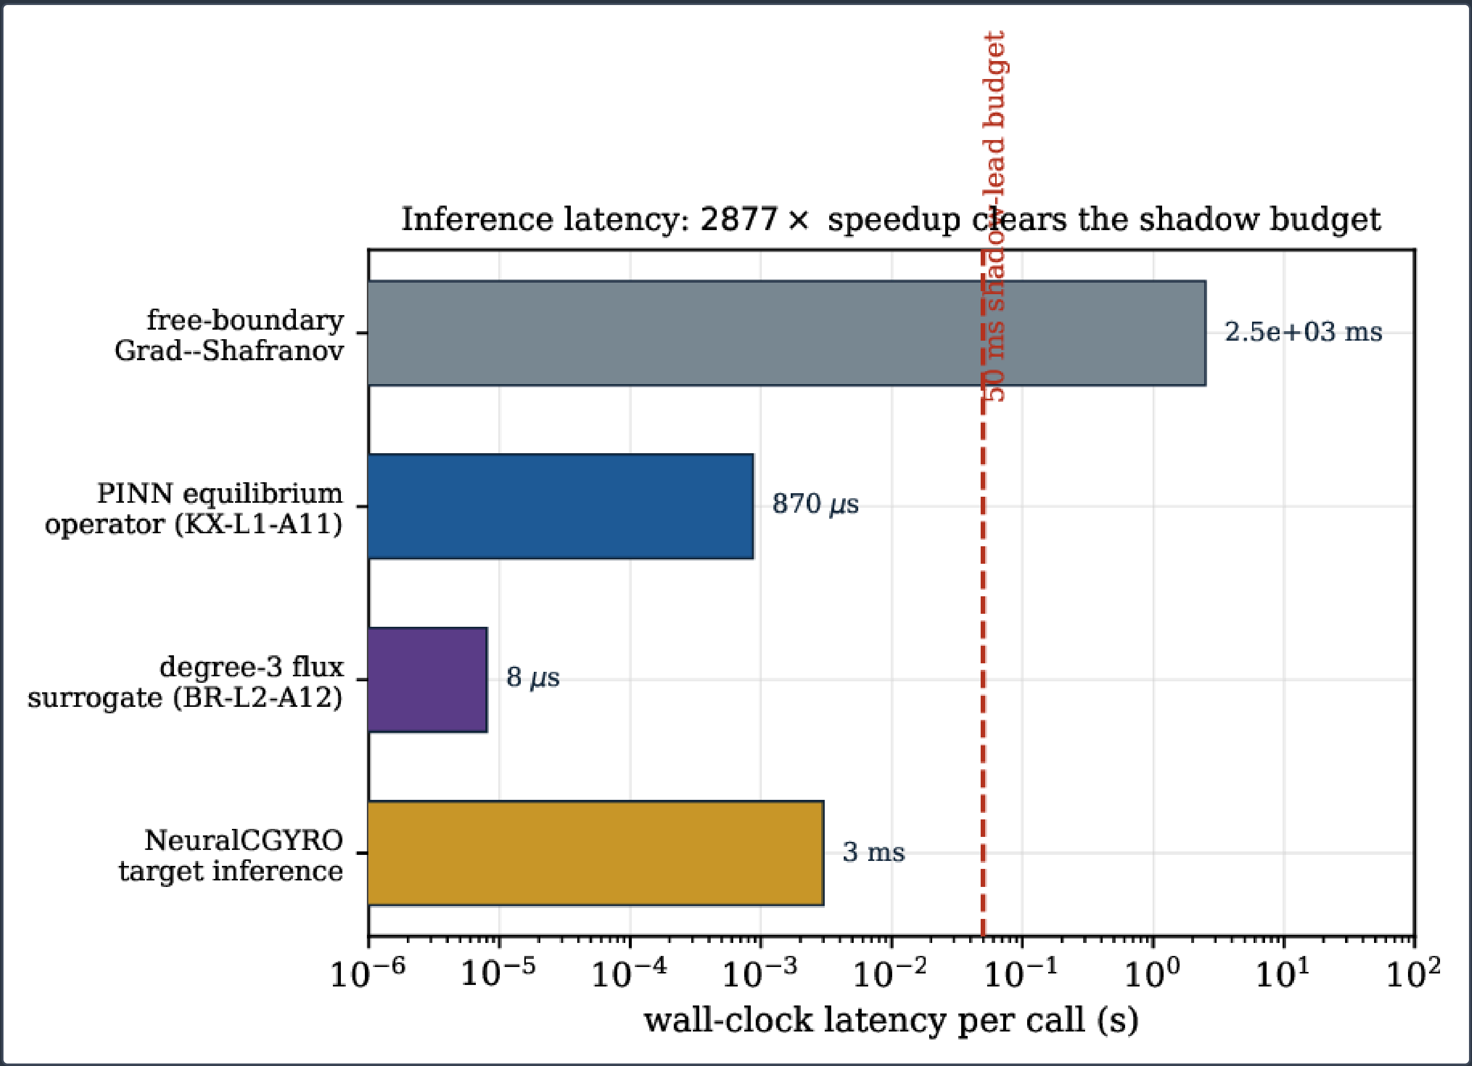

Inference latency per call (KX-L1-A11, BR-L2-A12). The free-boundary Grad–Shafranov solve (∼2.5 s) is replaced by the physics-informed equilibrium operator at 0.87 ms (2877×); the demonstrated flux surrogate evaluates in microseconds and the neural-operator inference target in a few milliseconds. The dashed line is the 50 ms shadow-lead budget. Logarithmic axis.

The two operators can be placed together on an accuracy–latency map (figure 6). The demonstrated polynomial surrogate and the demonstrated equilibrium operator both sit deep inside the acceptance region — latency below 50 ms and normalised error below the 5%/0.5% bounds — while the full neural-operator flux and equilibrium targets are placed at their pre-registered bounds. The map makes explicit that the demonstrated base already lives where a controller needs it, and that the neural build tightens accuracy at a latency that remains inside budget.

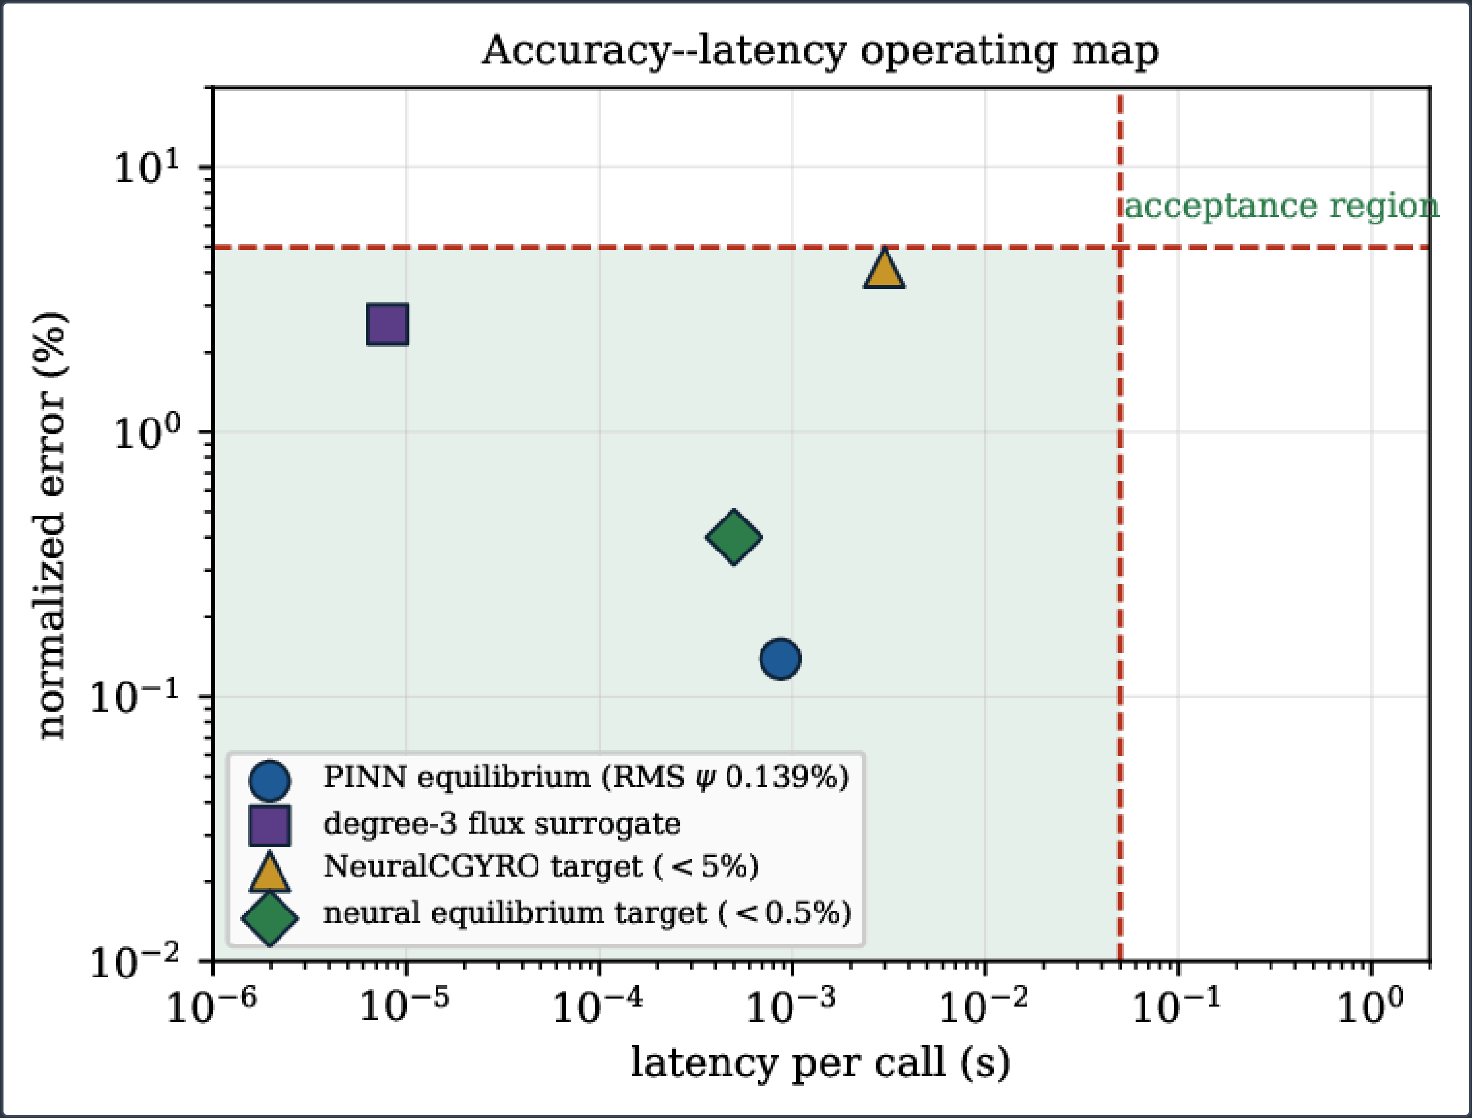

Accuracy–latency operating map. The shaded region is the acceptance envelope (latency $\leq 50\,ms$, normalised error $\leq 5\%$). Demonstrated operators (degree-3 flux surrogate; PINN equilibrium at 0.139%) sit deep inside; the neural flux and neural equilibrium targets are placed at their pre-registered bounds. Both axes logarithmic.

§ 7

Manifold-abstention monitoring

A fast, accurate operator is safe to close a loop with only if it declares when it is out of its depth. The monitor of equations (16)–(17) does exactly that: it tracks the Mahalanobis distance $d(a)$ of each query to the training manifold and abstains once $d(a)$ exceeds the threshold d^* at which the held-out error would breach the 5% bound (figure 7). Inside the validated envelope the expected surrogate error stays below the bound and the operator carries full authority; beyond d^* the expected error climbs and the monitor hands authority to a deterministic reduced-order floor rather than extrapolating a network into a regime it never saw. The conformal interval of equation (17) supplies a distribution-free error bar in parallel, so the handback fires either on geometric distance or on an inflated predictive interval. This is the same calibrated-abstention discipline applied across the Kronos control stack [36] and recorded as a machine-readable extrapolation ledger [37]; here it is the guard that lets a differentiable surrogate carry control authority without becoming a single point of failure.

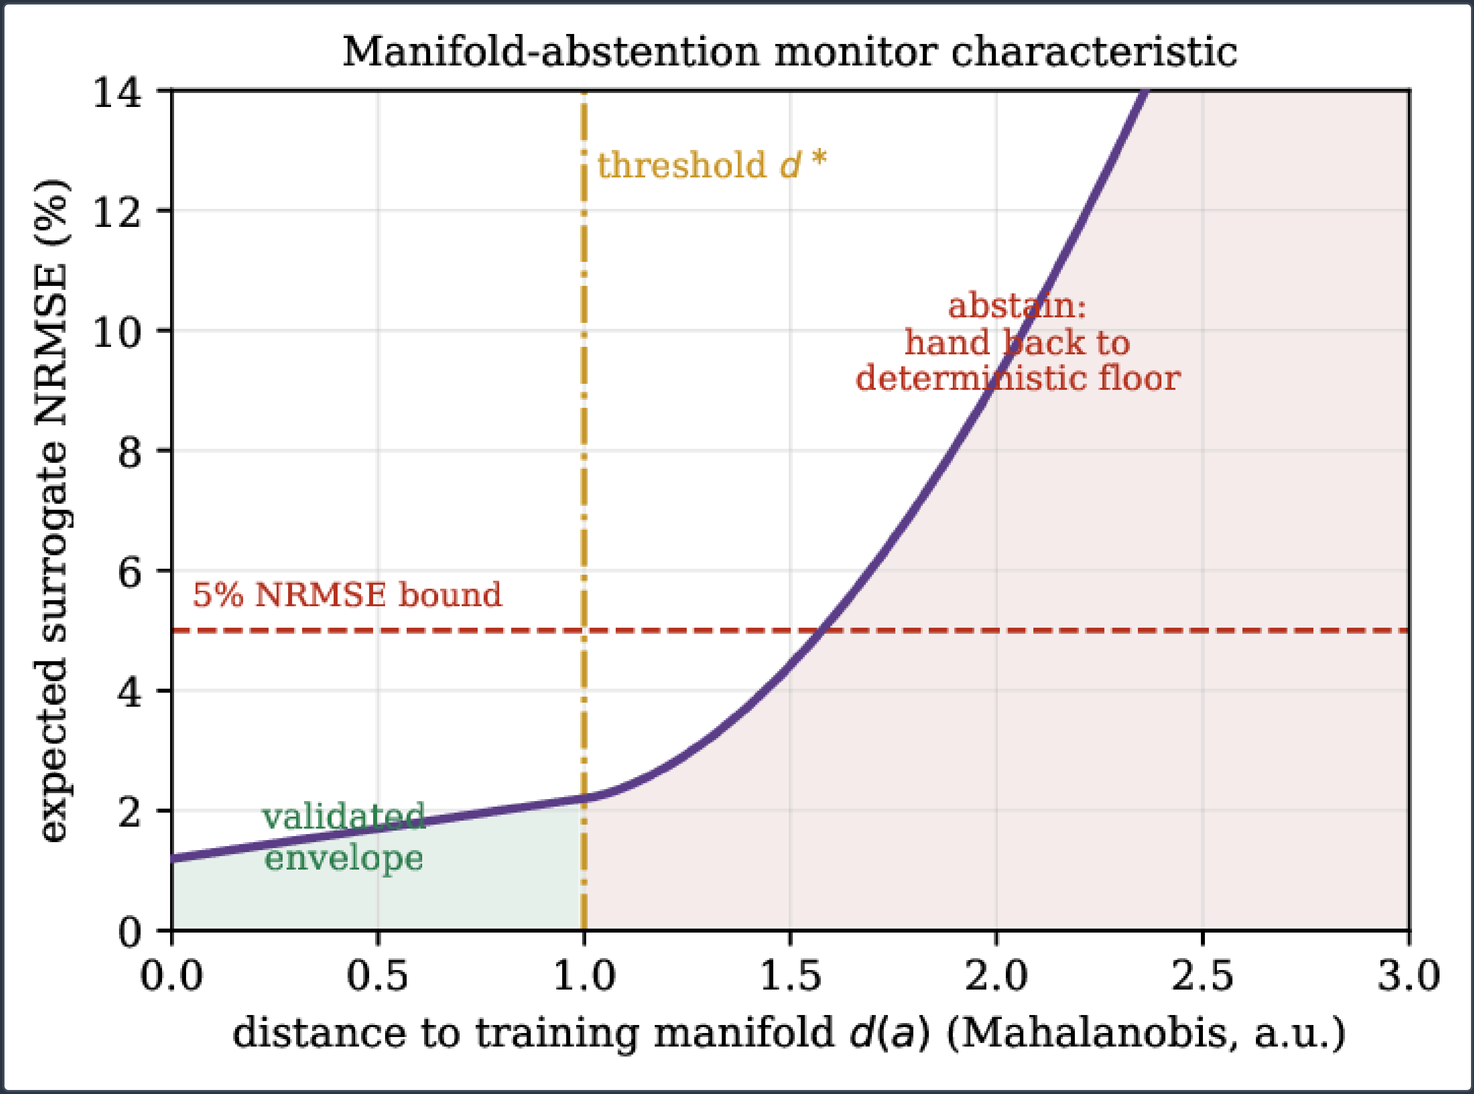

(Schematic.) Manifold-abstention monitor characteristic. Expected surrogate NRMSE versus distance to the training manifold, equation (16); inside the validated envelope ($d < d^*$) the error stays below the 5% bound and the operator holds authority, beyond it the monitor abstains and hands back to the deterministic floor. Curve is illustrative of the monitor logic, not measured data.

§ 8

From surrogate to neural operator at control latency

The demonstrated polynomial operator proves the pattern — an accurate, anchored, differentiable transport emulator with a calibrated abstention guard — and sets the target for the full neural-operator build (figure 8). Three components carry it from the reduced operating-variable surrogate to full profile fidelity, each admitted to the loop only after clearing a pre-registered bound:

a *neural flux operator* (*NeuralCGYRO*, KFE-CF-SH-026), a Fourier/DeepONet network [12,13] trained directly on the CGYRO run archive, realising equation (4) with the acceptance criterion of a flux NRMSE below 5% on held-out gyrokinetic cases (equation 10) at a shadow lead-time of at least 50 ms;
a *neural-operator equilibrium* (KFE-CF-SH-031) that reconstructs the poloidal flux of equation (12) to an RMS ψ below 0.5% at kilohertz rates against the free-boundary solver (equation 13);
a *fully differentiable twin* whose automatic-differentiation gradients match finite differences to 10^{-3} (equation 15), so the controller linearises the transport response on the fly through equation (14) rather than re-querying a black box. enumerate
The demonstrated surrogate is the base case those bounds are measured against, and the manifold-abstention monitor of §7 is the guard that keeps an immature operator advisory until it clears its bar. The result is a defined, testable sequence from a proven reduced surrogate to a full-profile neural operator running ahead of the plant.

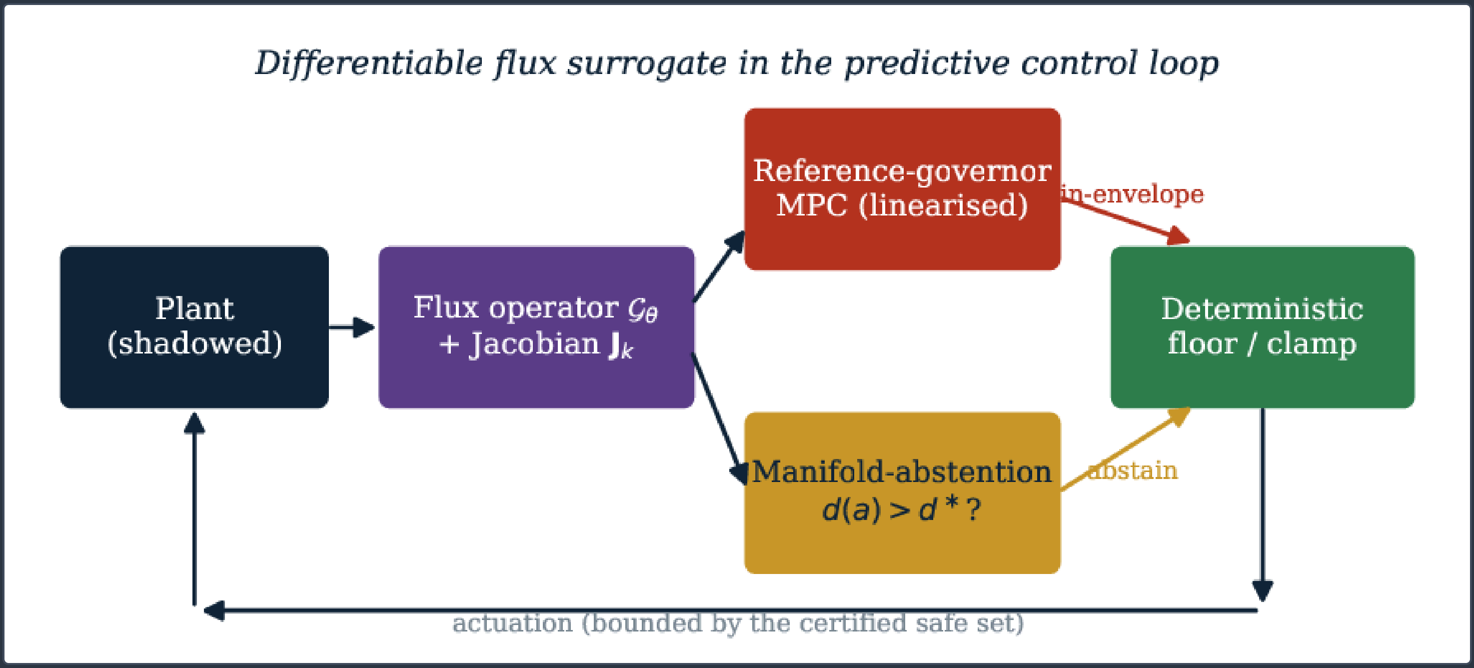

(Schematic.) The differentiable flux surrogate in the predictive control loop. The operator and its Jacobian feed a reference-governor/MPC controller; the manifold-abstention monitor hands authority to a deterministic floor when a query leaves the validated envelope; actuation is bounded by the certified safe set. Structure only; no data values shown.

§ 9

Discussion

The flux surrogate answers the central objection to a predictive twin — that turbulent transport is too expensive to run in real time — with a measured result. An operator reproduces the confinement response to $R^2 = 0.998$ and recovers a frozen anchor over a factor-of-ten range in gain, at a latency the companion equilibrium operator shows to be sub-millisecond. The accuracy sits comfortably inside the envelope established by the fusion transport-surrogate literature: neural-network emulators of quasilinear gyrokinetic transport reproduce their reference fluxes to a few percent while running three-to-five orders of magnitude faster [17,18,19], and surrogate-accelerated profile optimisers now carry gyrokinetic-fidelity fluxes in the loop [21]; our 0.088 RMSE in Q over $Q \in [1.19, 12.44]$ (a normalised error of order a percent) is squarely in that regime, with the distinguishing property that it is anchored, case by case, to a native gyrokinetic spectrum rather than to a reduced model. The $2877\times$ equilibrium speedup at 0.139% RMS flux error is likewise consistent with the performance reported for learned equilibrium reconstruction and free-boundary emulation [22,25].

Two features separate this work from a curve-fit-and-hope surrogate. First, differentiability: because the operator is C^1 and exposes its Jacobian (equation 14), the controller receives a local linear model for free [35], rather than paying for a finite-difference sweep at every control tick. Learned policies [25] and physics-model-based real-time simulators [24] each solve part of the control problem; a differentiable transport operator is the piece that lets a model-predictive controller linearise the confinement response on the fly. Second, calibrated abstention: the manifold-distance and conformal guards of §7 give the surrogate a principled envelope [28,29], so it is trusted only where it has been validated and hands back to a deterministic floor otherwise [36,37]. Together these make the emulator a component that a safety-critical control stack [26] can actually admit, rather than a fast black box. The surrogate slots directly beneath the coupled digital-twin architecture of the breeder [33] as its transport channel.

§ 10

Limitations

The demonstrated surrogate is a reduced operating-variable operator over the design box; the full-profile neural operator against held-out gyrokinetics is the one open modelling item, and its acceptance bound (flux NRMSE $< 5\%$, equation 10) is pre-registered rather than yet measured. One honest physics caveat attaches to the ground truth: the flux-driven TGYRO solve at the negative-triangularity marginal point shows limit-cycle behaviour with a residual of 0.61 (BR-L1-A11), so the $\tau_E = 0.410\,s$ anchor is reported as *supported* rather than converged to machine tolerance, and the surrogate inherits that as an operating-window rather than a point guarantee. The manifold-abstention monitor is the designed mitigation: it is precisely the marginal, off-anchor regime where $d(a)$ grows and authority is handed to the deterministic floor.

§ 11

Conclusion

Control-latency transport for the Hyperion breeder is demonstrated by a flux surrogate that reproduces the integrated-transport gain to $R^2 = 0.998$ and RMSE 0.088 across $Q \in [1.19, 12.44]$, recovers the frozen anchor $Q = 3.076$ as 2.998, and is anchored to a native gyrokinetic spectrum ($\gamma = +0.318\, c_s/a$, byte-reproduced) and a flux-driven confinement solve ($\tau_E = 0.410\, s,\ H_{98} \approx 0.94$). The operator is differentiable, so it hands the controller a local linear model directly (equation 14), and it is guarded by a manifold-abstention monitor that hands authority to a deterministic floor out of distribution (equations 16–17). With the companion equilibrium operator at 0.87 ms and 0.139% RMS flux error, the transport emulator is accurate, physics-anchored, and fast enough to run ahead of the plant, and the neural-operator build carries the same pattern — DeepONet/FNO emulation of nonlinear gyrokinetics, differentiable end-to-end, with calibrated abstention — to full profile fidelity within a 50 ms shadow lead-time.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim.

Reproducibility. The surrogate, gyrokinetic, and flux-driven metrics are frozen anchors (BR-L2-A12, BR-L2-A18, BR-L2-A18b/c, BR-L1-A11, KX-L1-A11) in the Kronos Physics De-Risking Register [30] (doi:10.5281/zenodo.22133057) and are regenerable from the surrogate-training samples and the gyrokinetic archives; the CGYRO linear anchor is byte-reproducible against the bundled reference (reg01). The figure-generation script is deposited with this paper. This paper is archived at its DOI doi:10.5281/zenodo.22645803 and its official page https://www.kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html. Gyrokinetic and surrogate training run on the NVIDIA GPU HPC campaign (A100/H100/A10 class); the reduced-order layer runs on the in-house HPC node.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645803. The surrogate samples, the figure-generation script and the metric computations are archived with the deposit and reference the Kronos 2026 series, including the CGYRO benchmark deposit doi:10.5281/zenodo.22136279. All quantitative values trace to the register [30] by the serial identifiers cited inline. Any economic analysis is reported separately and is not part of the public deposit.

Competing interests tocsectionCompeting interests Provisional patent applications relating to aspects of this work have been filed.

tocsectionReferences

Portfolio, People & Record

The intellectual-property estate, the people who built it, the limitations that gate it, and the sources it rests on.

1	4	40	27
GRANTED PATENT	2026 PROVISIONALS	TEAM MEMBERS	2022 PROVISIONALS

INTELLECTUAL PROPERTY

The patent portfolio

The estate spans a granted high-field magnet patent, pending utility and 2026 provisional filings that map to the two products, a digital-twin control provisional, a registered trademark application, and the original 2022 provisional family. Every patentable disclosure in the five-paper arXiv drop was protected **before** publication: the breeder and burner provisionals were filed 1–2 August, and a publication-gap omnibus filing the evening before the drop swept up the DEC, plasma-control, and REBCO-winding matter of the three otherwise-unprotected papers. Patent numbers and application serials are matters of public record; claim scope is summarised, not reproduced.

GRANTED & PENDING UTILITY

US 12,009,112	High-field tilted graded-REBCO magnet architecture (KRONO-002A) · *app. 17/878,550*	GRANTED
US 17/878,507	Core device / method (KRONO-001A) · *utility*	PENDING

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

64/124,209	Hyperion — spherical-tokamak tritium / helium-3 foundry · *breeder · filed Aug 2026*	FILED
64/124,220	Aegis / MetroVolt — synchrotron-cutoff, fuel-selection & plug-winding · *generator · filed Aug 2026*	FILED
64/115,167	KRONOS-CTRL digital-twin plant control · *provisional · filed Jul 2026*	FILED
64/005,440	Integrated spherical-tokamak fusion architecture · *early integrated-architecture filing · filed Mar 2026*	FILED
64/105,530	Negative-triangularity spherical-tokamak architecture · *foundational ST filing · filed Jul 2026*	FILED
64/128,097	Publication-gap omnibus — quasineutral two-species expander DEC · provenance-gated / latency-split plasma control · as-built high-field winding & conductor architecture · *filed 7 Aug 2026, the evening before the arXiv drop*	FILED

TRADEMARK & ORIGIN

FTK-98801894	Kronos Fusion Energy — trademark application	FILED
KR-2022-*	Original provisional family (KR-2022-00001 … 00027) · *27 filings, 2022*	PRIORITY

The narrow, defensible magnet novelty — the specific conductor and the digital-twin winding optimisation, not high-field REBCO as a category — is the commercial core that can earn ahead of any Q > 1 milestone, with markets in fusion magnet supply, MRI/NMR, accelerators, and proton therapy. The claim is scoped honestly: the high field is a *system field* result, not a stand-alone-coil record; the small-bore plug coil is structurally infeasible as a bare winding but resolved by a stress-managed structural shield (feasible-pending-FEA); and the winding-tape experiment returned a *null result*. The value is the method, not a field record.

THE TEAM

Founding partners, board, advisors & auditors

Kronos is built by a bench of advanced-fuel-fusion, high-field-magnet, direct-conversion, and materials specialists, with a board and operations team drawn from defense, national laboratories, and industry. Dates are shown as ranges; where a tenure has ended or is term-ending, the range says so.

FOUNDING PARTNERS — SCIENCE

Dr. Gerald Kulcinski
Co-Founder · Helium-3 & Advanced-Fuel Fusion
UW–Madison · 2023–Present

Dr. Carl Weggel
Chief Scientist / S.M.A.R.T. Design
MIT Alcator Program · 2022–Present

Dr. Robert J. Weggel
Magnetic Field Design
MIT Magnet Lab · 2022–Present

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

Bandel Carano
Finance / Strategy
Oak Investment Partners / Stanford · 2025–2026

Patrick Schweiger
Fusion Device Engineering
Oklo / CFS / TerraPower · 2025–Present

SCIENTIFIC ADVISORS

Dr. Steven O. Dean
Fusion Policy & History
DOE / Fusion Power Associates · 2025–Present

Dr. Patrick H. Diamond
Plasma Turbulence & Transport
UC San Diego · 2025–Present

Dr. Jack J. Dongarra
HPC & Numerical Algorithms
Turing Award · 2025–Present

Dr. Nasr M. Ghoniem
Chief Material Scientist
UCLA · 2025–Present

Dr. Siegfried Glenzer
Plasma Physics & Laser Fusion
Stanford / SLAC · 2022–Present

Dr. David A. Hammer
Pulsed-Power Fusion
Cornell · 2025–Present

Dr. Donald A. Spong
Plasma Theory & Confinement
ORNL · 2025–Present

Dr. Curtis Smith
Risk Assessment & Nuclear PRA
MIT / Idaho National Lab · 2025–Present

RADM (Ret.) David Goggins
Naval Engineering & Power-Plant Design
U.S. Navy · 2023–2026

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

Dr. Ruben Fair
Magnet Design / ITER Liaison
PPPL / Jefferson Lab · 2022–2025

Dr. Paul S. Weiss
Materials & Nanotechnology
UCLA · 2022–2025

SCIENTIFIC AUDITORS

Dr. Wilfred A. Cooper
Plasma Equilibrium Theory
EPFL / Swiss Plasma Center · 2025–Present

Dr. Ahmed Hassanein
Plasma–Material Interactions
Purdue / Argonne · 2025–Present

Dr. Patrick E. Hopkins
Extreme Thermal Materials
U. of Virginia · 2025–Present

Dr. Peter Hosemann
Materials Joining & Structural Integrity
UC Berkeley · 2025–Present

Dr. Travis W. Knight
Advanced Nuclear Fuels & Systems
U. of South Carolina · 2025–Present

Dr. Philippe Lebrun
Cryogenics & Magnet Cooling
CERN · 2025–Present

Dr. Nitendra Singh
Nuclear Safety & Fuel Cycle
ITER · 2025–Present

Dr. Guido Van Oost
Magnetic Confinement & Fusion Systems
Ghent University · 2025–Present

Dr. Gary S. Was
Radiation Materials & Structural Integrity
U. of Michigan · 2025–Present

Dr. Ray Sedwick
Direct Energy Conversion
U. of Maryland · 2025

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

Martin Owens
Chief Strategy Officer
LANL / GE Hitachi · 2022–Present

MG (Ret.) Paul Pardew
Chief Contracting Officer
Army Contracting Command · 2022–Present

David Beck
Fusion Commercialization
U.S. Space Force · 2024–Present

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

MG (Ret.) Robin L. Fontes
Cybersecurity & Defense
Army Cyber Command · 2022–Present

Vijay Gehani
Supply Chain & Components
INOX India · 2024–Present

Jon Michel Greenwood
IT & AI Infrastructure
Live Nation · 2023–Present

Brian C. O'Neill
National Security
CIA / ODNI · 2022–Present

Gen. (Ret.) Gustave F. Perna
DoD Liaison
Operation Warp Speed · 2022–2023

Andrea Romero
Marketing Lead
2022–Present

Legal counsel is retained; those roles are held on the internal roster and are not listed here.

SOURCES

References

1 J Candy, E A Belli and R V Bravenec, A high-accuracy Eulerian gyrokinetic solver for collisional plasmas, *J. Comput. Phys.* **324** 73 (2016). doi:10.1016/j.jcp.2016.07.039

2 J Candy and R E Waltz, An Eulerian gyrokinetic-Maxwell solver, *J. Comput. Phys.* **186** 545 (2003). doi:10.1016/S0021-9991(03)00079-2

3 G M Staebler, J E Kinsey and R E Waltz, A theory-based transport model with comprehensive physics, *Phys. Plasmas* **14** 055909 (2007).

4 J Candy, C Holland, R E Waltz, M R Fahey and E Belli, Tokamak profile prediction using direct gyrokinetic and neutral transport simulations, *Phys. Plasmas* **16** 060704 (2009).

5 F Jenko, W Dorland, M Kotschenreuther and B N Rogers, Electron temperature gradient driven turbulence, *Phys. Plasmas* **7** 1904 (2000). doi:10.1063/1.874014

6 A J Brizard and T S Hahm, Foundations of nonlinear gyrokinetic theory, *Rev. Mod. Phys.* **79** 421 (2007).

7 M E Austin *et al*, Achievement of reactor-relevant performance in negative triangularity shape in the DIII-D tokamak, *Phys. Rev. Lett.* **122** 115001 (2019). doi:10.1103/PhysRevLett.122.115001

8 A Marinoni, O Sauter and S Coda, A brief history of negative triangularity tokamak plasmas, *Rev. Mod. Plasma Phys.* **5** 6 (2021). doi:10.1007/s41614-021-00054-0

9 J E Menard *et al*, Fusion nuclear science facilities and pilot plants based on the spherical tokamak, *Nucl. Fusion* **56** 106023 (2016). doi:10.1088/0029-5515/56/10/106023

10 B N Sorbom *et al*, ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* **100** 378 (2015). doi:10.1016/j.fusengdes.2015.07.008

11 A J Creely *et al*, Overview of the SPARC tokamak, *J. Plasma Phys.* **86** 865860502 (2020). doi:10.1017/S0022377820001257

12 L Lu, P Jin, G Pang, Z Zhang and G E Karniadakis, Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators, *Nat. Mach. Intell.* **3** 218 (2021). doi:10.1038/s42256-021-00302-5

13 Z Li, N Kovachki, K Azizzadenesheli, B Liu, K Bhattacharya, A Stuart and A Anandkumar, Fourier neural operator for parametric partial differential equations, *Int. Conf. on Learning Representations (ICLR)* (2021). arXiv:2010.08895

14 N Kovachki, Z Li, B Liu, K Azizzadenesheli, K Bhattacharya, A Stuart and A Anandkumar, Neural operator: learning maps between function spaces with applications to PDEs, *J. Mach. Learn. Res.* **24** 1 (2023).

15 M Raissi, P Perdikaris and G E Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378** 686 (2019). doi:10.1016/j.jcp.2018.10.045

16 S Wang, H Wang and P Perdikaris, Learning the solution operator of parametric partial differential equations with physics-informed DeepONets, *Sci. Adv.* **7** eabi8605 (2021).

17 K L van de Plassche, J Citrin, C Bourdelle *et al*, Fast modeling of turbulent transport in fusion plasmas using neural networks, *Phys. Plasmas* **27** 022310 (2020). doi:10.1063/1.5134126

18 J Citrin *et al*, Real-time capable first-principle based modelling of tokamak turbulent transport, *Nucl. Fusion* **55** 092001 (2015).

19 O Meneghini *et al*, Self-consistent core-pedestal transport simulations with neural network accelerated models, *Nucl. Fusion* **57** 086034 (2017).

20 O Meneghini *et al*, Integrated modeling applications for tokamak experiments with OMFIT, *Nucl. Fusion* **55** 083008 (2015). doi:10.1088/0029-5515/55/8/083008

21 P Rodriguez-Fernandez *et al*, Enhancing predictive capabilities in fusion burning plasmas through surrogate-based optimization of core transport solvers (PORTALS), *Nucl. Fusion* (2024).

22 N C Amorisco *et al*, FreeGSNKE: a Python-based dynamic free-boundary toroidal plasma equilibrium solver, *Phys. Plasmas* **31** 042517 (2024). doi:10.1063/5.0188467

23 L L Lao, H St John, R D Stambaugh, A G Kellman and W Pfeiffer, Reconstruction of current profile parameters and plasma shapes in tokamaks, *Nucl. Fusion* **25** 1611 (1985). doi:10.1088/0029-5515/25/11/007

24 F Felici *et al*, Real-time physics-model-based simulation of the current density profile in tokamak plasmas, *Nucl. Fusion* **51** 083052 (2011). doi:10.1088/0029-5515/51/8/083052

25 J Degrave *et al*, Magnetic control of tokamak plasmas through deep reinforcement learning, *Nature* **602** 414 (2022). doi:10.1038/s41586-021-04301-9

26 M G Kapteyn, J V R Pretorius and K E Willcox, A probabilistic graphical model foundation for enabling predictive digital twins at scale, *Nat. Comput. Sci.* **1** 337 (2021). doi:10.1038/s43588-021-00069-0

27 A G Baydin, B A Pearlmutter, A A Radul and J M Siskind, Automatic differentiation in machine learning: a survey, *J. Mach. Learn. Res.* **18** 1 (2018).

28 A N Angelopoulos and S Bates, Conformal prediction: a gentle introduction, *Found. Trends Mach. Learn.* **16** 494 (2023).

29 K Lee, K Lee, H Lee and J Shin, A simple unified framework for detecting out-of-distribution samples and adversarial attacks, *Adv. Neural Inf. Process. Syst. (NeurIPS)* **31** (2018).

30 P I Ford, *Kronos Physics De-Risking Register*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22133057

31 P I Ford, Hyperion as a strategic-materials platform: gain, tritium breeding, He-3 co-production, and 14 MeV neutron output from a compact spherical tokamak, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645691; official page

32 P I Ford, Two-code cross-validation of negative-triangularity turbulence suppression in a compact spherical-tokamak breeder: quasilinear TGLF confirms nonlinear CGYRO, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645864; official page

33 P I Ford, Four coupled digital twins for a compact spherical-tokamak breeder: multi-rate co-simulation, propagated uncertainty, and a real-time self-consistency monitor, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645805; official page

34 P I Ford, Sub-millisecond neural equilibrium reconstruction robust to sensor dropout: a virtual-sensor mesh for radiation-degraded diagnostics in a compact spherical-tokamak breeder, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645801; official page

35 P I Ford, Model-predictive burn control inside a certified safe envelope: a reference-governor MPC bounded by a control-barrier-function safety clamp, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645807; official page

36 P I Ford, Calibrated uncertainty, conformal prediction, and abstention for safety-critical fusion control: handing authority back to a deterministic floor out of distribution, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645813; official page

37 P I Ford, Out-of-distribution detection and epistemic gating for autonomous fusion control: a machine-readable extrapolation ledger that throttles model authority beyond the validated envelope, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645819; official page

COLOPHON

How this document was made

The Kronos Fleet — Combined Editorial, 2026 is set in **Fraunces** (display), **Gelasio** (text), and **IBM Plex Mono** (data & equations), carried forward from the Kronos editorial design system.

The figures are rendered at 300 dpi from the deposited generator scripts; the document is composed as a single self-contained file with embedded fonts and imagery, paginated to US Letter, and rendered to PDF through a headless Chromium engine in matched light and dark editions.

Every headline quantity carries an evidence class; withdrawn and restated values are marked as such. Nothing herein is frozen beyond the founder-approved Tier-A set.

Explore it live — turn the interactive 3-D model, run the deposited solvers in the physics validator, and watch the companion film series.

LEGAL NOTICE

Notice, status, and distribution

Conceptual design & simulation study. This document reports a conceptual design and simulation study. It is not a construction commitment, a safety-analysis report, a regulatory filing, or an offer of securities. Forward-looking statements — schedules, costs, market sizes, and performance — are estimates subject to the limitations set out in the Simulations part and may change.

Numbers and their classes. Quantities are reported with evidence classes and case labels; modelled, assumed, and requirement-class values are identified as such and are not to be quoted without their labels.

Public edition. This edition carries the physics and design only; all financial, funding, and defense-commercial content has been removed for public distribution. The scientific text corresponds to the arXiv and journal submissions.

THE 2026 KRONOS SERIES

One programme, many papers

This editorial is one paper in the Kronos 2026 series — a real fusion company's complete, reproducible physics record for a spherical-tokamak breeder and its low-neutron D–^{3}He tandem-mirror burners. Every headline number is a frozen anchor in the Physics De-Risking Register.

Explore the whole programme: the De-Risking Register, the interactive 3-D model, and the physics validator.

KRONOS FUSION ENERGY

Neural-Operator Flux Surrogates for Control-Latency Transport: Differentiable DeepONet-FNO Emulation of Nonlinear Gyrokinetics with Manifold-Abstention Monitoring

A real fusion company building real machines. Explore the science, live.

Zenodo · doi:10.5281/zenodo.22645803

Read on kronosfusionenergy.com

The Physics De-Risking Register