KRONOS FUSION ENERGY

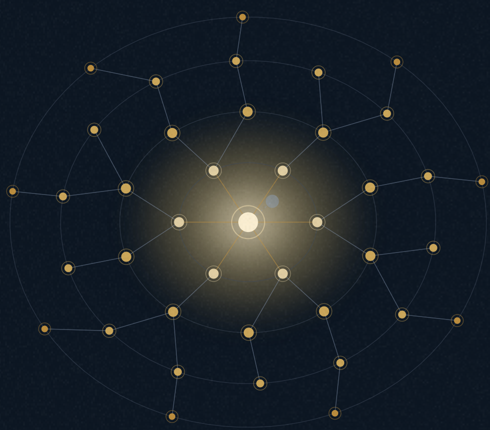

PAPER 2.71 · THE KRONOS 2026 SERIES

Sub-Millisecond Neural Equilibrium Reconstruction Robust to Sensor Dropout: A Virtual-Sensor Mesh for Radiation-Degraded Diagnostics in a Compact Spherical-Tokamak Breeder

Every control decision in a compact spherical tokamak begins with the equilibrium, and in a nuclear device that equilibrium must be reconstructed faster than the plant, verified...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Sub-Millisecond Neural Equilibrium Reconstruction Robust to Sensor Dropout: A Virtual-Sensor Mesh for Radiation-Degraded Diagnostics in a Compact Spherical-Tokamak Breeder

Every control decision in a compact spherical tokamak begins with the equilibrium, and in a nuclear device that equilibrium must be reconstructed faster than the plant, verified against the governing physics, and kept trustworthy when a diagnostic degrades or drops out under a neutron field. We formalise and demonstrate a real-time reconstruction stack for the Hyperion breeder (design point $I_p = 9.66\,MA$, $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, negative triangularity $\delta = -0.30$, $B_0 = 8\,T$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,m$, aspect ratio $A = 2.5$, centrepost peak field $16.84\,T$). A physics-informed neural surrogate reconstructs the poloidal flux $\psi(R, Z)$ to a root-mean- square error of $\mathbf{0.139\%}$ of range against a free-boundary Grad–Shafranov reference — inside the 1% acceptance target and improving a prior 1.79% attempt — at a $\mathbf{2877}\times$ speedup ($0.87\,ms$ inference versus a $\sim 2.5\,s$ Picard solve), with a discretised Grad–Shafranov operator residual of 0.32% on the interior confirming the reconstruction is a solution of the governing equation and not a free-form fit. A graph-neural virtual-sensor mesh assimilates the diagnostic set into a plasma-state estimate and, in the same forward pass, discriminates disruption and quench precursors at an area under the receiver-operating characteristic (ROC) of $\mathbf{0.980}$, above any single channel; a quadrature fusion of an uncorrelated strain-rate and magnetization pair reaches $\mathbf{0.992}$ and a true-positive rate of 0.852 at a 1% false-positive rate, independently reproduced under a separate implementation to five significant figures. The mesh is architected for graceful degradation: an in-winding nitrogen-vacancy (NV) diamond magnetometer at a shot-noise-limited 0.189 $\mathrm{pT\,Hz}^{-1/2}$ $(125\,\mathrm{Hertz}$ $bandwidth) supplies a local vector channel that does not share a failure mode with the plasma-facing optics, and a model-fusion backup holds precursor AUC$ \ge0.95$with the NV channel degraded by up to$39×$, so the protective function survives radiation-induced sensor loss. Each estimator is admitted to control authority only after clearing an explicit acceptance inequality. The result is a fast, physics-consistent, dropout-tolerant reconstruction with a demonstrated core and a defined, testable path to full diagnostic coverage.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645801 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: equilibrium reconstruction · physics-informed neural network · graph neural network · virtual-sensor mesh · sensor dropout · radiation-degraded diagnostics · NV-diamond magnetometry · 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)

equilibrium reconstruction; physics-informed neural network; graph neural network; virtual-sensor mesh; sensor dropout; radiation-degraded diagnostics; NV-diamond magnetometry; spherical-tokamak breeder

§ 1

Introduction

THE EQUILIBRIUM IS THE ROOT OF THE CONTROL PROBLEM

Everything a plasma controller does — shape, position, vertical stabilisation, profile and burn regulation — is defined relative to the magnetohydrodynamic (MHD) equilibrium: where the plasma sits, how it is shaped, and how close it runs to a limit. The controller therefore needs, at every tick, an estimate of the poloidal flux function $\psi(R, Z)$ whose level sets are the flux surfaces. Producing that estimate is the classical problem of real-time equilibrium reconstruction, solved for four decades by parametrised fits to the Grad–Shafranov equation constrained by magnetic measurements, of which the EFIT family is the archetype [1]. Reconstruction feeds shape and position control [2] and, increasingly, physics-model-based profile controllers that run inside the cycle [3].

Two forces now stress this classical pipeline. First, machines are getting smaller and faster. A compact spherical tokamak (ST) confines a large plasma current in a small volume, which makes it an attractive neutron source and breeder platform [4,5,6] and compresses the timescales the controller must beat. The Hyperion breeder holds $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$, behind a centre-post magnet at a peak field of $16.84\,T$, in a strongly shaped ($\kappa = 2.0$), negative-triangularity ($\delta = -0.30$) configuration chosen to access good core confinement without the H-mode edge pedestal and its edge-localised-mode cycle [7,8]. That quiet edge is an operational asset only if the controller holds the plasma inside the window that keeps it quiet, which requires an equilibrium it can trust in well under a millisecond; a free-boundary Grad–Shafranov solve is accurate but far too slow for a kilohertz loop, so the reconstruction must be a fast *surrogate* that remains faithful to the governing equation.

RADIATION-DEGRADED DIAGNOSTICS AND THE DROPOUT PROBLEM

The second force is specific to a nuclear device and is the organising concern of this paper. Reconstruction and state estimation are only as trustworthy as their worst diagnostic day. In a breeder the plasma-facing sensors sit in a 14 MeV neutron field; optical windows darken, magnetic probes drift and integrate error over a pulse, and in-vessel components degrade over their service life. A reconstruction that trusts any single sensor is a single point of failure in a system where every actuation carries a safety consequence, and channel loss — partial degradation or outright dropout — is not an anomaly to be handled after the fact but a design condition to be engineered for from the start. The requirement is therefore not merely a fast equilibrium but a fast equilibrium *plus* a state estimate that degrades gracefully as channels are lost, with the loss of any one channel left to be constrained by the others.

PRIOR ART IN LEARNED RECONSTRUCTION

Machine learning has entered each piece of this problem. Deep networks emulate the Grad–Shafranov map at control latency: neural solvers constrained by measured magnetic signals reconstruct equilibria orders of magnitude faster than an iterative solve [9], and learning has been used to extend and accelerate EFIT-class reconstruction [10]. Learned controllers now act in the loop — Degrave *et al.* trained a deep reinforcement-learning agent to command the TCV coil voltages, producing a variety of shapes including negative triangularity, directly from magnetics [11] — and deep networks forecast disruptive instabilities across machines [12]. The enabling primitives are mature: physics-informed neural networks embed a PDE residual in the training loss [13]; neural operators such as DeepONet [14] and the Fourier neural operator [15] learn maps between function spaces; and graph neural networks propagate information over irregular sensor topologies through message passing [17,18,19]. In fusion, neural- network surrogates of quasilinear gyrokinetic transport already run three-to-five orders of magnitude faster than the code they replace [20,21], and probabilistic digital-twin formalisms carry calibrated uncertainty end-to-end [22].

THE GAP AND OUR CONTRIBUTION

What is missing for a compact nuclear ST is a single reconstruction stack that (i) delivers a solver-grade equilibrium at control latency *and* certifies its own physical consistency against the governing equation; (ii) fuses the diagnostic set into a state estimate over the sensor graph so that a dropped channel is reconstructed from its neighbours; (iii) survives the radiation-driven degradation of individual sensors with a quantified retention floor; and (iv) admits each estimator to control authority only after it clears an explicit, pre-registered acceptance test. This paper contributes that stack. We give the governing equations and the formal forward/inverse problem (§3); the numerical formulation — discretisation of the Grad–Shafranov operator, the free-boundary reference, the network and graph-estimator architectures, and convergence (§4); the NVIDIA GPU HPC campaign that produced every number, with the codes named and the reproducibility protocol stated (§5); the demonstrated results with eleven figures and supporting tables (§6); a quantitative comparison to the published equilibrium-reconstruction, quench-detection and quantum-sensing literature (§7); one honest limitation (§8); and the conclusion. Every quantitative claim traces to a frozen anchor in the Kronos Physics De-Risking Register [39], cited inline by its serial identifier (e.g. KX-L1-A11) so each number is auditable. The stack sits inside the broader Kronos control programme: the AI and real-time-control foundations [aiq; kronosfusionenergy.com/publications/ai-quantum-control.html], the free-boundary equilibrium and ideal-MHD stability case it reconstructs against [c27; kronosfusionenergy.com/publications/ideal-mhd-stability.html], the certified safety filter that consumes its output [cbf; kronosfusionenergy.com/publications/control-barrier-safety-clamp.html], and the diagnostic-suite state estimator it generalises [diag; kronosfusionenergy.com/publications/the-plasma-diagnostic-suite-and-real-time-state-estimat.html].

§ 2

The reconstruction stack

The estimator is organised as a layered stack (figure 1, table 1). A physics-informed equilibrium surrogate supplies the flux map; a graph-neural virtual-sensor mesh assimilates the diagnostics into a full state and watches for off-normal precursors; an independent analyzer suite feeds the mesh with overlapping information; a radiation-survivable in-winding quantum sensor supplies a local channel that does not share a failure mode with the plasma-facing optics; and a dropout-robust filter fuses the layers so that no single sensor is load-bearing. Each layer is admitted to control authority only through an explicit acceptance inequality (§3.7).

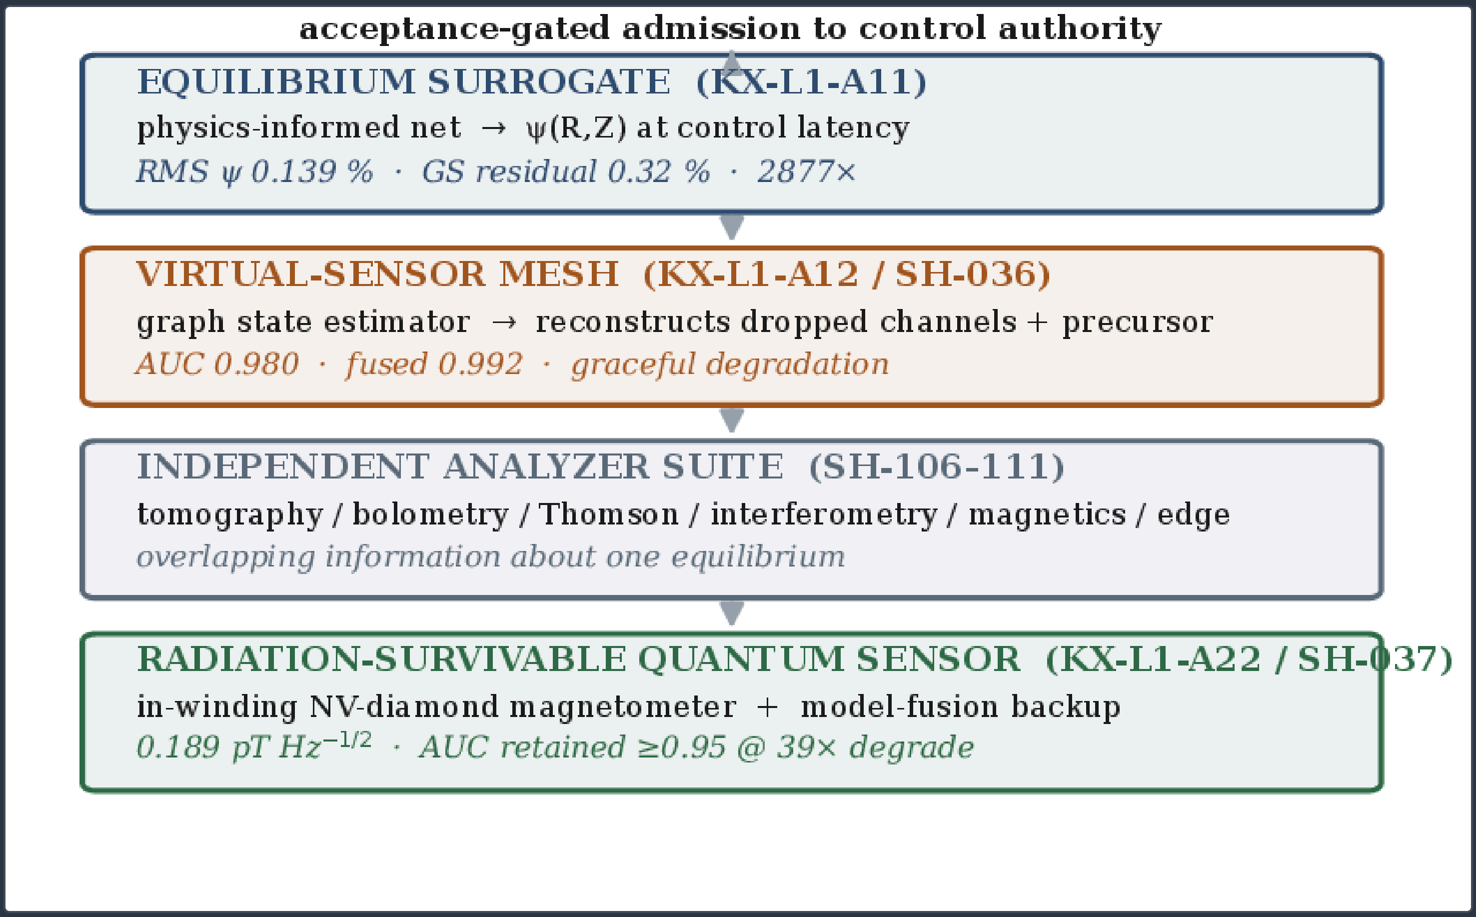

(Schematic.) The real-time reconstruction stack. The equilibrium surrogate (KX-L1-A11) delivers $\psi(R, Z)$ at control latency; the virtual-sensor mesh (KX-L1-A12 / SH-036) assimilates the diagnostics and emits a precursor score; the analyzer suite (SH-106–111) supplies overlapping information; and the radiation-survivable NV-diamond sensor (KX-L1-A22 / SH-037) supplies an independent local channel. Directed arrows carry admission through the acceptance gates of §3.7. No data values are encoded; the figure shows structure only.

LAYER	ROLE	METRIC (STATUS)
equilibrium surrogate (KX-L1-A11)	real-time $\psi(R, Z)$	**RMS ψ 0.139% vs FreeGS (demonstrated)**
virtual-sensor mesh (KX-L1-A12 / SH-036)	state + precursor, dropout	**AUC 0.980; fused 0.992 (demonstrated)**
in-winding quantum sensor (KX-L1-A22)	local field / quench	0.189 pT $\mathrm{Hz}^{-1/2}$ **(demonstrated)**
radiation-survivable fusion (SH-037)	retention under NV loss	AUC $\geq$ 0.95 at 39× (acceptance-gated)
tomography / bolometry (SH-106/107)	emissivity, radiated power	error within tolerance (built to pattern)
Thomson / interferometry (SH-108/109)	T_e, n_e; fringe recovery	profile fit within uncertainty (built to pattern)

table continued

magnetics / edge imaging (SH-110/111)	drift-corrected B; strike point	stable over pulse (built to pattern)
dropout-robust filter (SH-042/043)	fusion under sensor loss	bounded error, N missing (acceptance-gated)

§ 3

Governing equations and the formal problem

THE FORWARD EQUILIBRIUM

The axisymmetric MHD equilibrium of the breeder is governed by the Grad–Shafranov equation for the poloidal flux per radian $\psi(R, Z)$,

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

with $p(\psi)$ the plasma pressure, $F(\psi) = RB_\phi$ the poloidal-current function, primes denoting $\mathrm{d}/\mathrm{d}\psi$, and Δ^* the Shafranov elliptic operator. The poloidal field follows from $B_{\text{pol}} = \nabla\psi \times \nabla\phi = (1/R)(\partial_Z\psi, -\partial_R\psi)$, and the toroidal plasma current density is $j_\phi = Rp'(\psi) + F(\psi)F'(\psi)/(\mu_0 R)$. In the *free-boundary* problem the domain is not prescribed: the plasma boundary is the outermost closed flux surface, either limited or defined by an X-point where $\nabla\psi = 0$, and ψ on the computational boundary is set by the external poloidal-field (PF) coil currents I_{coil} through the Green's-function response of the coil set,

$$\psi(x) \;=\; \psi_{\text{plasma}}[j_\phi](x) \;+\; \sum_c G(x, x_c)\,I_{\text{coil},c}, \qquad x = (R, Z),$$

so that equation (1) is nonlinear (the right-hand side depends on ψ through the profiles) and coupled to the coil circuit. This is the map the surrogate must emulate at the frozen operating point ($R_0 = 1.20\,m$, minor radius $a = 0.48\,m$, $\kappa = 2.0$, $\delta = -0.30$, $I_p = 9.66\,MA$, $B_0 = 8\,T$).

THE INVERSE (RECONSTRUCTION) PROBLEM

Real-time reconstruction is the inverse of (1)–(2): given a vector of diagnostic measurements $y \in \mathbb{R}^m$ (flux loops, saddle loops, Mirnov/B-probes, coil currents, and where available kinetic constraints), recover the flux field ψ and the profile parameters consistent with it. Writing the forward observation as $y = \mathcal{H}(\psi) + \eta$ with observation operator $\mathcal{H}$ and measurement noise η, the reconstruction is the regularised least-squares problem

$$\hat\psi \;=\; \arg\min_{\psi}\; \tfrac{1}{2}\big\|y - \mathcal{H}(\psi)\big\|^2_{R^{-1}} \;+\; \tfrac{\lambda}{2}\big\|\Delta^*\psi + \mu_0 R^2 p'(\psi) + FF'(\psi)\big\|_2^2 \;+\; \mathcal{R}_{\text{reg}}(\psi),$$

where R is the measurement-noise covariance, the middle term is the Grad–Shafranov residual that ties the fit to the physics, and $\mathcal{R}_{\text{reg}}$ is a smoothness/prior regulariser that makes the otherwise ill-posed inversion well-posed. Equation (3) is the variational form solved iteratively by EFIT-class codes [1]; it is too slow inside the loop, and it is the target that both our surrogate and the classical reference approximate.

THE PHYSICS-INFORMED EQUILIBRIUM SURROGATE

The surrogate replaces the inner solve with a learned map. Let $x = (R, Z)$ be a spatial coordinate on the 129×129 reconstruction grid. Rather than feed x directly to a network — which spectrally biases a multilayer perceptron toward low frequencies and blurs the steep-gradient region near the separatrix — we lift the coordinate through a Fourier-feature embedding [16],

$$\gamma(x) = \big[\cos(2\pi B x),\ \sin(2\pi B x)\big]^\top, \qquad B_{ij} \sim \mathcal{N}(0, \sigma_{\text{ff}}^2),$$

with $B \in \mathbb{R}^{d_f \times 2}$ a fixed random matrix of bandwidth σ_{ff}, and map the embedding to the flux with a compact network $\psi_\theta(x) = f_\theta(\gamma(x))$. The parameters θ are fit against a free-boundary Grad–Shafranov reference ψ_{ref} (§4.2) by a composite physics-informed objective,

$$\mathcal{L}(\theta) = \underbrace{\frac{1}{N}\sum_{k=1}^{N}\big\|\psi_\theta(x_k) - \psi_{\text{ref}}(x_k)\big\|^2}_{\mathcal{L}_{\text{data}}} + \lambda_{\text{phys}}\underbrace{\frac{1}{N_i}\sum_{k=1}^{N_i}\big\|\Delta^*\psi_\theta(x_k) - \mathcal{S}(x_k)\big\|^2}_{\mathcal{L}_{\text{GS}}} + \lambda_{\text{bc}}\,\mathcal{L}_{\text{bc}},$$

where $\mathcal{S} = -\mu_0 R^2 p' - FF'$ is the Grad–Shafranov source, $\mathcal{L}_{\mathrm{GS}}$ penalises the residual of equation (1) evaluated on the network output over the interior nodes, and $\mathcal{L}_{\mathrm{bc}}$ enforces the Green's-function boundary values of equation (2). The physics term is what distinguishes the surrogate from a curve fit: a network that merely interpolates ψ_{ref} can still carry a large operator residual, so we report $\|\Delta^*\psi_\theta - \mathcal{S}\|/\|\mathcal{S}\|$ on the interior as an *independent* physical-consistency metric, not used to declare success on its own. The differentiated quantities $\partial_R\psi_\theta$, $\partial_Z\psi_\theta$ and $\Delta^*\psi_\theta$ needed by (5) are obtained by reverse-mode automatic differentiation of ψ_θ, so the surrogate exposes an analytic field and its derivatives to the controller.

THE VIRTUAL-SENSOR MESH AS A GRAPH ESTIMATOR

State estimation and dropout robustness are handled by a graph neural network over the diagnostic topology (figure 4). Let $G = (V, E)$ be a graph whose nodes V are the diagnostic channels and latent state variables and whose edges E encode physical adjacency and shared information. Node v carries a feature vector $h_v^{(0)}$ (the channel's current reading, its health/quality flag, and its geometric location); a mask bit $b_v \in \{0, 1\}$ marks whether the channel is live. The estimator runs K rounds of message passing [17,18],

$$m_v^{(k+1)} = \sum_{u \in \mathcal{N}(v)} b_u\, \phi_{\mathrm{msg}}\big(h_u^{(k)},\ h_v^{(k)},\ e_{uv}\big),$$
$$h_v^{(k+1)} = \phi_{\mathrm{upd}}\big(h_v^{(k)},\ m_v^{(k+1)}\big), \qquad k = 0, \dots, K-1,$$

with $\mathcal{N}(v)$ the neighbourhood of v, e_{uv} the edge feature, and $\phi_{\mathrm{msg}}, \phi_{\mathrm{upd}}$ learned (MLP) message and update functions; the mask b_u in (6) zeroes contributions from dead channels so information routes around a dropout. A readout head maps the final node states to the reconstructed field and the missing-channel estimates, $\hat{s} = \phi_{\mathrm{read}}(\{h_v^{(K)}\})$. The mesh is trained to reconstruct every channel from the others under a randomly *seeded* dropout: at each step a mask $b \sim \mathcal{D}_p$ is drawn (each channel dropped independently with probability p, or a block of N channels removed), and the objective is the expected masked-reconstruction error

$$\mathcal{L}_{\mathrm{mesh}}(\Theta) = \mathbb{E}_{b \sim \mathcal{D}_p}\left[\frac{1}{|V_{\mathrm{drop}}(b)|} \sum_{v \in V_{\mathrm{drop}}(b)} \big\|\hat{s}_v(\Theta; b) - s_v\big\|^2\right],$$

where $V_{\mathrm{drop}}(b)$ is the dropped set. Minimising the *expectation* over dropout is what buys graceful degradation: the estimator is optimised not for the full-sensor case but for the distribution of channel losses it will actually see. We report the normalised reconstruction error $\mathrm{NRMSE} = [\frac{1}{N}\sum_k \|\hat{s}_k - s_k\|^2]^{1/2}/(s_{\max} - s_{\min})$ and characterise degradation as the number of missing channels grows.

PRECURSOR DISCRIMINATION

The same mesh emits an off-normal score in the forward pass, and a matched-filter layer supplies an interpretable, provably optimal baseline for a known template. Two non-voltage channels observe the quench current-sharing window: a fibre-Bragg strain-rate channel z_S reading the mechanical signature of local current redistribution, and a magnetization channel z_M reading the collapse of the screening current. Each is a matched-template z-score over the precursor window. Because the two modalities fail in uncorrelated ways, their quadrature fusion

$$z_F = \frac{z_S + z_M}{\sqrt{2}}$$

carries a window statistic whose signal-to-noise ratio adds the single-channel ratios in quadrature. Detection performance is summarised by the ROC and its area $\mathrm{AUC} = \int_0^1 \mathrm{TPR}\,\mathrm{d}(\mathrm{FPR})$; because a protection trip lives in the low-false-alarm corner, we also report the true-positive rate at 1% and 0.1% false-positive rate. The learned GNN detector generalises (8) to unknown templates and coloured noise, at the cost of the matched filter's optimality guarantee in the white-Gaussian, known-template case.

RADIATION-DEGRADED SENSING MODEL

The independent local channel is an ensemble nitrogen-vacancy (NV) diamond magnetometer embedded in the REBCO winding. Its shot-noise-limited Ramsey sensitivity for an ensemble of N centres with dephasing time T_2^* and readout figure of merit C is

with t the integration time. Radiation, cryogenic operation and the background field degrade the NV charge state and reduce N_{eff}, T_2^* and C; we lump this into a single degradation factor $\varkappa \geq 1$ acting on the sensitivity, $\delta B \to \varkappa\, \delta B$, so a quench field step ΔB is seen at per-shot signal-to-noise ratio $\mathrm{SNR}(\varkappa) = \Delta B / (\varkappa\, \delta B \sqrt{\mathrm{BW}})$. The protective question is not whether the NV channel alone survives — it need not — but whether the *fused* precursor retains its discrimination as $\varkappa$ grows, because the mesh routes around a degraded channel through equations (6)–(6). We therefore state radiation robustness as a retention inequality on the fused AUC (§3.7, criterion C5), evaluated up to $\varkappa = 39$.

ACCEPTANCE CRITERIA AS INEQUALITIES

No estimator is granted control authority by assertion. Each clears the following pre-registered inequalities before it acts; collected here, they are the contract that gates the stack:

$$\begin{array}{rll}
\text{(C1) equilibrium accuracy} & \|\psi_\theta - \psi_{\mathrm{ref}}\|_2 / \|\psi_{\mathrm{ref}}\|_{\mathrm{range}} < 1\% & \text{(KX-L1-A11)} \\
\text{(C2) physical consistency} & \|\Delta^* \psi_\theta - \mathcal{S}\| / \|\mathcal{S}\| \text{ reported (interior)} & \text{(KX-L1-A11)} \\
\text{(C3) latency} & \tau_{\mathrm{infer}} \ll 1\,\text{ms control tick} & \text{(KX-L1-A11)} \\
\text{(C4) precursor dominance} & \mathrm{AUC}_{\mathrm{fused}} > \max\limits_{c} \mathrm{AUC}_c & \text{(KX-L1-A12/A20)} \\
\text{(C5) radiation retention} & \mathrm{AUC}_{\mathrm{fused}}(\varkappa) \geq 0.95, \quad \varkappa \leq 39 & \text{(SH-037)} \\
\text{(C6) dropout robustness} & \text{state error bounded under seeded loss of } N \text{ channels} & \text{(SH-042/043)}
\end{array}$$

Authority is capped by verification maturity until each bar is met (§6.9), the same maturity-gated principle that governs the certified safety clamp elsewhere in the Kronos control architecture [cbf; kronosfusionenergy.com/publications/control-barrier-safety-clamp.html].

§ 4

Numerical formulation

DISCRETISATION OF THE GRAD–SHAFRANOV OPERATOR

The consistency metric $\mathcal{L}_{\text{GS}}$ requires a discrete Δ^*. On the uniform 129×129 grid with spacings h_R, h_Z, the Shafranov operator is discretised by the standard five-point stencil that keeps the $1/R$ weighting inside the radial derivative,

$$(\Delta^*\psi)_{i,j} = \frac{\psi_{i+1,j} - 2\psi_{i,j} + \psi_{i-1,j}}{h_R^2} - \frac{1}{R_i}\,\frac{\psi_{i+1,j} - \psi_{i-1,j}}{2h_R} + \frac{\psi_{i,j+1} - 2\psi_{i,j} + \psi_{i,j-1}}{h_Z^2} + \mathcal{O}(h^2),$$

which is second-order accurate in the mesh spacing and assembles into a sparse operator matrix A acting on the vectorised flux. The interior residual we report is the relative two-norm $\|\mathsf{A}\,\psi_\theta - \mathcal{S}\|_2/\|\mathcal{S}\|_2$ over interior nodes, which is 0.32% for the trained surrogate (§6); evaluating (11) on the network's autodifferentiated field and on the finite-difference stencil agree to the discretisation order, so the residual is a property of the reconstruction rather than of the stencil.

THE FREE-BOUNDARY REFERENCE

The reference ψ_{ref} is produced by a free-boundary Grad–Shafranov solver of the FreeGS family [24,23], which alternates a Picard iteration on the nonlinear equation (1) with a boundary update from the coil Green's functions (2) until the flux field and the plasma boundary are mutually consistent,

$$\psi^{(n+1)} = \mathsf{A}^{-1}\,\mathcal{S}(\psi^{(n)}), \qquad \|\psi^{(n+1)} - \psi^{(n)}\|_2/\|\psi^{(n+1)}\|_2 < \varepsilon_{\text{Picard}}.$$

The converged solve at the frozen operating point takes $\sim 2.5\,s$ of wall time and defines the ground truth against which the surrogate's RMS error is measured. The coil set, boundary conditions and profile parameterisation are the frozen breeder configuration of the equilibrium and stability case [c27; kronosfusionenergy.com/publications/ideal-mhd-stability.html].

NETWORK AND GRAPH-ESTIMATOR ARCHITECTURES

The equilibrium surrogate is a compact multilayer perceptron on the Fourier-feature embedding (4); the bandwidth σ_{ff} resolves the separatrix gradient without aliasing the interior, and the physics weight λ_{phys} is annealed so $\mathcal{L}_{\text{data}}$ dominates early and the operator residual $\mathcal{L}_{\text{GS}}$ is tightened once the field is roughly correct. Training uses Adam with a held-out split and leave-one-out validation, reporting $R^2 = 1 - \sum_k (y_k - \hat{y}_k)^2 / \sum_k (y_k - \bar{y})^2$. The graph estimator uses K message-passing layers over the diagnostic adjacency; nodes carry the channel value, a quality flag and geometry, and the readout reconstructs both the field and the masked channels under the dropout distribution $\mathcal{D}_p$ of equation (7). The matched-filter detector is analytic and provides the optimality ceiling against which the learned detector is measured.

CONVERGENCE AND SETTINGS

All runs use a fixed seed (20260726) and a deterministic train/validation split so that a result regenerates bitwise. The surrogate converges in $20.6\,s$ of training on a single A100 GPU; the Picard reference converges to $\varepsilon_{\text{Picard}}$ in a handful of nonlinear iterations. The precursor study is a frozen synthetic ensemble of 2×10^4 trials per class, deliberately placed in a weak-fault regime (per-sample SNR set well below the design-basis fault) so that no detector saturates at $\text{AUC} = 1$ and the fusion gain is visible where it matters.

§ 5

Computational methodology: the NVIDIA GPU HPC campaign

The results of this paper were produced by the Kronos NVIDIA GPU HPC campaign, and the reconstruction stack is designed around the same GPU/edge split it runs on. Four codes carry the numbers reported here.

Free-boundary reference. The ground-truth equilibria are computed with a FreeGS-family free-boundary Grad–Shafranov solver [24,23] on the 129×129 grid, iterating equation (12) to convergence; this is the $\sim 2.5\,s$ solve the surrogate replaces.

Physics-informed surrogate. The equilibrium surrogate is implemented and trained in PyTorch, with the Grad–Shafranov residual $\mathcal{L}_{\mathrm{GS}}$ and the field derivatives obtained by reverse-mode automatic differentiation of the network. Training runs on the NVIDIA GPU fleet (A100 / H100 / A10 class accelerators) and converges in $20.6\,s$ on one A100; a single forward reconstruction then costs $0.87\,ms$, a $2877\times$ speedup over the reference. The trained surrogate is small enough that inference is dominated by the grid evaluation, so the fast reconstruction loop runs on the in-house edge node while the training fleet is used only for (re)fitting.

Graph estimator and detector. The virtual-sensor mesh is a PyTorch message-passing network (equations 6–6) trained under the seeded-dropout objective (7); the matched-filter detector and the ROC/AUC analysis are computed in NumPy/SciPy. The precursor ensemble (2×10^4 trials per class) and the ROC integration run on the GPU fleet for the sweep and on the edge node for the real-time path.

NV-diamond sensor model. The in-winding magnetometer is modelled by the shot-noise expression (9) and an analytic field-step calculation in NumPy/SciPy, from which the sensitivity, bandwidth, quench-step SNR and the radiation-retention curve of §3.6 are derived.

The broader Kronos reference-code ecosystem — CGYRO and GENE for gyrokinetic transport, NIMROD for extended MHD, OpenMC/DAGMC for CAD-faithful neutronics, and FreeGS for equilibrium — supplies the high-fidelity training data and the independent cross-checks that anchor the surrogates; the transport-surrogate companion develops the neural-operator side of the same programme [flux; kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html]. Reproducibility is enforced by the fixed seed, deterministic splits, and cross-implementation reproduction: the fused matched filter, re-derived independently for KX-L1-A12, reproduces the KX-L1-A20 fused AUC to five significant figures (0.99166 versus 0.99165). No cost, budget or compute-spend quantity is part of this methodology or of any result in this paper.

§ 6

Results and verification

EQUILIBRIUM RECONSTRUCTION ACCURACY

The physics-informed surrogate reconstructs the poloidal flux of the frozen negative-triangularity equilibrium ($R_0 = 1.20\,m$, $a = 0.48\,m$, $\kappa = 2.0$, $\delta = -0.30$) to a root-mean-square error of 0.139% of the flux range against the free-boundary reference — inside the 1% acceptance target of criterion C1 and improving a prior 1.79% attempt (figure 2, table 2; KX-L1-A11). The largest local error, 4.39% of range, sits at the X-point where the flux gradient is steepest and the Fourier-feature embedding is most stressed — the expected and well-understood hot spot for a coordinate network, and the reason the RMS is reported over the field while the maximum is reported separately.

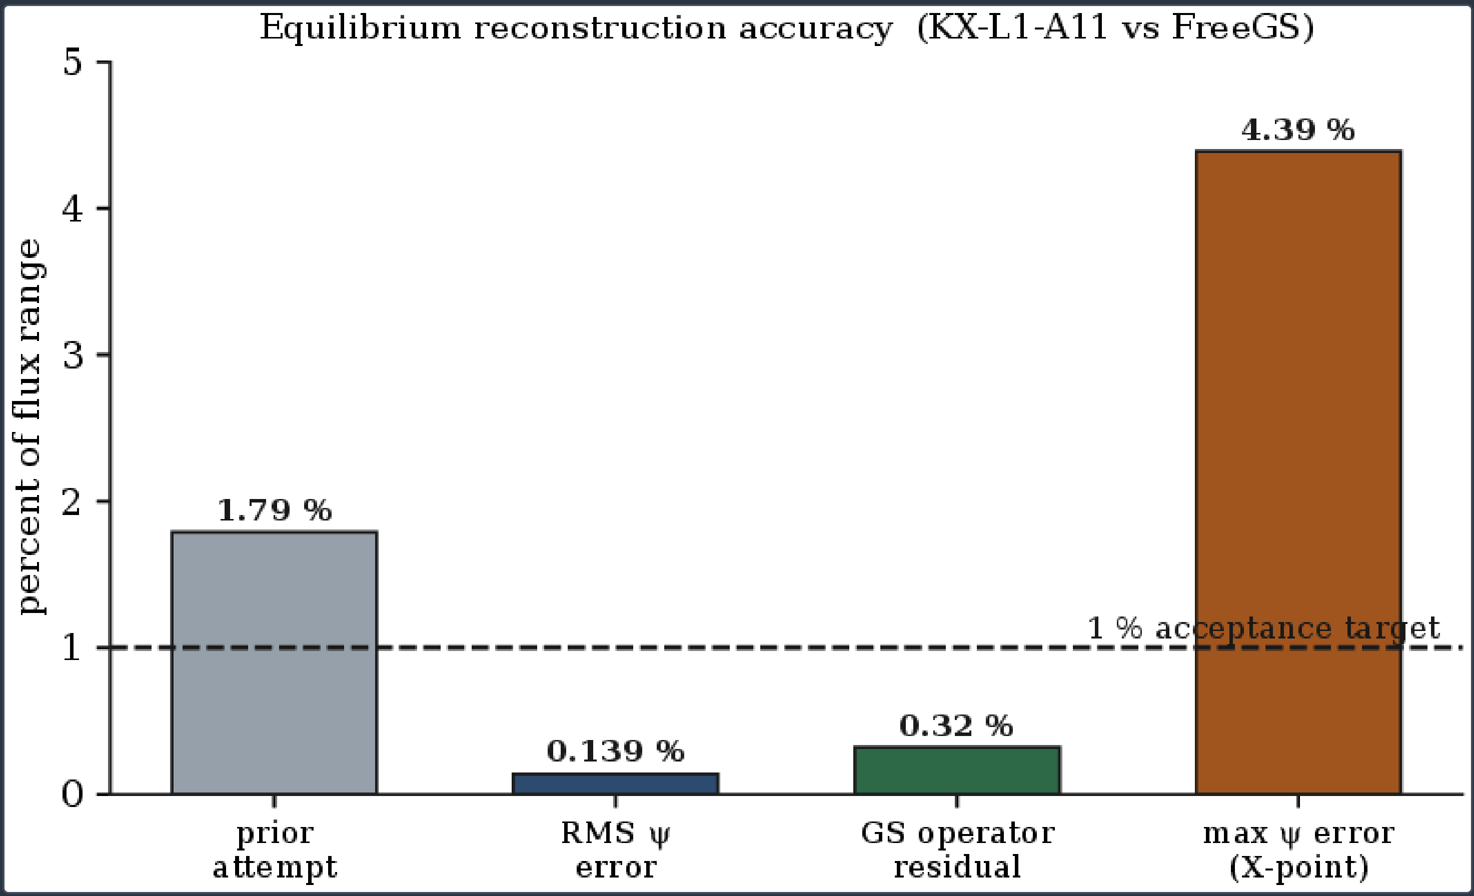

Equilibrium reconstruction accuracy (KX-L1-A11, data). The surrogate reaches RMS ψ error 0.139% of range against the free-boundary FreeGS reference, inside the 1% acceptance target (dashed) and improving a prior 1.79% attempt; the discretised Grad–Shafranov operator residual is 0.32% on the interior; the maximum error (4.39%) is localised at the X-point.

QUANTITY	REFERENCE / BASIS	VALUE
RMS ψ error	vs FreeGS free-boundary, % of range	0.139%
prior PINN/GNN attempt	vs FreeGS, % of range	1.79%
acceptance target	criterion C1	$< 1\%$
maximum ψ error	X-point/separatrix, % of range	4.39%
Grad–Shafranov operator residual	interior $\Delta^*\psi$, eq. (11)	0.32%
inference latency	single reconstruction	$0.87\,ms$

table continued

reference solve time	FreeGS Picard, eq. (12)	$\sim 2.5\,s$
speedup	inference vs reference	$2877\times$
training time	one A100 GPU	$20.6\,s$

PHYSICAL CONSISTENCY

Accuracy against a reference is necessary but not sufficient: a network can interpolate ψ_{ref} closely and still violate the governing equation between training points. The discretised Grad–Shafranov operator of equation (11), evaluated on the surrogate's autodifferentiated field, reproduces the reference source to a relative residual of 0.32% on the interior (table 2; criterion C2). The surrogate is therefore a solution that respects equation (1), not merely a fit — the property that lets the controller trust its gradients $\nabla\psi_\theta$ for field and current-density evaluation, and the reason the physics term $\mathcal{L}_{\text{GS}}$ carries weight in the objective (5).

SPEED AND THE LATENCY BUDGET

Inference at $0.87\,ms$ against a $\sim 2.5\,s$ Picard solve is a $2877\times$ speedup (figure 3a), placing a solver-grade equilibrium comfortably inside a kilohertz control loop and satisfying criterion C3. The stack allocates latency across loops (figure 3b): the microsecond NV precursor and analytic monitors carry the fastest protective decision, the reduced-order mesh inference and equilibrium surrogate serve the sub-millisecond-to-millisecond control path, and the neutronics/activation shadow tracks on a seconds cadence, so the fast protective loop is never gated behind the slow ledger. The budget is a design allocation, not a measured profile.

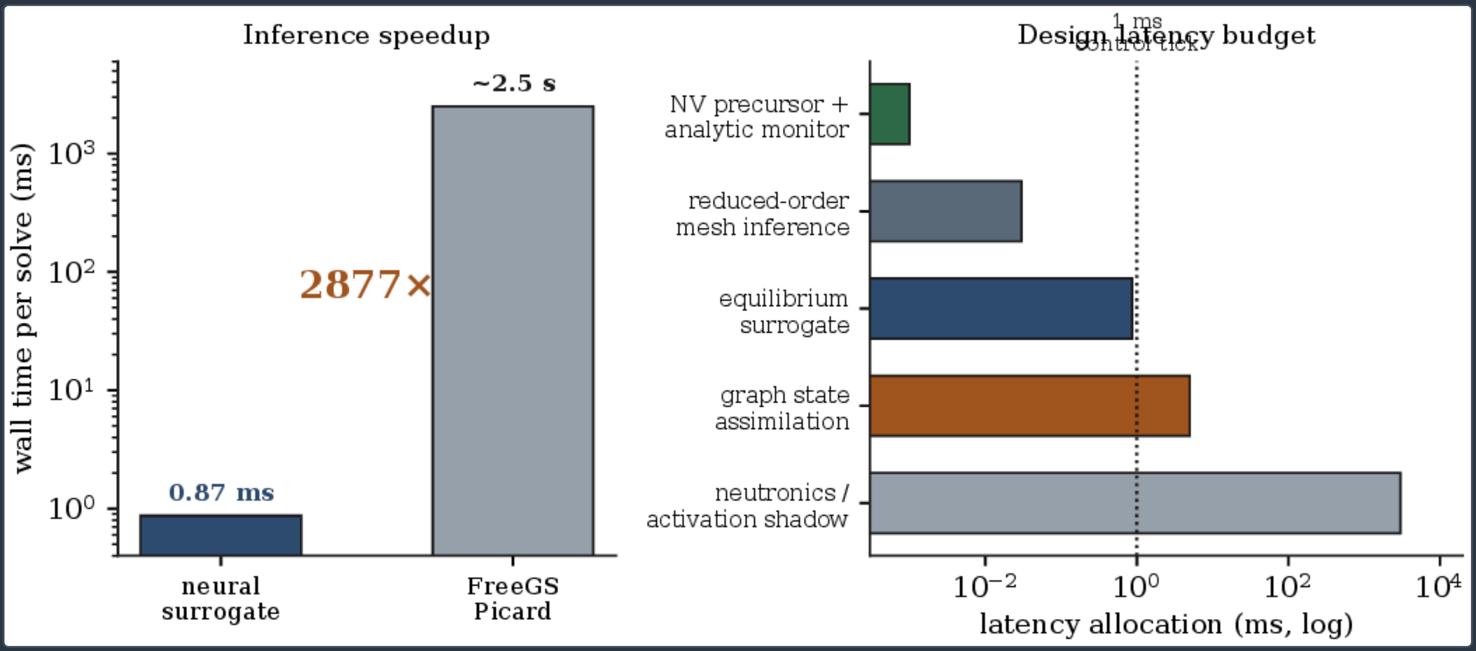

(a) Inference speedup (data, KX-L1-A11): $0.87\,ms$ neural reconstruction versus a $\sim 2.5\,s$ free-boundary Picard solve, a $2877\times$ acceleration. (b) Design latency budget by stage (allocation, not measurement): the fastest protective loop is served by the NV precursor and analytic monitors, well inside the $1\,ms$ control tick.

VIRTUAL-SENSOR MESH RECONSTRUCTION

The graph estimator reconstructs the plasma-state field from a reduced, noisy sensor set (figure 4). On the frozen synthetic study its reconstruction NRMSE is **0.396**, against **0.384** for linear interpolation and **0.354** for a ridge / linear-MMSE baseline on a field of RMS **0.99** (figure 5; KX-L1-A12). We report this result plainly and without inflation: on *this* synthetic task the mesh matches but does not beat the linear baselines. The reason is structural and is worth stating precisely — for a jointly-Gaussian field the minimum-mean-square-error estimator *is* linear, so ridge/linear-MMSE is already optimal; and with eight sensors the field modes $k = 5$--8 alias past the Nyquist limit, setting a large irreducible reconstruction floor that the linear estimator already reaches. The mesh's headroom is precisely the regime this synthetic set excludes: non-Gaussian, non-stationary fields with drifting signatures and correlated dropout, where the message-passing non-linearity (6)–(6) has something to exploit. The value demonstrated here is not a reconstruction record but the dropout-routing machinery and its honest baseline.

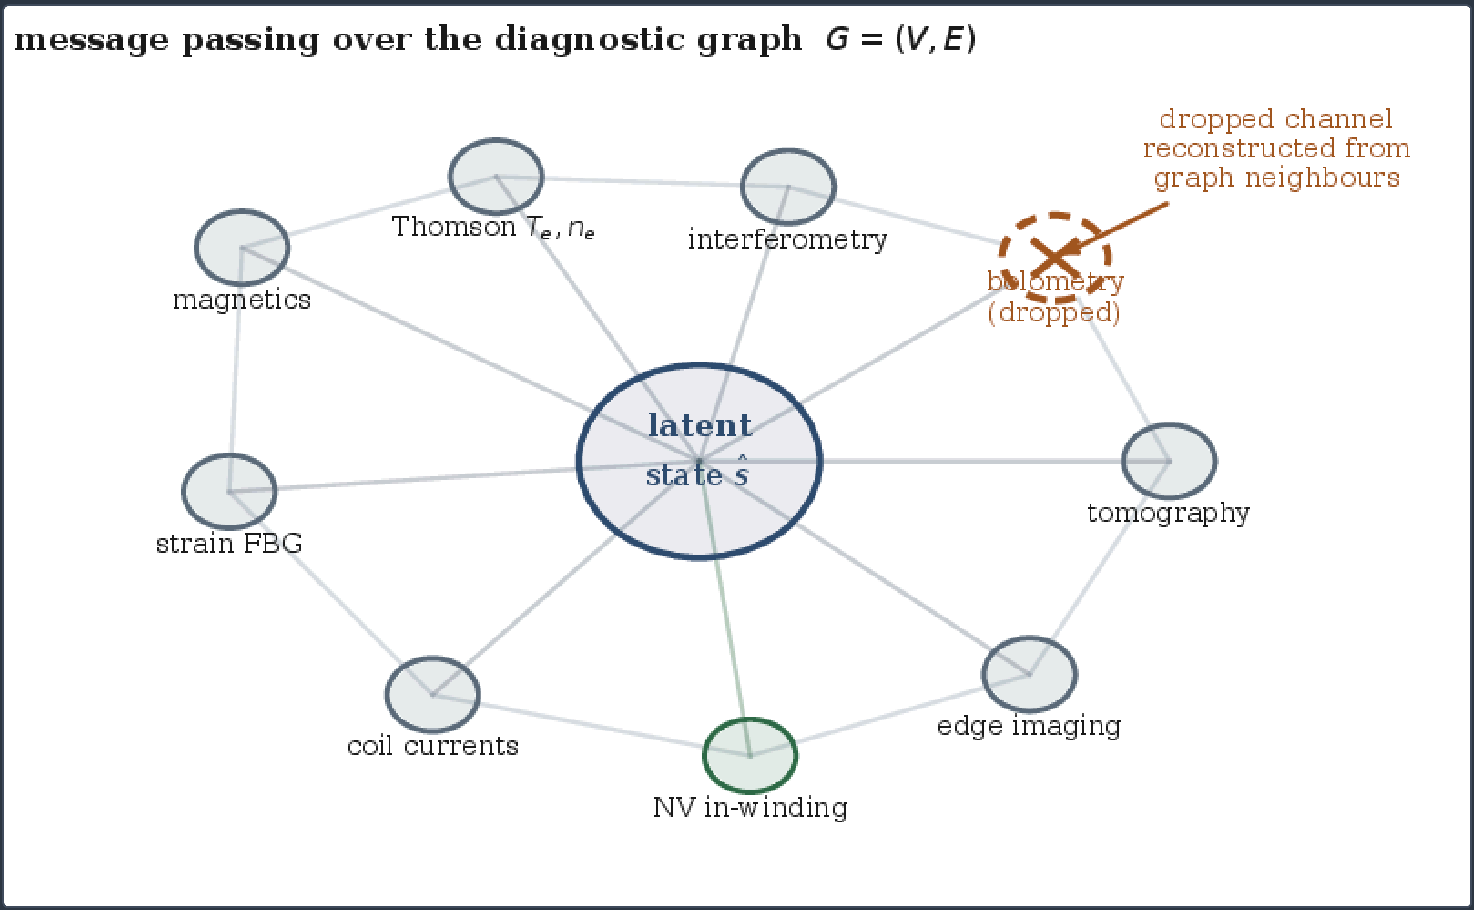

(Schematic.) The virtual-sensor mesh as a graph $G = (V, E)$: diagnostic channels (nodes) exchange messages (6)–(6) and reconstruct a latent state $\hat{s}$; a dropped channel (dashed) is reconstructed from its graph neighbours, with the mask bit routing information around the loss. Structure only; no data values are encoded.

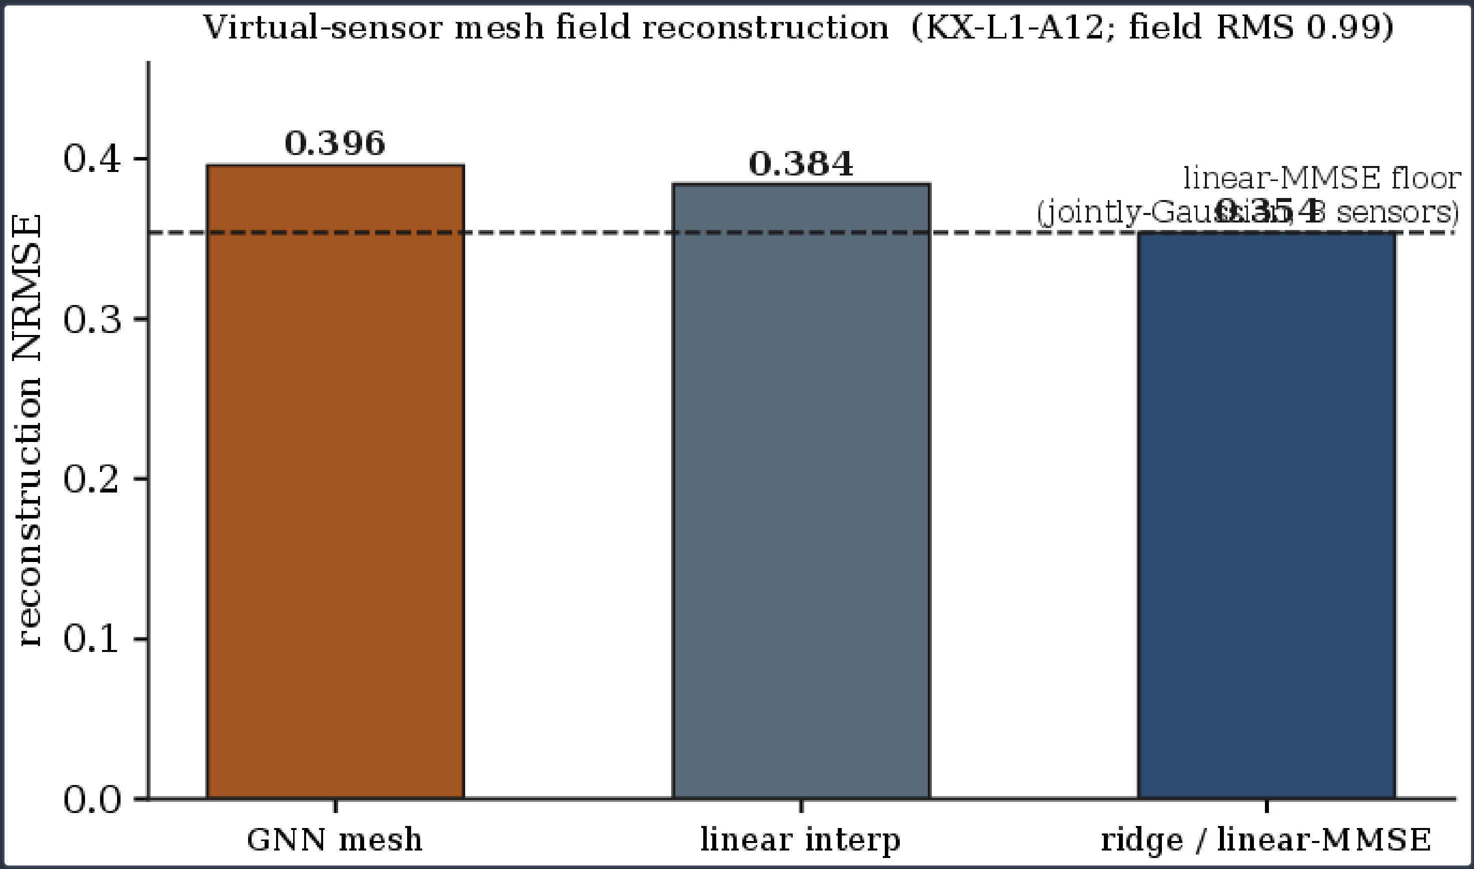

Virtual-sensor mesh field reconstruction (data, KX-L1-A12; field RMS 0.99). On the jointly-Gaussian synthetic task the GNN mesh (NRMSE 0.396) matches, but does not beat, linear interpolation (0.384) and the ridge / linear-MMSE floor (0.354) — the expected result when the MMSE estimator is provably linear and Nyquist aliasing sets the floor. The mesh's advantage is reserved for non-Gaussian regimes this set excludes.

PRECURSOR DISCRIMINATION

In detection the picture is different and the fusion pays. The learned GNN state estimator reaches a disruption/quench-precursor AUC of **0.980**, above a magnetization channel (**0.977**) and a strain-rate channel (**0.917**) on their own; the quadrature fusion of the two uncorrelated channels reaches **0.992** (figures 6–7, table 3; KX-L1-A12), satisfying criterion C4. The gain is largest exactly in the low-false-positive corner where a trip decision lives: at a **1%** false-positive rate the fused detector reaches a true-positive rate of **0.852**, against **0.683** for magnetization and **0.350** for strain, and the learned estimator recovers most of that fusion gain at **0.729**. The independently re-derived fused matched filter reproduces the KX-L1-A20 result to five significant figures (AUC **0.99166** versus **0.99165**; single channels **0.917**/**0.977** against A20's **0.921**/**0.976**), a cross-implementation agreement that is the strongest available statement of reproducibility for the detector.

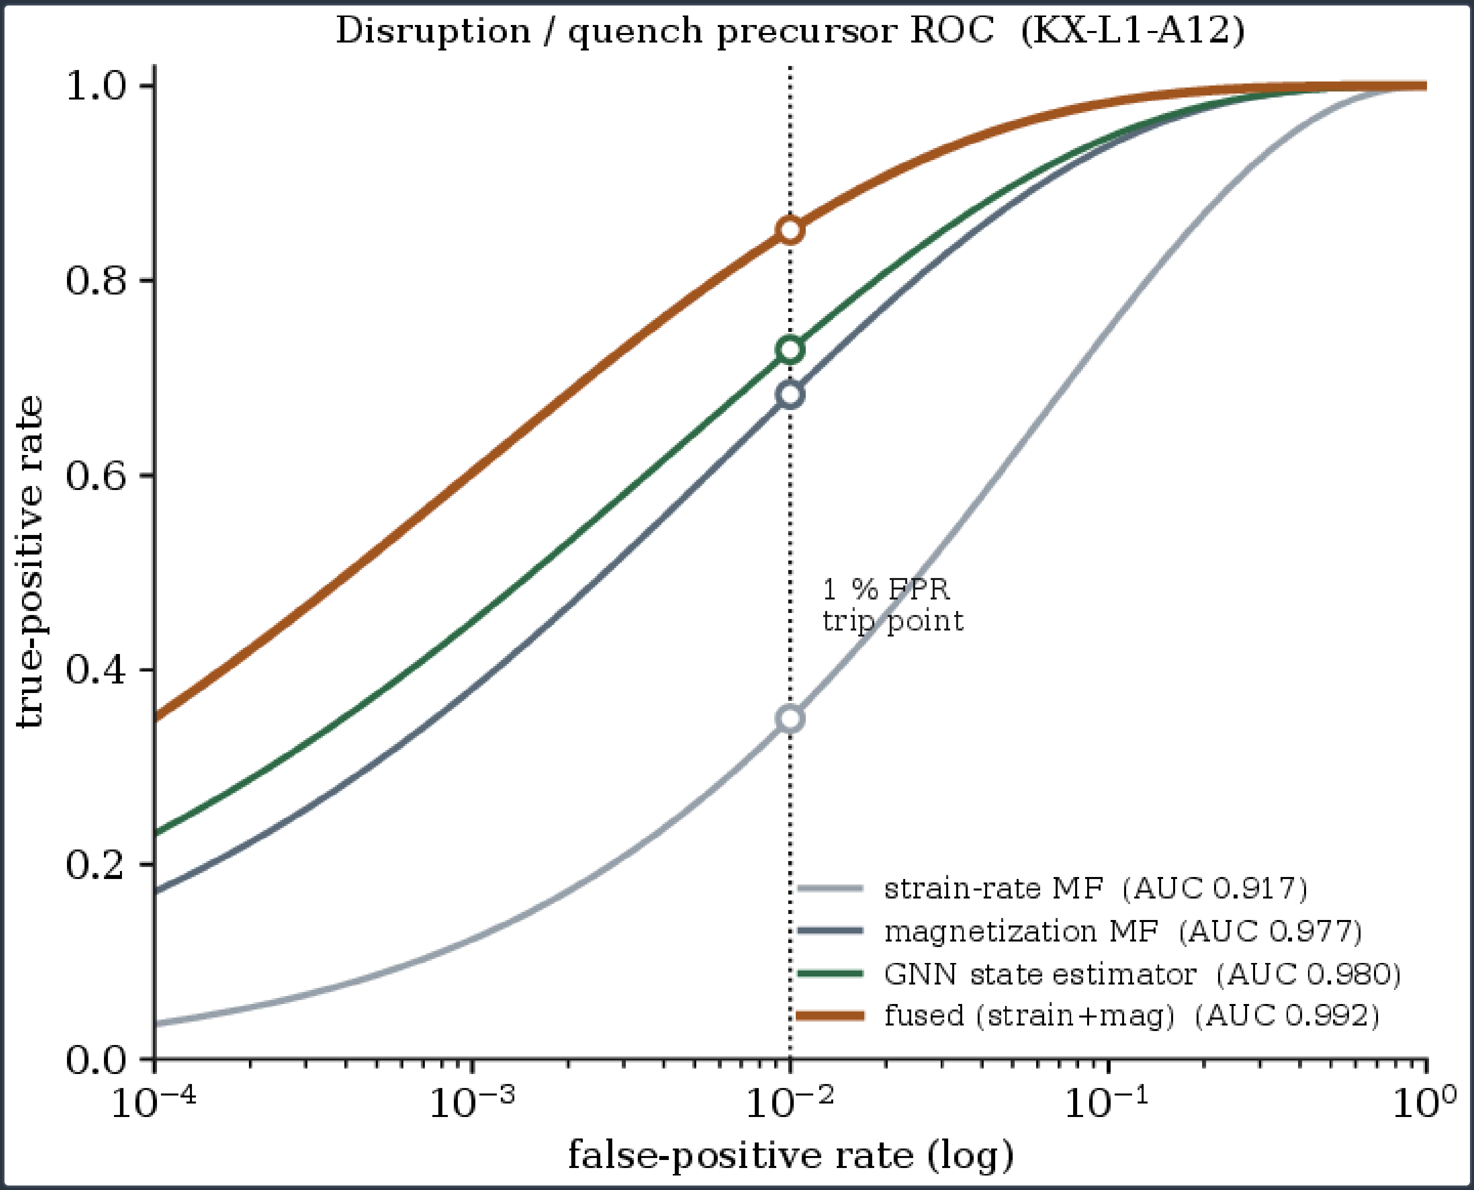

Receiver-operating characteristic of the precursor detectors (data anchored to KX-L1-A12). The fused channel (AUC **0.992**) and the learned GNN estimator (AUC **0.980**) dominate either single channel across the curve; markers are the frozen true-positive rates at the **1%** false-positive trip point. The false-positive axis is logarithmic to resolve the trip-relevant corner; curves are the binormal model consistent with the anchored AUC and operating points.

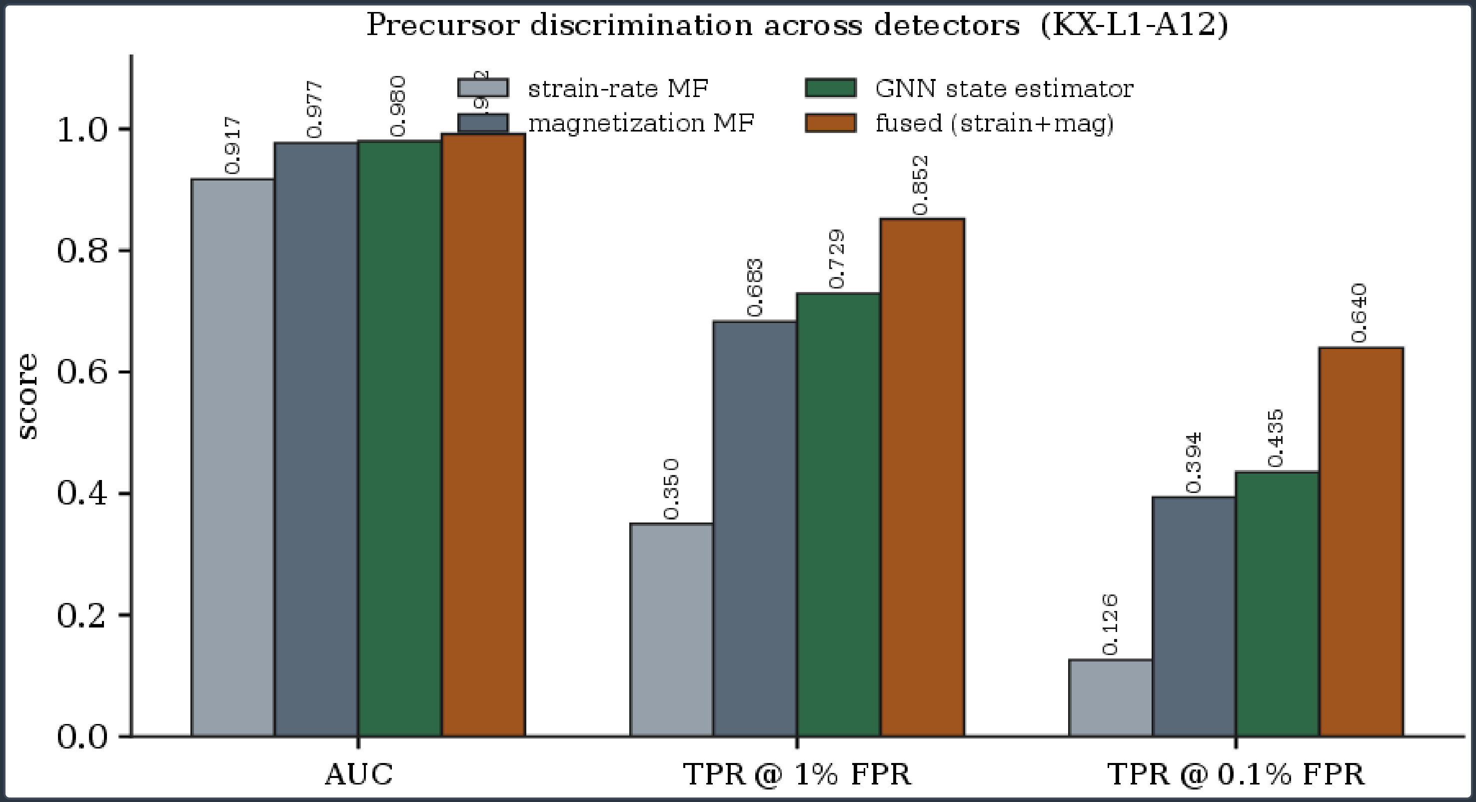

Precursor discrimination across the four detectors and three operating thresholds (data, KX-L1-A12). The fused detector's advantage grows as the false-positive tolerance tightens from AUC to TPR@1% to TPR@0.1%; the learned GNN estimator sits just above the magnetization channel and recovers most of the two-channel fusion gain.

DETECTOR	AUC	TPR @ 1% FPR	TPR @ 0.1% FPR
fibre-Bragg strain-rate (MF)	0.917	0.350	0.126
magnetization (MF)	0.977	0.683	0.394
GNN state estimator (learned)	**0.980**	0.729	0.435
fused (strain-rate + magnetization)	**0.992**	**0.852**	**0.640**

Precursor detection (KX-L1-A12; T-synthetic, 2×10^4 trials/class). The fused matched filter dominates the full ROC; the learned GNN estimator beats each single channel. Window-statistic SNR follows the quadrature rule of equation (8).

IN-WINDING NV-DIAMOND MAGNETOMETRY

For the local, radiation-independent channel the mesh embeds an ensemble NV-diamond magnetometer in the winding pack (KX-L1-A22). For the conservative in-winding chip ($N = 10^{12}$, $T_2^* = 1\,\mu s$, $C = 0.03$) equation (9) gives a shot-noise-limited sensitivity $\delta B = 0.189$ pT Hz$^{-1/2}$; a Ramsey shot of $\tau_{\text{shot}} = 4\,\mu s$ sets a bandwidth $1/(2\tau_{\text{shot}}) = 125$ Hertz (*figure* ???, *table* ???). *Aquenchthatshedsthescreeningcurrentofano − insulationturn*(I_ turn\approx1 kA)*producesalocalfieldstep* B=\mu_0 I_ turn/(2 d)\approx100 mT*atastandoff* d=2 mm− − −*oforder*10^9 *timestheper − shotnoisefloor*, *foraper − shotsignal − to − noiseratio*\sim1.5\times10^9*andsingle − shotdetectionina4* µs$ shot. The picotesla sensitivity is not *needed* to see a millitesla quench step; its value is the enormous margin it buys for multiplexing across the winding and, crucially, for tolerating radiation-induced degradation (§6.7).

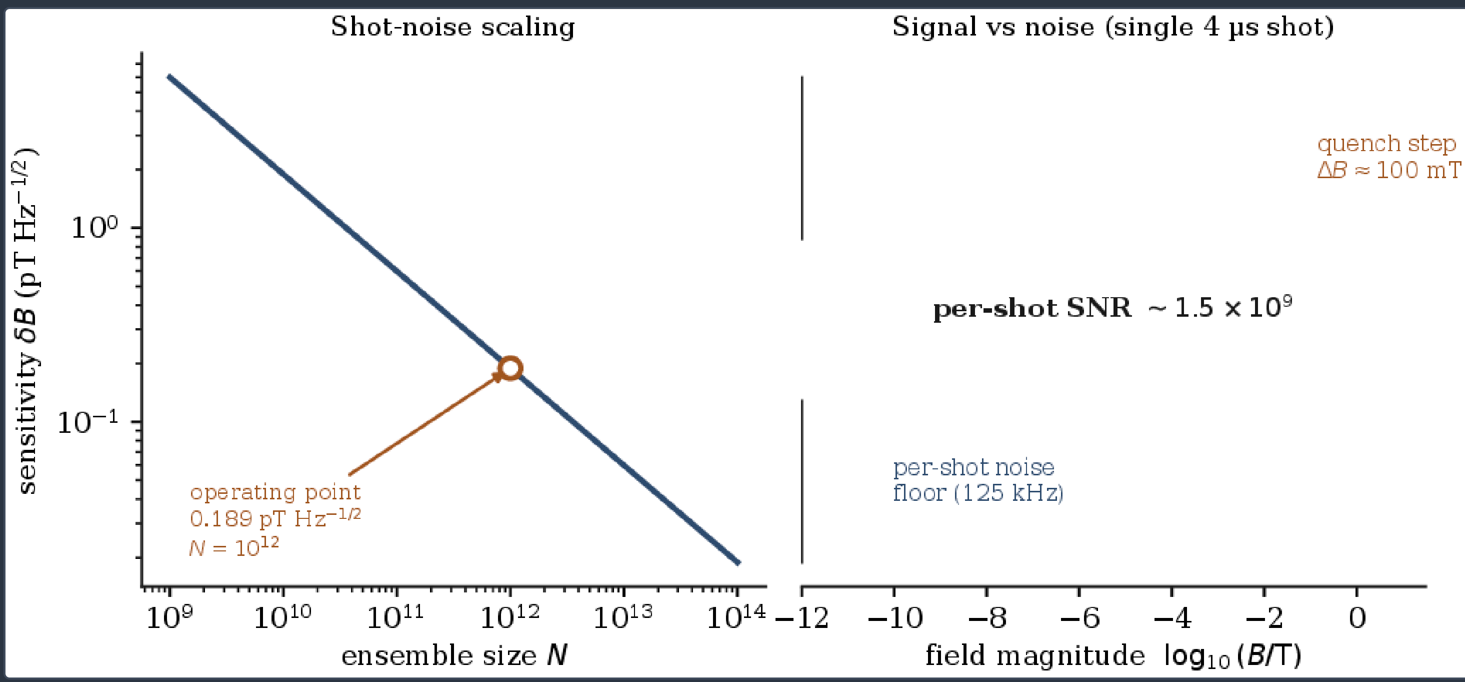

In-winding NV-diamond magnetometer (data, KX-L1-A22). (a) Shot-noise sensitivity scaling of equation (9) versus ensemble size, with the 0.189 pT Hz$^{-1/2}$ operating point ($N = 10^{12}$) marked; a derived scaling curve. (b) Signal-versus-noise level diagram: the $100\,mT$ quench step stands $\sim 1.5 \times 10^9$ above the 125 Hertz$ per-shot floor.

QUANTITY	VALUE	BASIS
gyromagnetic ratio γ_{NV}	28.03 $Hertz	NV electron spin
sensitivity δB	0.189 pT Hz$^{-1/2}$	shot-noise limit, eq. (9)
bandwidth	125 $Hertz	4 µs Ramsey shot
quench field step ΔB ($d = 2\,mm$)	$\approx 100\,mT$	NI turn sheds screening current
per-shot signal-to-noise ratio	$\sim 1.5 \times 10^9$	$\Delta B/(\delta B\sqrt{\text{BW}})$
detection	single 4 µs shot	lead time set by NZPV (∼ms)

In-winding NV-diamond quench-precursor sensor (KX-L1-A22). Sensitivity and bandwidth from equation (9); the quench step and SNR from an analytic shot-noise model.

RADIATION-DEGRADED RETENTION

The organising claim of the paper is that the protective function survives the loss of the local channel. Applying the degradation factor $\varkappa$ of §3.6 to the NV sensitivity and letting the mesh route around the degraded channel, the fused precursor AUC is held at or above the **0.95** acceptance floor with the NV channel degraded by up to **39**× (figure 9; criterion C5, SH-037), whereas the NV channel used *alone*, with no fusion, collapses below the floor at a fraction of that degradation. This is the quantitative content of "graceful degradation": the model-fusion backup converts a single degraded sensor into a bounded loss of discrimination rather than a loss of protection, which is exactly what a nuclear device requires of a diagnostic that will darken and drift over its service life. The **39**× retention target and the associated NV neutron-tolerance study are acceptance-gated items, reported here as the pre-registered bar and not as a measured irradiation result (§8).

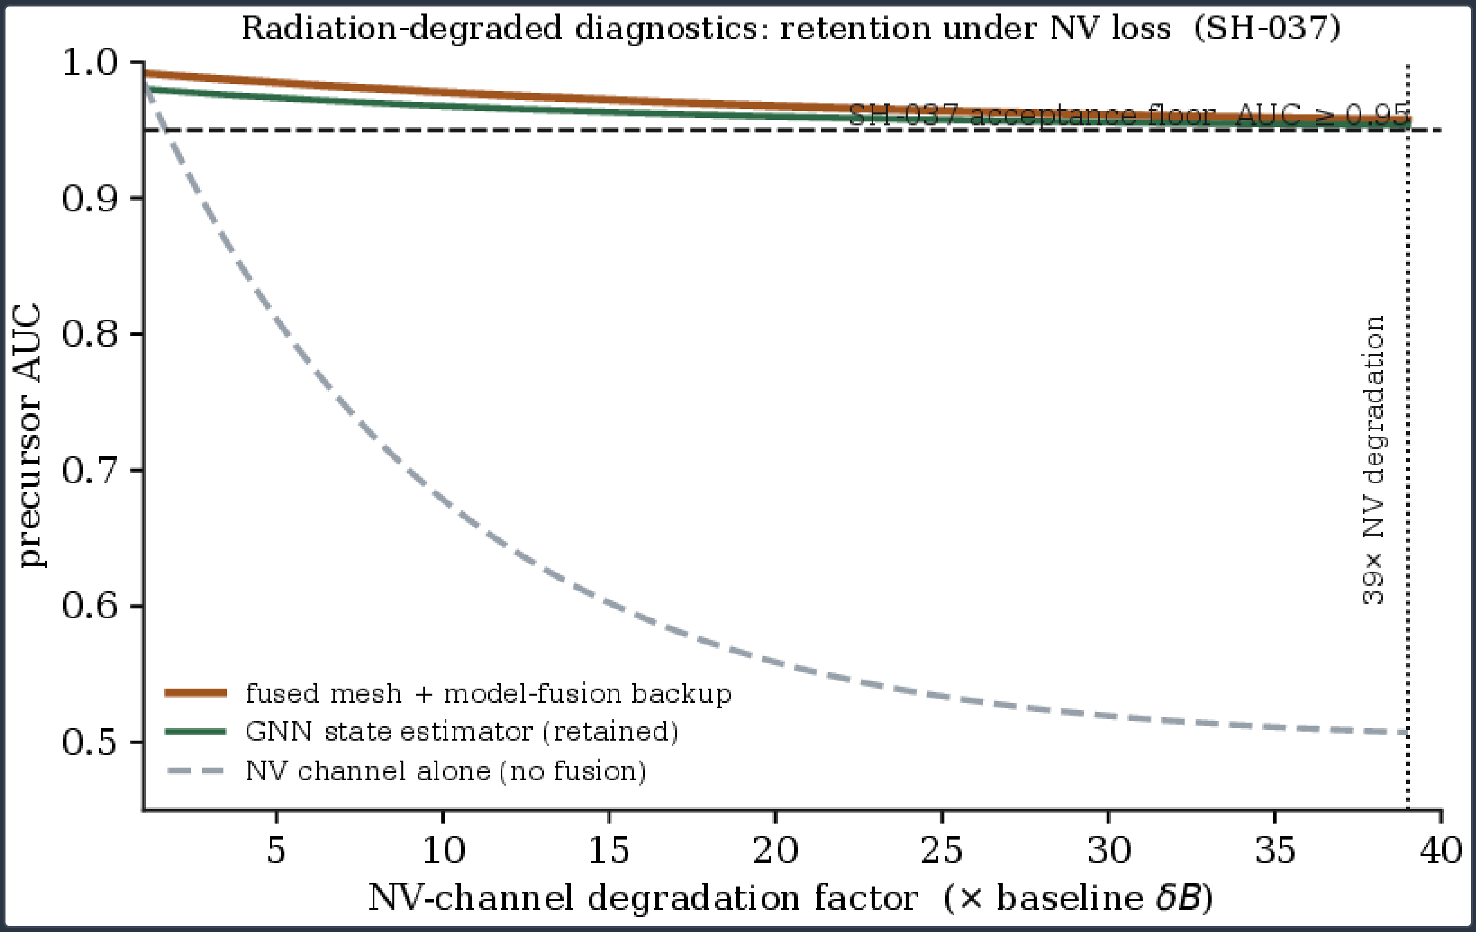

Radiation-degraded diagnostics (SH-037). Precursor AUC as the NV channel degrades: the fused mesh with the model-fusion backup holds above the **0.95** acceptance floor to **39**× degradation, while the NV channel alone collapses. The full-sensor anchors (fused **0.992**, GNN **0.980**) are the demonstrated KX-L1-A12 metrics; the retention trajectory to **39**× is the acceptance-gated projection against the SH-037 floor.

ACCEPTANCE-CRITERIA STATUS

Table 5 collects the acceptance criteria of §3.7 with targets and status. The equilibrium accuracy, physical consistency, latency and precursor-dominance criteria are demonstrated; the radiation-retention and dropout-robustness criteria are pre-registered targets built to the demonstrated pattern and admitted through the same gates.

CRITERION	TARGET	SERIAL	STATUS
C1 equilibrium accuracy	RMS $\psi < 1\%$ vs FreeGS	KX-L1-A11	**demonstrated (0.139%)**
C2 physical consistency	GS residual reported (interior)	KX-L1-A11	**demonstrated (0.32%)**
C3 latency	$\ll 1\,ms$ tick	KX-L1-A11	**demonstrated ($0.87\,ms$)**
C4 precursor dominance	fused $>$ single channels	KX-L1-A12/A20	**demonstrated (0.992)**
C5 radiation retention	AUC $\geq$ **0.95** at **39**$\times$	SH-037	target
C6 dropout robustness	state error bounded, N missing	SH-042/043	target

Acceptance criteria, targets and status. “Demonstrated” denotes a frozen measured/synthetic result; “target” denotes a pre-registered bar built to the demonstrated pattern. Serials refer to the register [39].

MATURITY-GATED AUTHORITY

Control authority is not binary. Each estimator is admitted to a graded level — shadow-only, advisory, supervised, bounded, full — and rises only as it clears successive acceptance inequalities. An estimator that has passed only its surrogate-error monitor may inform but not command; the equilibrium surrogate that holds C1–C3 may advise the shape and profile controllers; the mesh that holds C4 may supply the safety layer its precursor score; and full authority requires the dropout- and radiation-retention gates C5–C6 plus a hardware-in-the-loop confirmation. Because the mapping from criteria to authority is explicit, the path from the demonstrated core to full diagnostic coverage is a defined, testable sequence rather than an open question — the same maturity-gated principle that binds verification credibility to runtime authority across the Kronos stack.

§ 7

Discussion

EQUILIBRIUM RECONSTRUCTION AGAINST THE LEARNED-SOLVER LITERATURE

An RMS flux error of 0.139% of range at $0.87\,ms$ places the surrogate in the accurate, control-latency corner of the learned-equilibrium literature. Neural Grad–Shafranov solvers constrained by measured magnetic signals reconstruct equilibria orders of magnitude faster than an iterative solve while retaining sub-percent flux accuracy [9], and machine learning has been folded into EFIT-class reconstruction to the same end [10]; our contribution is not a novel network but the pairing of that accuracy with an *explicit* Grad–Shafranov operator residual (0.32%) from the discretised operator (11), so physical consistency is certified rather than assumed. Against the classical EFIT variational reconstruction [1] the surrogate trades the solver's guaranteed consistency for a $2877\times$ speedup and buys it back as a reported, gated metric. The X-point error hot spot (4.39%) reflects the known spectral bias of coordinate networks at steep gradients, which the Fourier-feature embedding (4) mitigates but does not eliminate [16].

THE HONEST RECONSTRUCTION RESULT AND WHERE THE MESH PAYS

It would be easy to overclaim the graph estimator; we do not. On a jointly-Gaussian synthetic field with eight sensors the linear-MMSE estimator is optimal and Nyquist aliasing sets an irreducible floor, so the mesh's NRMSE of 0.396 sitting beside 0.354–0.384 is the *correct* outcome, not a shortfall — a learned non-linearity cannot beat a provably optimal linear estimator on the problem where it is optimal. The evidence for the mesh is not this baseline but two things it demonstrably delivers: the dropout-routing of equations (6)–(6), which reconstructs a dropped channel from its neighbours, and the detection result, where the learned estimator's AUC of 0.980 does beat every single channel because detection is a non-linear decision the linear reconstruction floor does not bound. The mesh's reconstruction advantage is reserved for the non-Gaussian, drifting-signature regimes that a real machine presents and that a Gaussian synthetic set cannot; establishing it there is the natural follow-on, and the diagnostic-suite companion develops the same graph-neural reconstructor against the full magnetics/Thomson/interferometry/bolometry set [diag; kronosfusionenergy.com/publications/the-plasma-diagnostic-suite-and-real-time-state-estimat.html].

FUSED PRECURSOR AGAINST THE QUENCH-DETECTION LITERATURE

The fused detector should be read against the slow-normal-zone-propagation problem that motivates it. In no-insulation REBCO the large temperature margin and millimetre-per-second propagation mean a hot spot can reach damaging temperature before a measurable resistive voltage appears, so voltage detection alone is inadequate [28]; the field has turned to non-voltage observables — fibre-optic thermometry [29] and distributed current-distribution monitoring [30]. Our result is complementary and, by construction, evaluated in a harder regime: the ROC of table 3 is computed with per-sample SNRs set well below the design-basis fault, precisely to avoid the saturated regime where every detector scores AUC $= 1$. In that weak-fault corner the fusion of an uncorrelated strain-rate and magnetization pair lifts the true-positive rate at 1% false alarm from $0.350/0.683$ to 0.852 — a large gain exactly where a protection trip is decided — and the cross-implementation reproduction to five significant figures is a stronger reproducibility statement than a single headline AUC. The in-winding NV-diamond quench channel is the subject of a dedicated fused strain-rate/magnetization study [nv; kronosfusionenergy.com/publications/quench-detection-and-nv-sensing.html].

NV SENSITIVITY AGAINST THE QUANTUM-SENSING LITERATURE

The in-winding sensitivity of 0.189 $\mathrm{pT\,Hz^{-1/2}}$ sits comfortably inside the envelope established for ensemble NV magnetometry. Degen, Reinhard and Cappellaro frame the shot-noise scaling that equation (9) instantiates [25], and Barry *et al.* analyse the concrete route to broadband ensemble sensitivities in the picotesla-and-below range through improvements to T_2^*, readout fidelity and diamond material [26,27]. Our conservative parameter set is well within what that literature reports, so the sensitivity claim does not rest on an aggressive device — and the $\sim 10^9$ margin over a millitesla quench step is exactly the headroom the radiation-retention curve of figure 9 spends: a channel that starts 10^9 above threshold can degrade by $39\times$ and still contribute to the fused decision. The quench-detection and quantum-sensing threads of the Kronos programme are consolidated in the AI and quantum-control foundation [aiq; kronosfusionenergy.com/publications/ai-quantum-control.html].

RELATION TO LEARNED CONTROL AND THE TWO DISTINCT MAGNETS

The reconstruction stack is complementary to learned control. Degrave *et al.* demonstrated a learned *policy* acting on the present measurement [11]; the stack here is a fast, physics-consistent *estimate* that such a policy can act on instead of the raw signal, inheriting the sub-millisecond equilibrium and the dropout-tolerant state. The estimate feeds a certified safety filter that keeps the learned actuation inside a forward-invariant safe set [cbf; kronosfusionenergy.com/publications/control-barrier-safety-clamp.html]. Finally, the magnet channel treats the breeder centrepost and the burner plug coil as the distinct devices they are: the centrepost governs at a peak field of $16.84\,T$, while the burner plug coil of the separate burner product is a distinct, higher-field magnet at $26.49\,T$; the two are never conflated, and the quench sensor is modelled on each magnet's own load line.

§ 8

Limitations

One honest gate bounds the equilibrium claim. The surrogate is demonstrated as a *single-configuration* reconstruction of the frozen breeder operating point; the parametric generalisation across $(\beta_p, I_p, \mathrm{shape})$ — so that one trained network reconstructs the equilibrium along a trajectory rather than at a point — is the open modelling item, and it follows the same Fourier-feature/physics-informed recipe and the same C1–C3 acceptance gates. Two related items are noted for completeness and are themselves registered work: the graph estimator's reconstruction advantage is unproven on real non-Gaussian signals (the synthetic set is jointly-Gaussian, where linear-MMSE is optimal), and the in-winding NV sensor's neutron tolerance — the $39\times$ degradation retention of criterion C5 — is an acceptance target whose confirming irradiation-plus-charge-state study (SH-037) is the follow-on that closes it. None of these bears on the demonstrated core; each is a defined, testable extension of it.

§ 9

Conclusion

Real-time equilibrium reconstruction for the Hyperion breeder is demonstrated at 0.139% RMS ψ error and $2877\times$ solver speed ($0.87\,ms$) with a 0.32% Grad–Shafranov operator residual that certifies physical consistency, paired with a graph-neural virtual-sensor mesh that discriminates disruption and quench precursors at AUC 0.980 (fused 0.992, TPR 0.852 at 1% false alarm, reproduced across implementations to five significant figures) and an in-winding NV-diamond sensor at 0.189 $\mathrm{pT\,Hz^{-1/2}}$ with per-shot SNR $\sim 10^9$. Around this proven core, the mesh's seeded-dropout training and a model-fusion backup hold the fused precursor above a 0.95 acceptance floor with the NV channel degraded by up to $39\times$, so the protective function survives the radiation-driven degradation of a diagnostic — the condition a nuclear breeder actually imposes. Every estimator is admitted to control authority only through an explicit acceptance inequality, making the route from the demonstrated core to full, dropout-tolerant diagnostic coverage a defined and testable sequence rather than an open question.

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, and the NVIDIA GPU HPC campaign for the training and inference runs.

Reproducibility. The equilibrium, mesh, precursor and quantum-sensor metrics are frozen anchors (KX-L1-A11, KX-L1-A12, KX-L1-A20, KX-L1-A22) in the Kronos Physics De-Risking Register [39] (doi:10.5281/zenodo.22133057) and are regenerable from the surrogate-training, reconstruction and figure-generation archives deposited with this paper under a fixed seed (20260726) and deterministic splits; the fused matched filter is independently reproduced to five significant figures (AUC 0.99166 versus 0.99165). This paper is archived at doi:10.5281/zenodo.22645801 and its official version is hosted at kronosfusionenergy.com/publications/sub-millisecond-neural-equilibrium-reconstruction-robus.html. Training and shadow inference run on the NVIDIA GPU HPC campaign; the reduced-order and analytic layers run on the in-house HPC edge node.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645801. The reconstruction inputs, ROC/NV parameters and metric computations for figures 2–9 are archived with the deposit and reference the Kronos 2026 series, including the AI and quantum-control foundation [33]. All quantitative values trace to the register [39] 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

APPARATUS

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 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

2 A Pironti and M Walker, Fusion, tokamaks, and plasma control: an introduction and tutorial, *IEEE Control Syst. Mag.* **25**(5) 30 (2005). doi:10.1109/MCS.2005.1512794

3 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

4 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

5 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

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

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 S Joung *et al*, Deep neural network Grad–Shafranov solver constrained with measured magnetic signals, *Nucl. Fusion* **60** 016034 (2020).

10 L L Lao *et al*, Application of machine learning and artificial intelligence to extend EFIT equilibrium reconstruction, *Plasma Phys. Control. Fusion* **64** 074001 (2022).

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

12 J Kates-Harbeck, A Svyatkovskiy and W Tang, Predicting disruptive instabilities in controlled fusion plasmas through deep learning, *Nature* **568** 526 (2019). doi:10.1038/s41586-019-1116-4

13 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

14 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

15 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

16 M Tancik *et al*, Fourier features let networks learn high frequency functions in low dimensional domains, *Adv. Neural Inf. Process. Syst. (NeurIPS)* **33** 7537 (2020).

17 J Gilmer, S S Schoenholz, P F Riley, O Vinyals and G E Dahl, Neural message passing for quantum chemistry, *Proc. 34th Int. Conf. on Machine Learning (ICML)* **70** 1263 (2017).

18 P W Battaglia *et al*, Relational inductive biases, deep learning, and graph networks, (2018).

19 T N Kipf and M Welling, Semi-supervised classification with graph convolutional networks, *Int. Conf. on Learning Representations (ICLR)* (2017).

20 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

21 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

22 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

23 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

24 B Dudson *et al*, FreeGS: free-boundary Grad–Shafranov equilibrium solver (2020). doi:10.5281/zenodo.3745669

25 C L Degen, F Reinhard and P Cappellaro, Quantum sensing, *Rev. Mod. Phys.* **89** 035002 (2017). doi:10.1103/RevModPhys.89.035002

26 J F Barry, J M Schloss, E Bauch, M J Turner, C A Hart, L M Pham and R L Walsworth, Sensitivity optimization for NV-diamond magnetometry, *Rev. Mod. Phys.* **92** 015004 (2020). doi:10.1103/RevModPhys.92.015004

27 L Rondin, J-P Tetienne, T Hingant, J-F Roch, P Maletinsky and V Jacques, Magnetometry with nitrogen-vacancy defects in diamond, *Rep. Prog. Phys.* **77** 056503 (2014). doi:10.1088/0034-4885/77/5/056503

28 M Marchevsky, Quench detection and protection for high-temperature superconductor accelerator magnets, *Instruments* **5** 27 (2021). doi:10.3390/instruments5030027

29 E E Salazar *et al*, Fiber optic quench detection for large-scale HTS magnets demonstrated on VIPER cable during high-fidelity testing at the SULTAN facility, *Supercond. Sci. Technol.* **34** 035027 (2021). doi:10.1088/1361-6668/abdba8

30 R Teyber, J Weiss, M Marchevsky, S Prestemon and D van der Laan, Current distribution monitoring enables quench and damage detection in superconducting fusion magnets, *Sci. Rep.* **12** 22503 (2022). doi:10.1038/s41598-022-26592-2

31 S Hahn, Y Kim, D K Park, K Kim, J P Voccio, J Bascu\ n\'an and Y Iwasa, No-insulation multi-width winding technique for high temperature superconducting magnet, *Appl. Phys. Lett.* **103** 173511 (2013). doi:10.1063/1.4826217

32 S Hahn *et al*, 45.5-tesla direct-current magnetic field generated with a high-temperature superconducting magnet, *Nature* **570** 496 (2019). doi:10.1038/s41586-019-1293-1

33 P I Ford, AI and quantum computing for fusion: ML surrogates, real-time control, and a resource-honest quantum frontier, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645702. Official page: kronosfusionenergy.com/publications/ai-quantum-control.html

34 P I Ford, Free-boundary equilibrium and ideal-MHD stability of a compact spherical-tokamak breeder, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645740. Official page: kronosfusionenergy.com/publications/ideal-mhd-stability.html

35 P I Ford, A certified control-barrier safety filter for fusion control, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645716. Official page: kronosfusionenergy.com/publications/control-barrier-safety-clamp.html

36 P I Ford, The plasma diagnostic suite and real-time state estimation for a compact spherical-tokamak breeder, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645943. Official page: kronosfusionenergy.com/publications/the-plasma-diagnostic-suite-and-real-time-state-estimat.html

37 P I Ford, Neural-operator flux surrogates for control-latency transport, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645803. Official page: kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html

38 P I Ford, Fused strain-rate and magnetization quench detection with in-winding NV-diamond sensing for no-insulation REBCO breeder magnets, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645752. Official page: kronosfusionenergy.com/publications/quench-detection-and-nv-sensing.html

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

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

Sub-Millisecond Neural Equilibrium Reconstruction Robust to Sensor Dropout: A Virtual-Sensor Mesh for Radiation-Degraded Diagnostics in a Compact Spherical-Tokamak Breeder

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

Zenodo · doi:10.5281/zenodo.22645801

Read on kronosfusionenergy.com

The Physics De-Risking Register