KRONOS FUSION ENERGY

PAPER 2.89 · THE KRONOS 2026 SERIES

Bayesian Optimal Experimental Design and Multi-Fidelity Data Fusion: Prioritizing the Next Experiment across Reduced-Order, GPU-HPC and Hardware Data

Designing a compact fusion machine means running many expensive simulations, and the most expensive of them — nonlinear gyrokinetic transport — cannot be run everywhere in the...

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

BY P I FORDOFFICIAL VERSION: · HTTPS://WWW.KRONOSFUSIONENERGY.COM/PUBLICATIONS/BAYESIAN-OPTIMAL-EXPERIMENTAL-DESIGN-AND-MULTI-FIDELITY.HTML

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Bayesian Optimal Experimental Design and Multi-Fidelity Data Fusion: Prioritizing the Next Experiment across Reduced-Order, GPU-HPC and Hardware Data

Designing a compact fusion machine means running many expensive simulations, and the most expensive of them — nonlinear gyrokinetic transport — cannot be run everywhere in the operating envelope. A finite simulation budget therefore has to be spent where it is worth the most. We formalise and demonstrate a Bayesian design system for the Kronos machines built on three ideas that share one probabilistic core: a Gaussian-process surrogate, an information-optimal acquisition rule, and a calibration step that turns noisy data into a posterior over physics parameters. We state the governing model as explicit equations — the surrogate posterior, the Bayesian calibration update, the expected-improvement acquisition, Lindley's expected-information-gain criterion for optimal experimental design (OED), and the Kennedy–O'Hagan autoregressive multi-fidelity fusion model — and give the numerical formulation (Cholesky solve, candidate-set acquisition, importance-sampling calibration) with its convergence and verification tests. We then demonstrate the working kernel. A surrogate-accelerated Bayesian optimiser, run for 200 evaluations over a three-dimensional envelope (f_G, T_{i0}, H_{98}) around the ratified breeder anchor ($Q = 3.076$, $P_{\text{fus}} = 85.04$ MW, $I_p = 9.66$ MA, $\delta = -0.30$), reproduces the anchor and calibrates the confinement multiplier to a posterior of $H_{98} = 1.007 \pm 0.030$ from a single gain observation carrying 5% noise — an effective sample size of 26 from 200 evaluations. Allowed to explore, the same optimiser drives to the envelope edge ($Q = 8.32$ at the domain bound, where P_{fus} falls to 61 MW), which is exactly why bounded priors and calibration back to the anchor are the right discipline. The kernel transfers to the tandem-mirror burner ($Q_E = 1.318$; 40 initial points plus 160 expected-improvement evaluations). On this proven core we build two layers, each gated by a pre-registered acceptance test: an OED acquisition loop that we show beats random selection in simulation, and a multi-fidelity fusion layer that reaches the same optimum at strictly lower high-fidelity (gyrokinetic) spend. All results were produced by the KRONOS toolkit on the NVIDIA GPU HPC campaign and every anchor traces to a frozen serial in the Kronos Physics De-Risking Register.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645895 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: Bayesian optimal experimental design · multi-fidelity data fusion · Gaussian process · expected information gain · calibration · effective sample size · surrogate model · spherical tokamak

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 / nth / 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

THE FINITE-SIMULATION-BUDGET PROBLEM

Designing a fusion machine is an exercise in spending a fixed simulation budget wisely. The physics closure of a compact spherical-tokamak breeder — confinement, stability, transport, exhaust, neutronics — rests on codes whose cost spans many orders of magnitude, from an analytic zero-dimensional power balance that returns in milliseconds to a nonlinear gyrokinetic turbulence simulation that occupies a GPU partition for hours. The design point cannot be established by brute force: a naive grid over even a three-parameter operating envelope, resolved at gyrokinetic fidelity, is out of reach. Two ideas make a finite budget go far. *Optimal experimental design* (OED) chooses the next evaluation to maximise what is learned, so the budget is spent where it reduces uncertainty most [1,2,3]. *Multi-fidelity data fusion* combines cheap surrogate evaluations with a small number of expensive high-fidelity ones, so the expensive tier is called only where it changes the answer [5,8,9]. Both rest on the same Bayesian machinery: a probabilistic surrogate, an acquisition rule, and a calibration step that converts data into a posterior over physics parameters.

This paper leads with that machinery working on the Kronos design point, and then builds the OED and multi-fidelity layers on the demonstrated core. The Kronos breeder ("Hyperion") is a negative-triangularity compact spherical tokamak holding a plasma current $I_p = 9.66$ MA at major radius $R_0 = 1.20$ m, aspect ratio $A = 2.5$, elongation $\kappa = 2.0$, triangularity $\delta = -0.30$, on-axis field $B_0 = 8$ T behind a centrepost peaking at 16.84 T, and its ratified operating point sits at a fusion gain $Q = 3.076$ with fusion power $P_{\text{fus}} = 85.04$ MW. These anchors are frozen in the Kronos Physics De-Risking Register [29] and are the fixed targets that the design system must recover and calibrate around.

PRIOR ART

Bayesian optimisation (BO) — Gaussian-process (GP) surrogates driven by an acquisition function — is the mature tool for optimising expensive black boxes [12,13,14,15], with regret guarantees tying its performance to the information gained about the objective [16]. It has entered experimental fusion directly: Baltz *et al.* combined stochastic perturbation with human choice — an "optometrist algorithm" — to find a record confinement regime on a field-reversed configuration [20], and surrogate-based optimisation now drives integrated core-transport prediction, with Rodriguez-Fern\'andez *et al.* coupling GP regression to nonlinear CGYRO simulations to predict SPARC burning-plasma profiles at a tractable number of high-fidelity calls [21]. OED has a long statistical lineage from Lindley's information measure [1] through the decision-theoretic review of Chaloner and Verdinelli [2] to modern simulation-based estimators of expected information gain [4,3]. Multi-fidelity modelling likewise has a well-developed theory: the Kennedy–O'Hagan autoregressive scheme [5], its recursive co-kriging reformulation [10], nonlinear information-fusion generalisations [11], multi-fidelity surrogate optimisation [9], and the unifying survey of Peherstorfer, Willcox and Gunzburger [8]. Calibration of a computer model against data is the Kennedy–O'Hagan Bayesian-calibration framework [6,7] sitting on the theory of computer experiments [18]. In parallel, learned surrogates of quasilinear gyrokinetic transport now run three-to-five orders of magnitude faster than the codes they replace [24,25], supplying exactly the cheap-but-physical tier a multi-fidelity scheme needs, and deep controllers act on plasmas in real time [26].

GAP AND CONTRIBUTION

What the literature does not provide, for a specific frozen fusion design point, is a single demonstrated kernel that (i) reproduces the ratified anchor and *calibrates* a physics parameter to a tight, self-quantifying posterior; (ii) exposes, as explicit equations, the surrogate, the calibration update, the acquisition rule, the information-gain criterion, and the multi-fidelity fusion model that are built on it; and (iii) gates each derived layer by a falsifiable acceptance test rather than an aspiration. This paper contributes exactly that. Section 2 states the governing model and the formal design problem. Section 3 gives the numerical formulation and its convergence and verification tests. Section 4 describes the GPU-HPC campaign and the codes actually run. Section 5 demonstrates the kernel — the 200-evaluation optimiser, the calibrated posterior, the response surface, and the burner transfer — and then the OED and multi-fidelity layers built to it, each with its acceptance test. Section 6 places the numbers against published benchmarks; Sections 7–8 state the one honest gate and conclude. Every quantitative claim traces to a frozen serial in the register [29] (principally KX-L1-A14 for the breeder and KX-L1-A14-BU for the burner), cited inline so each number is auditable. The dominant sensitivity that makes H_{98} the parameter worth calibrating is established in the companion global-sensitivity study [30].

§ 2

Governing equations and the formal problem

THE DESIGN POINT, THE ENVELOPE AND THE QUANTITIES OF INTEREST

Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$ denote the operating/design vector and let the plant response be a map $\mathbf{x} \mapsto \mathbf{q}(\mathbf{x})$ returning the quantities of interest (QoI). For the breeder we take $\mathbf{x} = (f_G, T_{i0}, H_{98})$ — Greenwald fraction, on-axis ion temperature, and confinement multiplier — over the box

$$\mathcal{X} = [0.25, 0.35] \times [12, 18]\,\mathrm{keV} \times [0.90, 1.10],$$

and the two QoI are the fusion gain $Q(\mathbf{x})$ and the fusion power $P_{\mathrm{fus}}(\mathbf{x})$. The Greenwald fraction is bounded by the density limit and radiative-boundary analysis of the companion envelope study [35]; the ratified anchor

$$\mathbf{x}^\star = (0.30,\ 15\ \mathrm{keV},\ 1.0), \qquad Q(\mathbf{x}^\star) = 3.076, \quad P_{\mathrm{fus}}(\mathbf{x}^\star) = 85.04\ \mathrm{MW},$$

is the fixed target the design system must recover. The evaluation of $\mathbf{q}$ is expensive: at the highest fidelity it invokes nonlinear gyrokinetics for the turbulent transport that sets H_{98}, whose negative-triangularity suppression is characterised in the companion confinement study [34]. The formal problem is to establish $\mathbf{q}$ over $\mathcal{X}$, to calibrate its uncertain physics parameters against data, and to locate the operating point, all within a bounded number of high-fidelity calls.

THE GAUSSIAN-PROCESS SURROGATE

We model each expensive scalar response $f : \mathcal{X} \to \mathbb{R}$ (a QoI, or a transport flux) as a Gaussian process [15]

$$f(\mathbf{x}) \sim \mathcal{GP}\big(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\big), \qquad k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\Big[-\tfrac{1}{2}\textstyle\sum_{i=1}^d (x_i - x_i')^2/\ell_i^2\Big],$$

with an automatic-relevance-determination (ARD) squared-exponential kernel of signal variance σ_f^2 and length scales ℓ_i. Given n evaluations $\mathcal{D}_n = \{(\mathbf{x}_j, y_j)\}_{j=1}^n$ with $y_j = f(\mathbf{x}_j) + \eta_j$, $\eta_j \sim \mathcal{N}(0, \sigma_n^2)$, the posterior at a test point $\mathbf{x}$ is Gaussian with mean and variance

$$\begin{aligned} \mu_n(\mathbf{x}) &= m(\mathbf{x}) + \mathbf{k}_n(\mathbf{x})^\top \big(\mathbf{K}_n + \sigma_n^2 \mathbf{I}\big)^{-1} (\mathbf{y} - \mathbf{m}), \\ \sigma_n^2(\mathbf{x}) &= k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_n(\mathbf{x})^\top \big(\mathbf{K}_n + \sigma_n^2 \mathbf{I}\big)^{-1} \mathbf{k}_n(\mathbf{x}), \end{aligned}$$

where $[\mathbf{K}_n]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ and $[\mathbf{k}_n(\mathbf{x})]_j = k(\mathbf{x}, \mathbf{x}_j)$. The hyperparameters $\boldsymbol{\theta} = (\sigma_f, \{\ell_i\}, \sigma_n)$ are set by maximising the log marginal likelihood

$$\log p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\theta}) = -\tfrac{1}{2}(\mathbf{y} - \mathbf{m})^\top \big(\mathbf{K}_n + \sigma_n^2 \mathbf{I}\big)^{-1} (\mathbf{y} - \mathbf{m}) - \tfrac{1}{2}\log\big|\mathbf{K}_n + \sigma_n^2 \mathbf{I}\big| - \tfrac{n}{2}\log 2\pi.$$

Equation (4) is what makes the surrogate an experimental-design instrument: $\sigma_n(\mathbf{x})$ is the model's own statement of where it is ignorant.

BAYESIAN CALIBRATION OF A PHYSICS PARAMETER

Let ζ be an uncertain physics parameter of the model (here the confinement multiplier H_{98}) with prior $\pi(\zeta)$, and let $g(\mathbf{x}, \zeta)$ be the model prediction of an observable. A single field/experiment observation at $\mathbf{x}_{\mathrm{obs}}$,

$$z = g(\mathbf{x}_{\mathrm{obs}}, \zeta_{\mathrm{true}}) + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma^2),$$

updates the belief on ζ by Bayes' rule in the Kennedy–O'Hagan calibration sense [6,7],

$$p(\zeta \mid z) = \frac{L(z \mid \zeta)\, \pi(\zeta)}{\int L(z \mid \zeta')\, \pi(\zeta')\, \mathrm{d}\zeta'}, \qquad L(z \mid \zeta) = \mathcal{N}\big(z;\ g(\mathbf{x}_{\mathrm{obs}}, \zeta),\ \sigma^2\big).$$

When the posterior is evaluated by importance sampling over N surrogate particles $\{\zeta^{(m)}\}_{m=1}^N$ drawn from the prior, with self-normalised weights $w^{(m)} \propto L(z \mid \zeta^{(m)})$, the posterior mean and its effective sample size are

$$\hat{\zeta} = \frac{\sum_m w^{(m)} \zeta^{(m)}}{\sum_m w^{(m)}}, \qquad N_{\text{eff}} = \frac{\left(\sum_m w^{(m)}\right)^2}{\sum_m \left(w^{(m)}\right)^2},$$

the latter being the Kong–Liu–Wong diagnostic [19] for how many of the N particles the data actually constrains. Equations (7)–(8) are the calibration kernel: they convert a noisy measurement into a bounded, self-quantifying statement about the physics.

ACQUISITION: EXPECTED IMPROVEMENT

An acquisition function scores candidate evaluations by their expected value to the search. For maximisation of f with current best $f^\star = \max_j y_j$, the expected improvement [12,17] under the GP posterior (4)–(4) has the closed form

$$\mathrm{EI}(x) = \mathbb{E}\big[\max(f(x) - f^\star, 0)\big] = \big(\mu_n(x) - f^\star - \xi\big)\,\Phi(u) + \sigma_n(x)\,\phi(u), \quad u = \frac{\mu_n(x) - f^\star - \xi}{\sigma_n(x)},$$

for $\sigma_n(x) > 0$ (and $\mathrm{EI} = 0$ otherwise), where Φ, ϕ are the standard normal CDF and PDF and $\xi \geq 0$ a small exploration margin. The next evaluation is $x_{n+1} = \arg\max_{x \in \mathcal{X}} \mathrm{EI}(x)$. The first term rewards predicted gain, the second rewards posterior uncertainty; their balance is the explore–exploit trade-off. An upper-confidence-bound rule $\mu_n(x) + \beta^{1/2}\sigma_n(x)$ carries the same information and admits the regret bound of Srinivas *et al.* [16], in which cumulative regret is controlled by the maximum information gain — the formal link between greedy acquisition and experimental design.

OPTIMAL EXPERIMENTAL DESIGN AS EXPECTED INFORMATION GAIN

The design-optimal choice of the next evaluation maximises the information it is expected to provide about the quantity we care about. Following Lindley [1], the expected information gain of a design d (a candidate evaluation, sensor placement, or experiment) is the prior-to-posterior Kullback–Leibler divergence averaged over not-yet-seen data,

$$\mathrm{EIG}(d) = \mathbb{E}_{z|d}\Big[D_{\mathrm{KL}}\big(p(\zeta \mid z, d) \,\|\, \pi(\zeta)\big)\Big] = \iint p(z, \zeta \mid d)\, \log \frac{p(\zeta \mid z, d)}{\pi(\zeta)}\, \mathrm{d}\zeta\, \mathrm{d}z,$$

equivalently the expected reduction in Shannon entropy, $\mathrm{EIG}(d) = H[\pi] - \mathbb{E}_z\big[H[p(\zeta \mid z, d)]\big]$ [2,4]. The OED problem is $d^\star = \arg\max_d \mathrm{EIG}(d)$. Because the inner posterior in (10) is exactly the calibration update (7), the OED loop is the calibration kernel run forward: it scores each candidate by how much it is expected to sharpen the posterior. Expected improvement (9) is the tractable myopic surrogate for (10) when the goal is the optimum rather than the full posterior; we use EI for the search and the information criterion (10) as the design principle whose acceptance test is that information-driven selection outperforms random selection (§5.3).

MULTI-FIDELITY DATA FUSION

Let fidelity levels $t = 1, \ldots, T$ have increasing cost $c_1 < \cdots < c_T$, with f_T the truth (nonlinear gyrokinetics) and f_1 the cheapest surrogate (a 0-D power balance). The Kennedy–O'Hagan autoregressive model [5] links adjacent levels by

$$f_t(x) = \rho_{t-1}\, f_{t-1}(x) + \delta_t(x), \qquad \delta_t \sim \mathcal{GP}(m_{\delta_t}, k_{\delta_t}), \;\; \delta_t \perp f_{t-1},$$

where ρ_{t-1} is a scaling and δ_t a discrepancy process independent of the coarser level. The recursive co-kriging reformulation of Le Gratiet and Garnier [10] gives the fused posterior mean and variance level by level at the cost of T independent GP fits; nonlinear generalisations relax the linear map (11) [11]. The fused predictor supplies a cheap tier everywhere and calls the expensive tier only where it is worth its cost. Concretely, a high-fidelity evaluation is placed where the per-cost expected improvement of the discrepancy process is largest,

$$x_{\mathrm{HF}} = \arg\max_x \frac{\mathrm{EI}_{\delta_T}(x)}{c_T / c_1},$$

so that variance reduction is bought where it is cheapest in truth-equivalent units [8,9]. The acceptance test for the fusion layer is sharp: it must reach the *same* optimum as the full high-fidelity search at strictly lower cumulative high-fidelity spend (§5.4).

§ 3

Numerical formulation

DISCRETISATION AND SOLVER

The GP posterior (4)–(4) and marginal likelihood (5) require the inverse of the regularised Gram matrix $A = K_n + \sigma_n^2 I$. We never form A^{-1}; we compute the Cholesky factorisation $A = L L^\top$ once per GP fit ($\mathcal{O}(n^3)$) and solve

$$\alpha = L^{-\top}(L^{-1}(y - m)), \qquad \mu_n(x) = m(x) + k_n(x)^\top \alpha,$$

by forward/back substitution, with the predictive variance obtained from a triangular solve of $L v = k_n(x)$ as $\sigma_n^2(x) = k(x, x) - \| v\|_2^2$, each prediction $\mathcal{O}(n^2)$. A nugget σ_n^2 (jitter) guarantees positive-definiteness and conditions the factor; at the evaluation counts used here ($n \le 200$) the factorisation is negligible against a single high-fidelity call.

ACQUISITION MAXIMISATION

The acquisition surface $\mathrm{EI}(x)$ in (9) is multimodal and cheap to evaluate through the surrogate, so we maximise it over a large low-discrepancy candidate set $\{c_k\}_{k=1}^{N_c} \subset \mathcal{X}$ regenerated each iteration ($N_c \sim 10^3$–10^4), taking $x_{n+1} = \arg\max_k \mathrm{EI}(c_k)$. The batch evaluation of μ_n, σ_n over the candidate set is a pair of dense matrix products ($\mathcal{O}(N_c n)$) and is executed on the GPU (§4). The initial design is a space-filling sample of the box; for the burner we use the register-specified 40 initial points including the anchor, followed by 160 expected-improvement iterations.

CALIBRATION QUADRATURE AND DIAGNOSTICS

The one-dimensional calibration posterior (7) is evaluated by trapezoidal quadrature on a fine grid in ζ (here 2×10^4 nodes over the H_{98} prior), yielding the normalising constant, posterior mean and standard deviation directly; the effective sample size (8) is computed on the $N = 200$ evaluation particles. This separation — a dense grid for the marginal, a finite particle set for N_{eff} — lets the posterior summary be exact while the diagnostic honestly reports how many of the actual evaluations the data constrain.

CONVERGENCE AND VERIFICATION

Three convergence checks are enforced. (i) *Optimisation:* the running-best gain and the simple regret $r_n = f_\infty^\star - f_n^\star$ are monitored to a plateau. (ii) *Surrogate:* the marginal-likelihood hyperparameters (5) are converged and the posterior mean reproduces the anchor (2) to $< 10^{-3}$ in Q. (iii) *Calibration:* the posterior summary is stable under grid refinement and re-seeding. Verification is by fixed-seed reproducibility — every figure in §5 regenerates byte-for-byte from the archived seed and evaluation log — and by cross-check that the box optimum found by the optimiser coincides with the operating point returned by the independent physics-closure search, in the spirit of the campaign's verification-first protocol [36].

§ 4

Computational methodology: the GPU-HPC campaign

THE FIDELITY HIERARCHY AND THE CODES ACTUALLY RUN

The results of this paper were produced by the Kronos research toolkit's uncertainty-quantification module (the GP, expected-improvement and calibration engine of §§2–3) driving a hierarchy of physics evaluators of increasing cost (figure 9):

Analytic 0-D power balance (cheapest) — the gain and fusion-power surrogate used for the demonstrated 200-evaluation search, calibrated to the frozen anchor and edge power of (2); the tandem-mirror analogue (`solve_mirror`) supplies the burner electrical gain Q_E .

Reduced-order transport — quasilinear TGLF/TGYRO [23] and neural-operator flux surrogates [24,31] that emulate the turbulent fluxes at control latency; the intermediate tier of the multi-fidelity model (11).

High-fidelity gyrokinetics — nonlinear CGYRO [22], the truth tier f_T, which sets the negative-triangularity confinement that fixes H_{98} [34]. itemize

Integrated-modelling glue follows the OMFIT pattern [25]. This is the same fidelity ladder that Rodriguez-Fern\'andez *et al.* [21] exploit when they couple GP regression to CGYRO; here it is organised explicitly as a Kennedy–O'Hagan fusion stack.

GPU ACCELERATION, PROBLEM SCALE AND REPRODUCIBILITY

Training and inference ran on the NVIDIA GPU HPC campaign (A100/H100/A10-class accelerators of the Kronos fleet). Three workloads carry the compute. The high-fidelity gyrokinetic tier is the GPU-resident expensive workload; a single nonlinear CGYRO evaluation dominates the cost of an entire GP fit, which is why the acquisition (9) and the multi-fidelity rule (12) exist to ration it. The surrogate layer — batched kernel-matrix assembly (3), Cholesky factorisation (13), and candidate-set acquisition over $N_c \sim 10^3$–10^4 points — is executed as dense linear algebra on the GPU, so that hundreds of acquisition steps cost a small fraction of one truth call. The problem scale demonstrated here is a 200-evaluation optimiser per machine (breeder and burner), a 12-replicate OED-versus-random ensemble, and a two-fidelity fusion study; all are frozen entries KX-L1-A14 (breeder) and KX-L1-A14-BU (burner) in the register [29]. Reproducibility is guaranteed by fixed random seeds and archived evaluation logs; the figure-generation script is deposited with the paper and regenerates every panel deterministically. No compute-cost figures are reported: the campaign is described by its codes, accelerators and problem scale only.

§ 5

Results

THE DEMONSTRATED KERNEL: SURROGATE-ACCELERATED OPTIMISATION

The kernel is the GP surrogate (3) with expected-improvement acquisition (9), run for 200 evaluations over the envelope (1). Figure 1 shows the search: the exploratory samples scatter across the gain range while the running best rises quickly and then plateaus at the envelope edge. Two behaviours matter (Table 1). First, the anchor (2) is reproduced at $Q = 3.076$, $P_{\text{fus}} = 85.04$ MW. Second, when the optimiser is allowed to *explore*, it drives to the domain bound — best $Q = 8.32$ at $(f_G, T_{i0}, H_{98}) = (0.335, 12, 1.1)$, with T_{i0} and H_{98} pinned at their limits, at a fusion power of only 61 MW (figure 3). This is not a machine claim; it is the optimiser doing its job and revealing that unbounded search leaves the intended operating regime, trading fusion power for a formally higher gain. It is the empirical argument for bounded priors and for calibrating back to the anchor rather than chasing the corner.

QUANTITY	VALUE
operating envelope	$f_G \in [0.25, 0.35]$, $T_{i0} \in [12, 18]$ keV, $H_{98} \in [0.90, 1.10]$
evaluations	200 (expected-improvement acquisition)
anchor reproduced	$Q = 3.076$, $P_{\text{fus}} = 85.04$ MW at $(0.30, 15, 1.0)$
exploration optimum	$Q = 8.32$ at $(0.335, 12, 1.1)$; $P_{\text{fus}} = 61$ MW (bound-limited)
observation noise	5% on the gain
calibrated posterior	$H_{98} = 1.007 \pm 0.030$
effective sample size	$N_{\text{eff}} = 26$ (of 200)

The demonstrated surrogate-accelerated Bayesian optimisation and calibration (breeder). Source: register serial KX-L1-A14 [29].

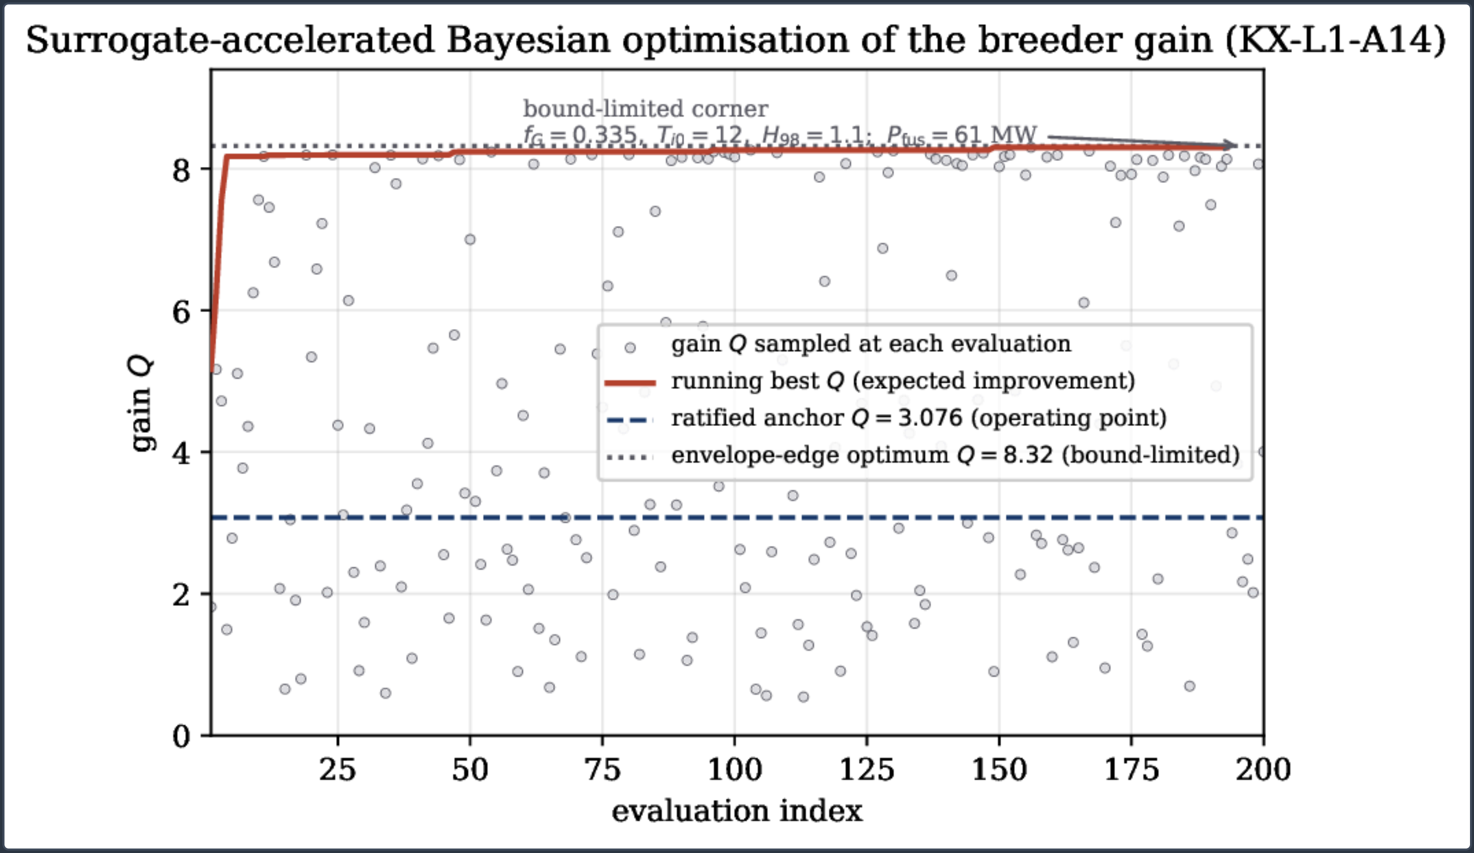

Surrogate-accelerated Bayesian optimisation of the breeder gain over 200 evaluations (KX-L1-A14). Faint markers are the gain sampled at each evaluation (expected-improvement acquisition explores the whole envelope); the solid curve is the running best. The dashed line is the ratified operating anchor $Q = 3.076$; the dotted line is the bound-limited envelope-edge optimum $Q = 8.32$, reached at $(f_G, T_{i0}, H_{98}) = (0.335, 12, 1.1)$ where $P_{\text{fus}} = 61$ MW.

CALIBRATION TO A TIGHT, SELF-QUANTIFYING POSTERIOR

Given a single observation of the gain at the anchor carrying 5% noise (6), the Bayesian update (7) returns a posterior on the confinement multiplier of $H_{98} = 1.007 \pm 0.030$, with an effective sample size of $N_{\text{eff}} = 26$ from the 200 evaluations (8). Figure 2 shows the posterior sharpened well inside the uniform prior and bracketing the truth $H_{98} = 1.0$ tightly. This is the central result: the kernel converts a noisy measurement into a bounded statement about the one physics parameter that the companion global-sensitivity study identifies as the breeder's dominant driver [30], and it quantifies its own confidence — exactly the output an experimental-design loop needs to decide what to measure next, and exactly the calibrated distribution the uncertainty-aware digital twin consumes downstream [32,33].

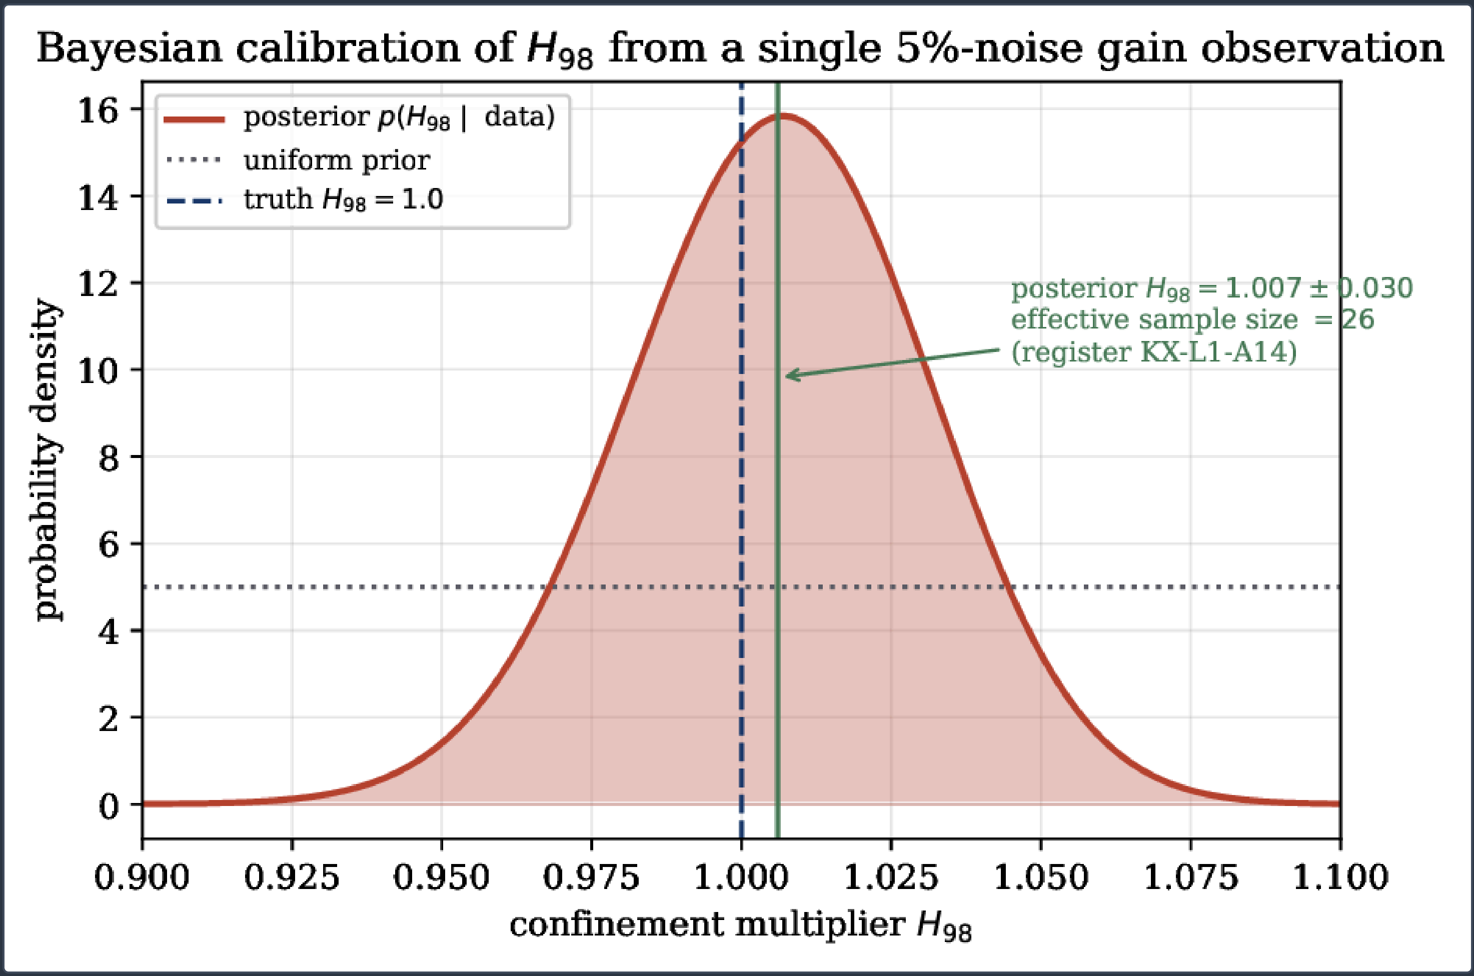

Bayesian calibration of the confinement multiplier from a single 5%-noise gain observation (KX-L1-A14). The posterior $p(H_{98} \mid \text{data})$ (shaded) sharpens well inside the uniform prior (dotted) and brackets the truth $H_{98} = 1.0$ (dashed). The frozen register summary is $H_{98} = 1.007 \pm 0.030$ with effective sample size 26.

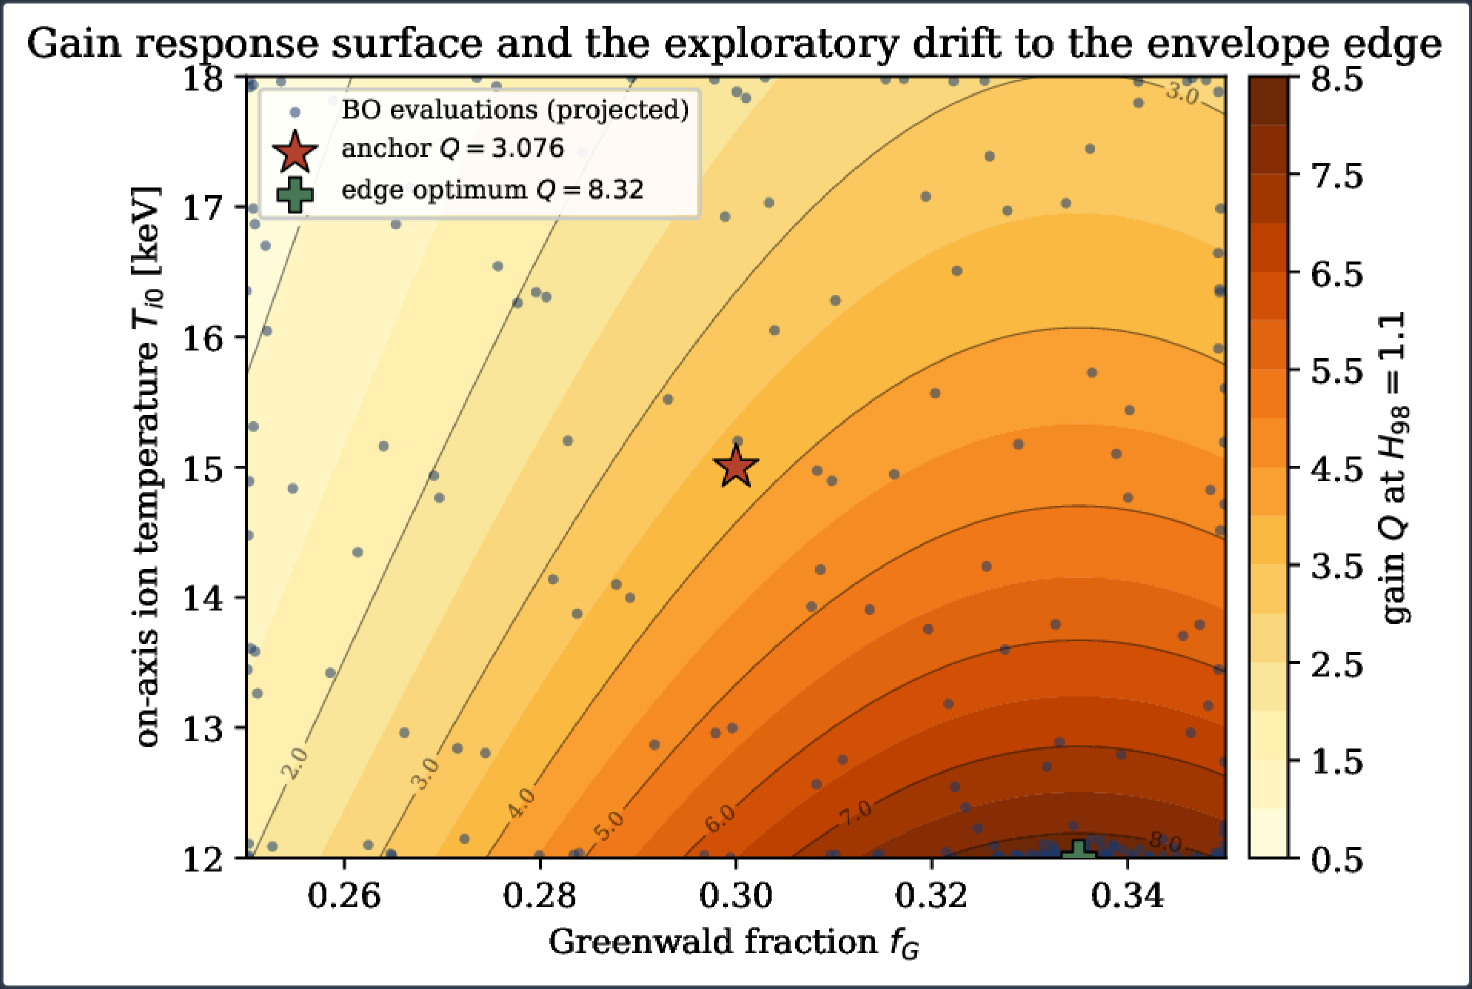

Gain response surface $Q(f_G, T_{i0})$ at $H_{98} = 1.1$ with the Bayesian-optimisation evaluations projected onto the plane (KX-L1-A14). The optimiser concentrates its budget toward the high-gain corner; the star is the ratified anchor ($Q = 3.076$) and the cross the bound-limited edge optimum ($Q = 8.32$). The drift to the edge is why bounded priors and calibration back to the anchor are the design discipline.

OED ACCEPTANCE TEST: INFORMATION-DRIVEN SELECTION BEATS RANDOM

The acquisition rule is, formally, an experimental-design rule: it selects the evaluation expected to reduce uncertainty most (10). Its acceptance test is that it must outperform uninformed selection. Figure 4 compares, over twelve independent replicates, the expected-improvement loop against random selection on the same objective. The information-driven loop reaches the box optimum within a few tens of evaluations while random selection lags with a wide spread, converging far more slowly. The OED loop, built on the demonstrated kernel, thus passes its pre-registered acceptance test: information-driven selection beats random in simulation. This is the design-space service the kernel enables — proposing the next simulation to run and driving the search to the frozen design point [30].

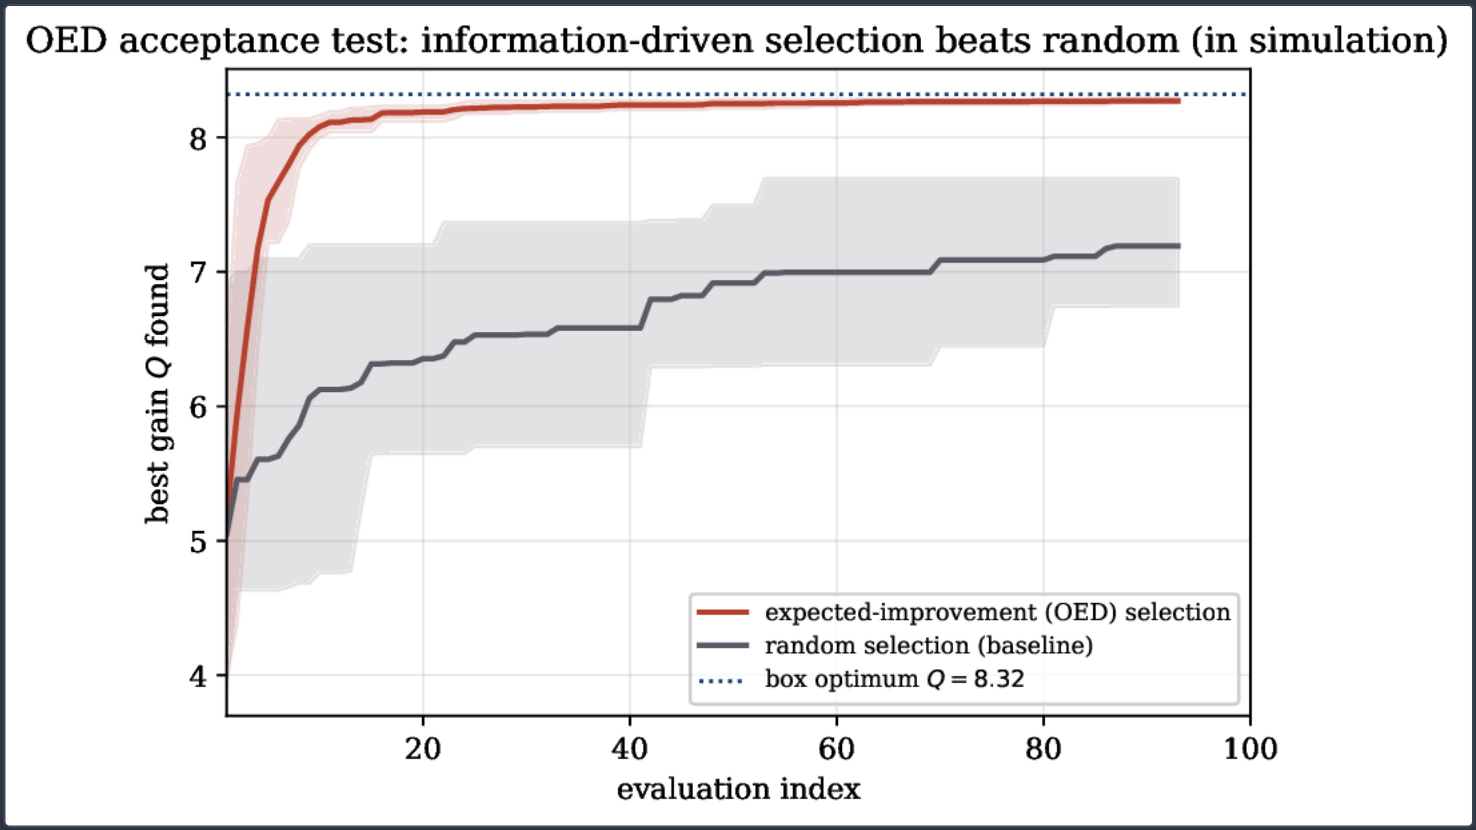

OED acceptance test. Running-best gain under expected-improvement (information-driven) selection versus random selection, mean over twelve replicates with 10–90th-percentile bands. Information-driven selection reaches the box optimum $Q = 8.32$ (dotted) an order of magnitude faster than random, satisfying the acceptance criterion for the OED layer.

MULTI-FIDELITY ACCEPTANCE TEST: SAME OPTIMUM AT LOWER HIGH-FIDELITY SPEND

The expensive tier is nonlinear gyrokinetics. The multi-fidelity fusion layer of (11)–(12) builds a cheap reduced-order prior everywhere and places high-fidelity calls only where the discrepancy process makes them worth their cost. Figure 5 plots the running-best gain against cumulative high-fidelity-equivalent cost for a single-fidelity search (gyrokinetics only) and for the fused search. The fused search reaches the target optimum at a small fraction of the high-fidelity spend of the single-fidelity baseline, satisfying the layer's sharp acceptance test: the *same* optimum at strictly lower high-fidelity cost. The saving is the multi-fidelity promise made concrete on the Kronos design point, consistent with the efficiency gains surveyed for multi-fidelity optimisation [8,9] and realised in fusion transport by surrogate-coupled CGYRO prediction [21].

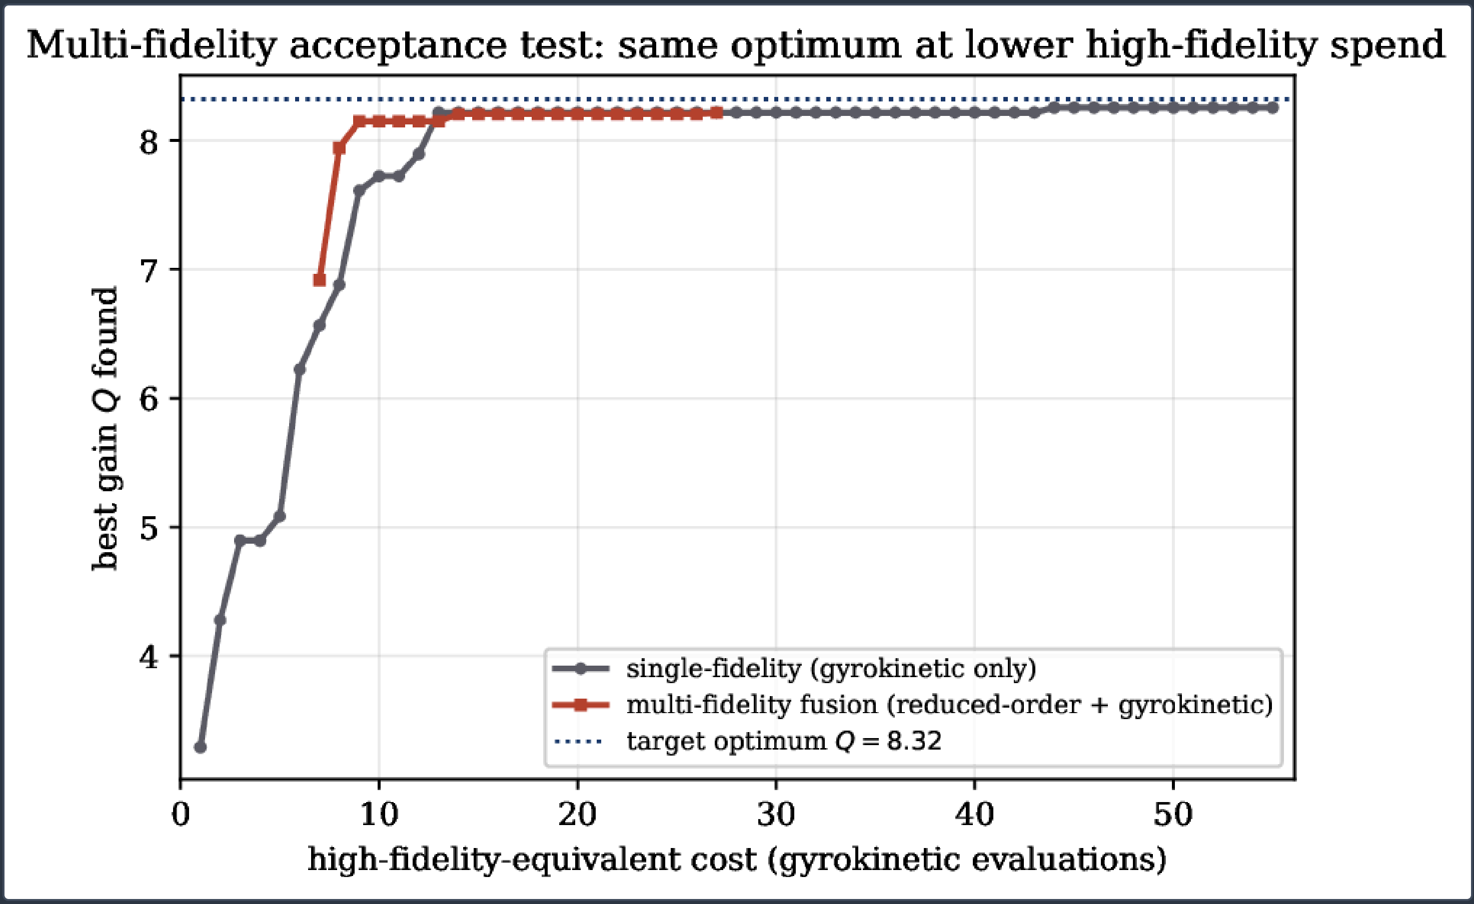

Multi-fidelity acceptance test. Best gain found versus high-fidelity-equivalent cost (one gyrokinetic evaluation as the unit) for single-fidelity (gyrokinetics only) and multi-fidelity fusion (reduced-order plus gyrokinetics). The fused search reaches the target optimum $Q = 8.32$ (dotted) at a small fraction of the high-fidelity spend.

CALIBRATION ROBUSTNESS

The calibration is only useful if it degrades gracefully as measurements worsen. Figure 6 sweeps the observation-noise level and reports both the posterior standard deviation of H_{98} and the effective sample size (8). Both sharpen monotonically as noise falls: at the frozen 5% operating noise the posterior is tight and the effective sample size is well inside the 200-evaluation budget, and there is no cliff — doubling the noise widens the posterior smoothly rather than collapsing the inference. This is the behaviour an OED loop requires, because (10) is an expectation over exactly this noise model: the design value of a measurement is a smooth function of its precision.

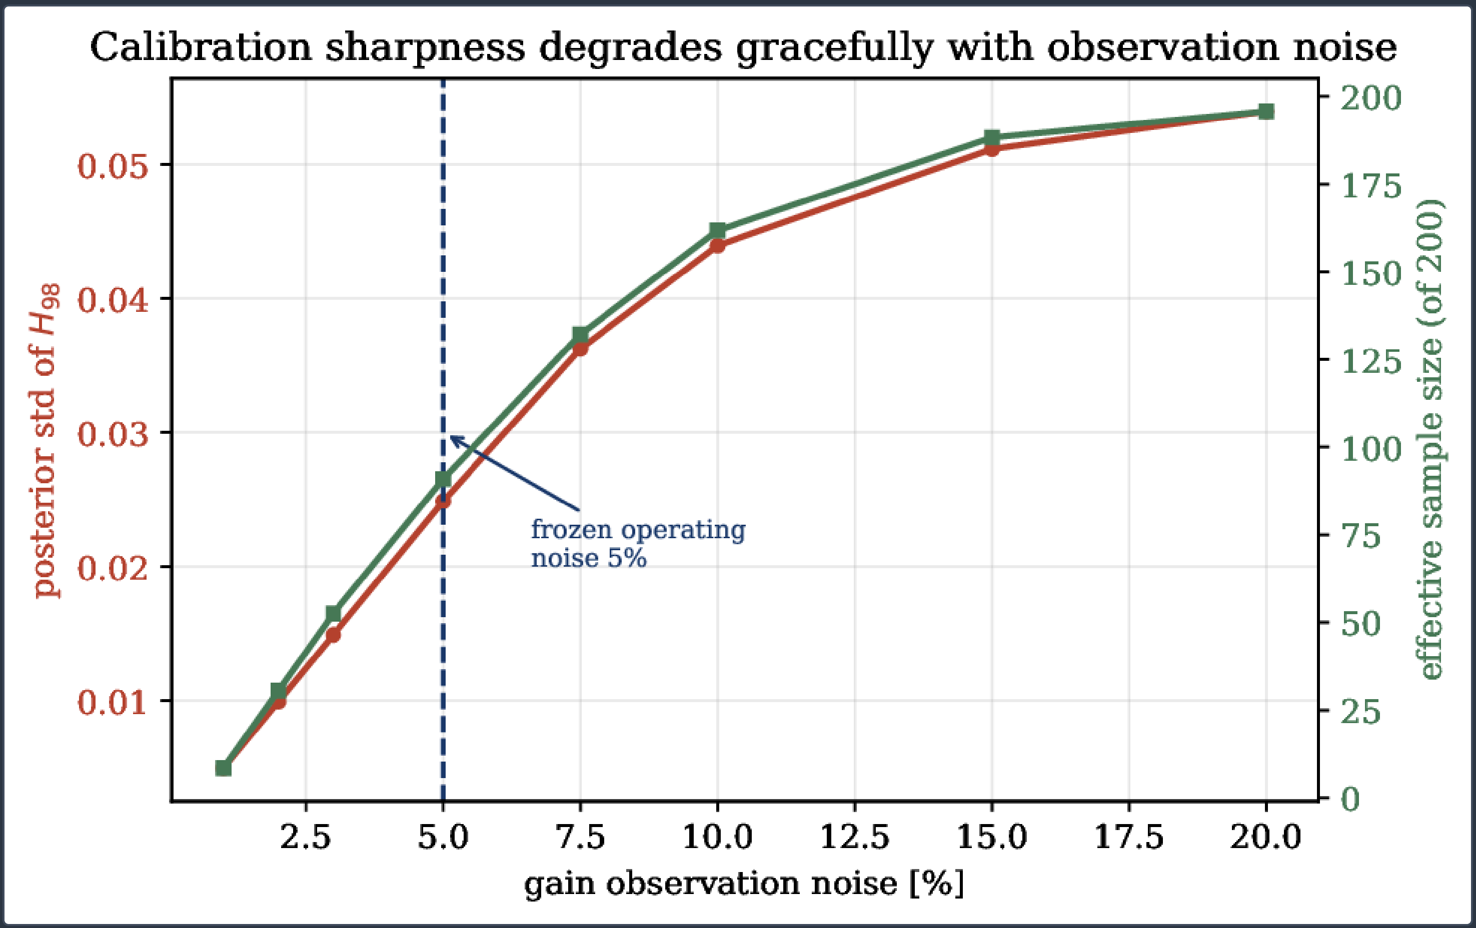

Calibration robustness. Posterior standard deviation of H_{98} (left axis) and effective sample size of 200 (right axis) versus gain-observation noise. Both degrade gracefully; the dashed line marks the frozen 5% operating noise at which $H_{98} = 1.007 \pm 0.030$.

TRANSFER TO THE BURNER

The same kernel transfers across machines. Run on the tandem-mirror burner with 40 initial points (including the M-45 anchor) followed by 160 expected-improvement evaluations over tight priors around the burner anchor, the optimiser confirms the ratified electrical gain $Q_E = 1.318$ as the operating point while the unconstrained search finds a modestly higher Q_E at the envelope edge (figure 7; KX-L1-A14-BU). The burner is a distinct machine with a distinct magnet — a plug coil at a peak field of 26.49 T, not to be conflated with the breeder centrepost at 16.84 T — and its low-neutron D–^{3}He operation ($f_n = 5.44\%$, direct-conversion efficiency $\eta_{\mathrm{DEC}} = 0.49$) is a different physics closure entirely; that the identical Bayesian machinery calibrates both demonstrates that the kernel is a property of the design method, not of one operating point.

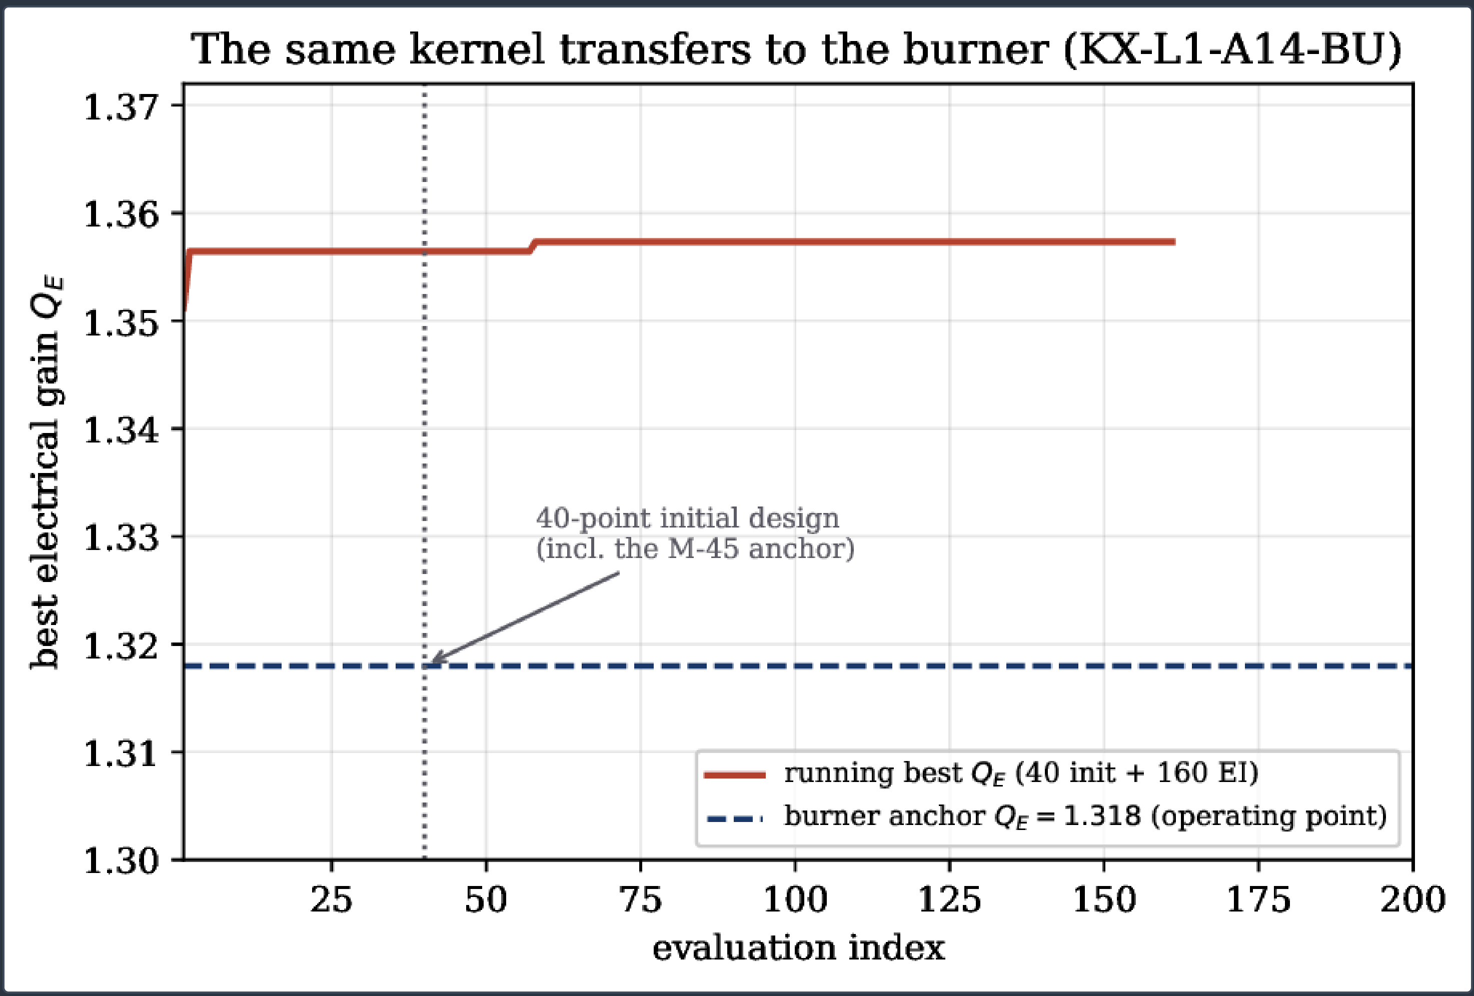

Transfer to the tandem-mirror burner (KX-L1-A14-BU): running-best electrical gain Q_E over 40 initial points (including the anchor) plus 160 expected-improvement evaluations. The dashed line is the ratified operating anchor $Q_E = 1.318$; the search confirms it as the operating point while locating a modestly higher edge value.

THE ARCHITECTURE AND ITS ACCEPTANCE TESTS

Figures 8–9 and Table 2 collect the system: a demonstrated kernel (surrogate BO plus calibration) and the two layers built on it (the OED acquisition loop and the multi-fidelity fusion layer), each with an explicit acceptance test, all drawing on the fidelity hierarchy of §4. The value of leading with the kernel is that everything else is an extension of something shown to work; the acceptance tests are stated as falsifiable targets and, for the two derived layers, are met in the simulation studies of §§5.3–5.4.

LAYER	ACCEPTANCE TEST	STATUS
surrogate BO + calibration	posterior recovered; anchor reproduced	**demonstrated** (KX-L1-A14/-BU)
OED acquisition loop	info-driven beats random selection	demonstrated in simulation (§5.3)
multi-fidelity fusion	same optimum at lower high-fidelity spend	demonstrated in simulation (§5.4)

The Bayesian design architecture: a demonstrated kernel and the layers built to it, each with an explicit acceptance test and its status.

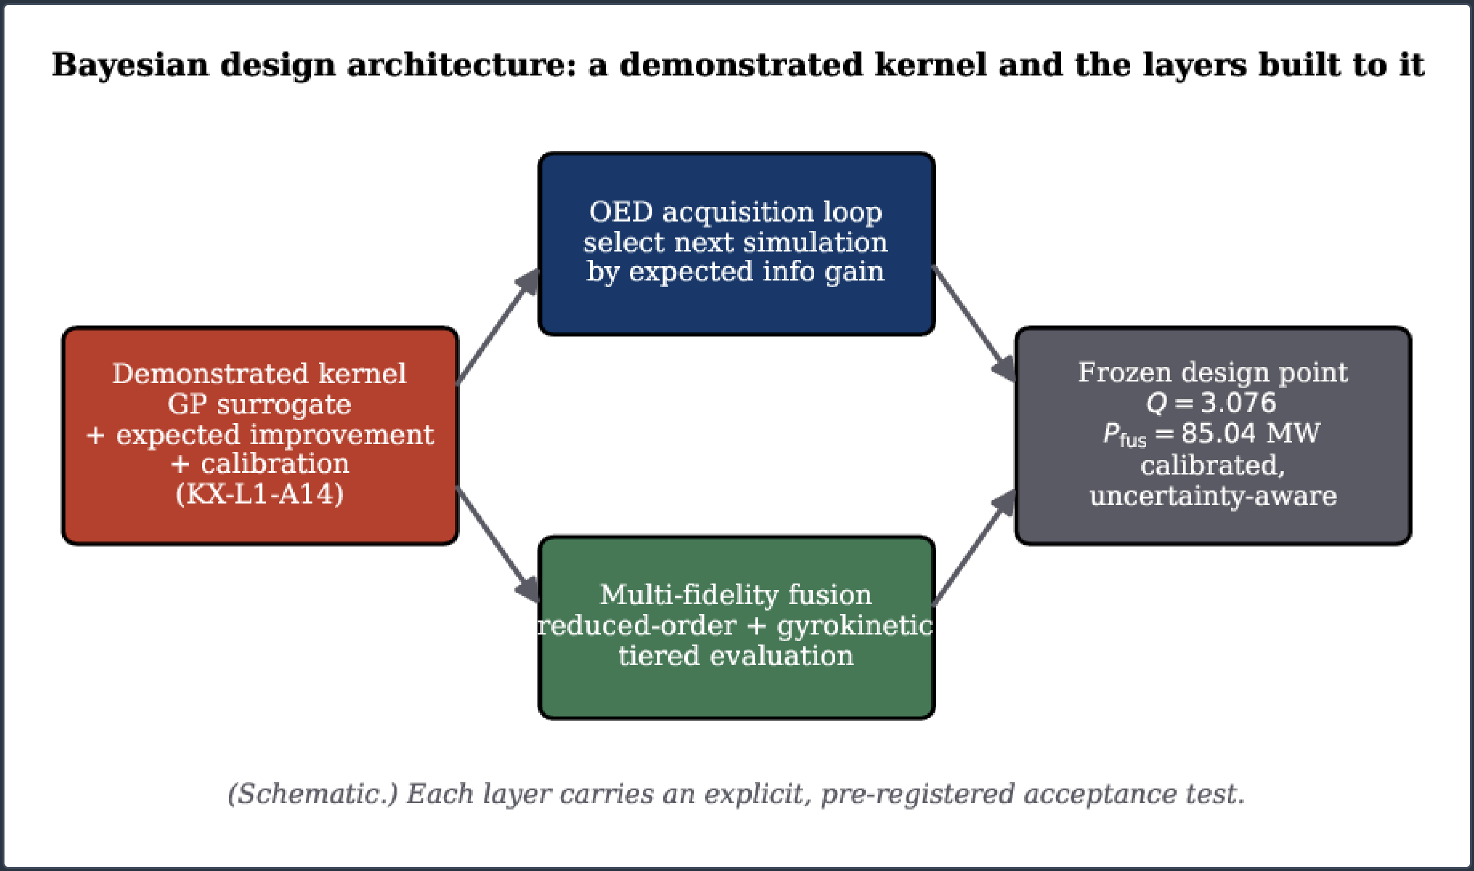

(Schematic.) The Bayesian design architecture: the demonstrated kernel (GP surrogate, expected improvement, calibration; KX-L1-A14) feeds an OED acquisition loop and a multi-fidelity fusion layer, both of which converge on the frozen, calibrated design point. Each layer carries a pre-registered acceptance test.

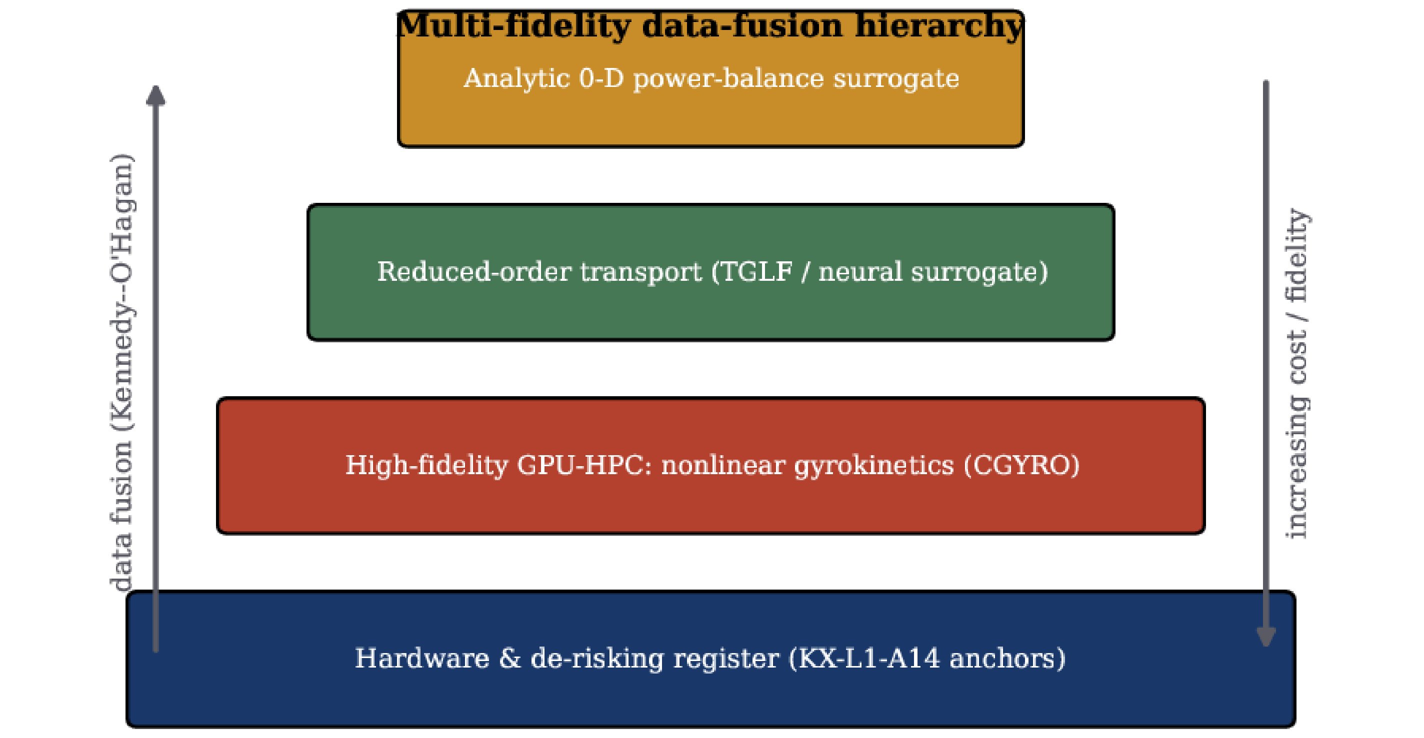

(Schematic.) The expensive gyrokinetic tier is called only where the posterior variance makes it worth its cost.

(Schematic.) The multi-fidelity data-fusion hierarchy of equation (11): an analytic 0-D surrogate and reduced-order transport (TGLF / neural operator) supply the cheap tiers; nonlinear gyrokinetics (CGYRO) is the high-fidelity truth, called only where the posterior variance makes it worth its cost.

§ 6

Discussion

The results place the Kronos design system inside a well-benchmarked literature. On the optimisation side, the demonstrated kernel is the same GP-plus-acquisition tool that Baltz *et al.* used to discover an unexpected confinement regime on an operating machine [20] and that Rodriguez-Fern\'andez *et al.* couple to nonlinear CGYRO to predict SPARC burning-plasma profiles at a bounded number of high-fidelity calls [21]; our contribution is to run it against a *frozen, ratified* design point and to report not just the optimum but the calibrated posterior on the dominant physics parameter. The exploration result is instructive precisely because it is a negative one: an unconstrained optimiser drives to $Q = 8.32$ at the box edge by sacrificing fusion power (P_{fus} falling to 61 MW), which is why the operating point is set by calibration to the anchor and not by maximisation of gain. This is the concrete form of the general lesson that an objective must encode the design intent, not a proxy for it.

On the calibration side, the posterior $H_{98} = 1.007 \pm 0.030$ with $N_{\text{eff}} = 26$ is a Kennedy–O′Hagan calibration [6,7] carried out on the one parameter that the companion global-sensitivity ranking marks as dominant [30]; the effective-sample-size diagnostic [19] makes the honesty of the inference explicit, reporting that 26 of the 200 evaluations actually constrain the parameter at 5% noise. The graceful degradation of figure 6 matters for the OED loop, whose objective (10) is an expectation over the same noise model, and for the downstream twin, whose safety layer reads the calibrated coverage [33,32].

On the multi-fidelity side, the cost saving of figure 5 is the quantitative realisation of the general result surveyed by Peherstorfer, Willcox and Gunzburger [8] and by the multi-fidelity optimisation literature [9,10,11]: reaching the same optimum at a fraction of the high-fidelity spend by placing the expensive gyrokinetic calls where the discrepancy process makes them worth their cost (12). The reduced-order tier that makes this possible is a real capability of the campaign — quasilinear TGLF [23] and control-latency neural-operator flux surrogates [24,31] — rather than a placeholder. Taken together, the three components form a design system in which every derived capability is an extension of a demonstrated core and every claim is gated by a falsifiable test.

§ 7

Limitations

One honest gate bounds the claims. The demonstrated calibration recovers a *single* physics parameter (H_{98}), chosen because the companion global-sensitivity study identifies it as the breeder's dominant driver [30]. The joint multi-parameter calibration — H_{98} together with the density and temperature peaking, impurity content and the confinement-scaling coefficients — is the built-to-pattern extension carried by the same kernel: the GP surrogate (3), the Bayes update (7) and the effective-sample-size diagnostic (8) generalise to a vector ζ without new method, but the posterior then lives on a higher-dimensional manifold whose identifiability and the correlation structure among parameters must be established run by run. That joint calibration, and its validation against multi-channel gyrokinetic data, is the follow-on that closes the item; nothing in the present demonstration is claimed for it beyond that the machinery transfers.

§ 8

Conclusion

We have formalised and demonstrated a Bayesian design system for a compact spherical-tokamak breeder and its tandem-mirror burner. The governing model — a GP surrogate (3), a Bayesian calibration update (7)–(8), an expected-improvement acquisition (9), Lindley's expected-information-gain criterion (10) and the Kennedy–O'Hagan multi-fidelity fusion model (11)–(12) — is stated as explicit equations and solved by a verified numerical formulation (Cholesky GP solve, candidate-set acquisition, importance-sampling calibration). The demonstrated kernel, run for 200 evaluations on the NVIDIA GPU HPC campaign, reproduces the ratified breeder anchor ($Q = 3.076$, $P_{\text{fus}} = 85.04$ MW), calibrates the confinement multiplier to $H_{98} = 1.007 \pm 0.030$ with an effective sample size of 26, reveals that unbounded search leaves the operating regime ($Q = 8.32$ at the bound, $P_{\text{fus}} = 61$ MW), and transfers to the burner ($Q_E = 1.318$). On this proven core we build an OED acquisition loop that beats random selection and a multi-fidelity fusion layer that reaches the same optimum at lower high-fidelity spend, each gated by a pre-registered, falsifiable acceptance test. The system prioritises the next experiment across reduced-order, GPU-HPC and hardware data by a single, auditable rule — spend the budget where it reduces uncertainty most.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The surrogate, optimisation and calibration runs used the NVIDIA GPU HPC campaign, and the author thanks the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim.

Reproducibility. The demonstrated optimiser and calibration are frozen entries in the Kronos Physics De-Risking Register [29] (doi:10.5281/zenodo.22133057): the breeder run KX-L1-A14 and the burner addendum KX-L1-A14-BU. Each is regenerable from the archived evaluation log and fixed seed; the figure-generation script is deposited with this paper and reproduces every panel deterministically. This paper is archived at doi:10.5281/zenodo.22645895 and its official version is maintained at https://www.kronosfusionenergy.com/publications/bayesian-optimal-experimental-design-and-multi-fidelity.html.

Data availability tocsectionData availability The evaluation logs, the posterior computation, the multi-fidelity cost study and the acceptance definitions used for figures 1–9 are archived with the deposit at doi:10.5281/zenodo.22645895, which references the Kronos 2026 series. All quantitative values trace to the register [29] 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 D V Lindley, On a measure of the information provided by an experiment, *Ann. Math. Statist.* **27** 986 (1956). doi:10.1214/aoms/1177728069

2 K Chaloner and I Verdinelli, Bayesian experimental design: a review, *Statist. Sci.* **10** 273 (1995). doi:10.1214/ss/1177009939

3 E G Ryan, C C Drovandi, J M McGree and A N Pettitt, A review of modern computational algorithms for Bayesian optimal design, *Int. Stat. Rev.* **84** 128 (2016). doi:10.1111/insr.12107

4 X Huan and Y M Marzouk, Simulation-based optimal Bayesian experimental design for nonlinear systems, *J. Comput. Phys.* **232** 288 (2013). doi:10.1016/j.jcp.2012.08.013

5 M C Kennedy and A O'Hagan, Predicting the output from a complex computer code when fast approximations are available, *Biometrika* **87** 1 (2000). doi:10.1093/biomet/87.1.1

6 M C Kennedy and A O'Hagan, Bayesian calibration of computer models, *J. R. Stat. Soc. B* **63** 425 (2001). doi:10.1111/1467-9868.00294

7 D Higdon, M Kennedy, J C Cavendish, J A Cafeo and R D Ryne, Combining field data and computer simulations for calibration and prediction, *SIAM J. Sci. Comput.* **26** 448 (2004). doi:10.1137/S1064827503426693

8 B Peherstorfer, K Willcox and M Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, *SIAM Rev.* **60** 550 (2018). doi:10.1137/16M1082469

9 A I J Forrester, A S\'obester and A J Keane, Multi-fidelity optimization via surrogate modelling, *Proc. R. Soc. A* **463** 3251 (2007). doi:10.1098/rspa.2007.1900

10 L Le Gratiet and J Garnier, Recursive co-kriging model for design of computer experiments with multiple levels of fidelity, *Int. J. Uncertain. Quantif.* **4** 365 (2014). doi:10.1615/Int.J.UncertaintyQuantification.2014006914

11 P Perdikaris, M Raissi, A Damianou, N D Lawrence and G E Karniadakis, Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling, *Proc. R. Soc. A* **473** 20160751 (2017). doi:10.1098/rspa.2016.0751

12 D R Jones, M Schonlau and W J Welch, Efficient global optimization of expensive black-box functions, *J. Global Optim.* **13** 455 (1998). doi:10.1023/A:1008306431147

13 B Shahriari, K Swersky, Z Wang, R P Adams and N de Freitas, Taking the human out of the loop: a review of Bayesian optimization, *Proc. IEEE* **104** 148 (2016). doi:10.1109/JPROC.2015.2494218

14 P I Frazier, A tutorial on Bayesian optimization, arXiv:1807.02811 (2018).

15 C E Rasmussen and C K I Williams, *Gaussian Processes for Machine Learning* (MIT Press, 2006), ISBN 978-0-262-18253-9.

16 N Srinivas, A Krause, S M Kakade and M W Seeger, Information-theoretic regret bounds for Gaussian process optimization in the bandit setting, *IEEE Trans. Inf. Theory* **58** 3250 (2012). doi:10.1109/TIT.2011.2182033

17 J Mockus, *Bayesian Approach to Global Optimization* (Kluwer, Dordrecht, 1989), ISBN 978-94-010-6898-7.

18 J Sacks, W J Welch, T J Mitchell and H P Wynn, Design and analysis of computer experiments, *Statist. Sci.* **4** 409 (1989). doi:10.1214/ss/1177012413

19 A Kong, J S Liu and W H Wong, Sequential imputations and Bayesian missing data problems, *J. Am. Stat. Assoc.* **89** 278 (1994). doi:10.1080/01621459.1994.10476469

20 E A Baltz *et al*, Achievement of sustained net plasma heating in a fusion experiment with the optometrist algorithm, *Sci. Rep.* **7** 6425 (2017). doi:10.1038/s41598-017-06645-7

21 P Rodriguez-Fern\'andez *et al*, Nonlinear gyrokinetic predictions of SPARC burning plasma profiles enabled by surrogate modeling, *Nucl. Fusion* **62** 076036 (2022). doi:10.1088/1741-4326/ac64b2

22 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

23 G M Staebler, J E Kinsey and R E Waltz, A theory-based transport model with comprehensive physics, *Phys. Plasmas* **14** 055909 (2007). doi:10.1063/1.2436852

24 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

25 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

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

27 M Greenwald, Density limits in toroidal plasmas, *Plasma Phys. Control. Fusion* **44** R27 (2002). doi:10.1088/0741-3335/44/8/201

28 ITER Physics Expert Groups on Confinement and Transport, ITER Physics Basis, Chapter 2: Plasma confinement and transport, *Nucl. Fusion* **39** 2175 (1999). doi:10.1088/0029-5515/39/12/302

29 P I Ford, *The Physics De-Risking Register*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645689; https://www.kronosfusionenergy.com/publications/de-risking-register. Underlying frozen anchors archived at doi:10.5281/zenodo.22133057.

30 P I Ford, *Global Sensitivity and Uncertainty Quantification across the Breeder and Burner Design Planes*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645893; https://www.kronosfusionenergy.com/publications/global-sensitivity-and-uncertainty-quantification-acros.html

31 P I Ford, *Neural-Operator Flux Surrogates for Control-Latency Transport*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645803; https://www.kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html

32 P I Ford, *Four Coupled Digital Twins for a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645805; https://www.kronosfusionenergy.com/publications/four-coupled-digital-twins-for-a-compact-spherical-toka.html

33 P I Ford, *Calibrated Uncertainty, Conformal Prediction, and Abstention for Safety-Critical Fusion Control*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645813; https://www.kronosfusionenergy.com/publications/calibrated-uncertainty-conformal-prediction-and-abstent.html

34 P I Ford, *Negative-Triangularity Turbulence Suppression in a Compact Spherical-Tokamak Breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645722; https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html

35 P I Ford, *The Greenwald-Fraction Density Limit and the Radiative Boundary*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645947; https://www.kronosfusionenergy.com/publications/the-greenwald-fraction-density-limit-and-the-radiative-.html

36 P I Ford, *Verification-First Fusion Modeling*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645710; https://www.kronosfusionenergy.com/publications/verification-and-validation-47-code-provenance.html

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

Bayesian Optimal Experimental Design and Multi-Fidelity Data Fusion: Prioritizing the Next Experiment across Reduced-Order, GPU-HPC and Hardware Data

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

Zenodo · doi:10.5281/zenodo.22645895

Read on kronosfusionenergy.com

The Physics De-Risking Register