plasmafast

Fast peeling-ballooning growth-rate inference from a magnetic equilibrium.

A product of dataquanta LLC · info@dataquanta.io

What it is

plasmafast is a neural-operator surrogate that takes a magnetic equilibrium (G-EQDSK file) plus two transport parameters — resistivity η and perpendicular viscosity ν — and returns the linear peeling-ballooning growth-rate field γ(n, ρ_pol) across toroidal mode numbers and radial position.

End-to-end inference is **71 ms on CPU** (36 ms equilibrium parse, 35 ms model forward — measured on an M3 Pro at batch size 1). This is roughly two to three orders of magnitude faster than the ideal-MHD solvers used for peeling-ballooning stability — ELITE, the inner stability step in EPED-class workflows, and the 3-D code PB3D — which is the latency budget that makes plasmafast suitable for real-time tokamak control loops rather than offline scenario analysis.

Where plasmafast sits among existing tools

TOOL	OUTPUT	INFERENCE TIME	UNCERTAINTY / OUT-OF-DISTRIBUTION DETECTION
ELITE (Wilson et al., PoP 9, 1277, 2002)	γ scalar	~tens of s	—
EPED (Snyder et al., NF 51, 103016, 2011)	pedestal height / width	~min	—
PB3D (Weyens et al., JCP, 2017)	γ scalar	~min	—
KARHU (Bruncrona et al., PoP 32, 092501, 2025)	γ_{max} scalar	ms-class	—
EuroPED-NN (Panera Alvarez et al., PPCF 66, 095012, 2024)	pedestal scalars (not γ)	ms-class	Bayesian NN
Kim, Kolemen et al. (Nat. Commun. 15, 2024)	RMP control scores (not γ)	ms-class	not stated
plasmafast (this work)	**γ(n, ρ_pol) field**	**71 ms**	split-conformal CI · k-NN OOD

ELITE and PB3D compute γ from first-principles ideal MHD, and the EPED workflow wraps ELITE for pedestal prediction — all at seconds-to-minutes timescales built for design-cycle iteration. **KARHU** (Bruncrona et al. 2025) is the closest prior ML surrogate: it is fast (sub-second — its paper reports under 1 s for a nine-equilibrium scan), but it regresses a single scalar, the maximum growth rate γ_max, from Europed-format profiles, and ships no out-of-distribution detection. EuroPED-NN predicts pedestal scalars rather than γ. plasmafast's position is the combination that none of these hold at once: a full γ(n, ρ_pol) field, computed directly from a raw G-EQDSK equilibrium, in 71 ms, with a calibrated OOD detector.

Current validation

TEST	THRESHOLD	RESULT	STATUS
γ on cbm18 reference equilibrium, vs PB3D γ = 0.245	within 30 %	0.2115 (13.7 %)	passes
Closed-loop rollout γ-ratio (model used recursively)	≤ 1.10	0.72	passes
SPARC classifier on 10-shot benchmark	≥ 9 / 15	14 / 15	passes
Out-of-distribution detector flags real-machine input	flag every case	6 of 6 cases flagged	passes
Cross-machine sign correctness (5 real-machine cases)	4 / 5	5 / 5	passes
Cross-machine magnitude calibration (γ ratio)	2× window	2× ≤ ratio ≤ 14×	partial

The closed-loop rollout gate moved from 1.261 (fail) to 0.72 (pass) on 2026-05-22 by switching from a state-MSE proxy loss to a differentiable port of the bench's γ-extractor (FFT slope on the rollout's ψ trajectory) plus pushforward sampling (Brandstetter et al. 2022, ICLR). The naive proxy made the gate *worse* (1.47) before the principled fix closed it. Ratio 0.72 means the model now slightly under-predicts γ in rollout — conservative for control triggers, still one-sided; tightening to a symmetric 0.9–1.1 window is stated Phase I scope.

Five real-machine equilibria are now in inference, from the SCOREC `Fusion_Public` archive: DIII-D 145419, JET 79692, NSTX 132588, ITER 15MA, KSTAR 018451. All five are correctly flagged out-of-distribution by the k-NN detector (scores in the narrow range 2,900–3,100 vs threshold 1.0), and the model predicts a growing (positive γ) ideal-MHD mode on all five — sign correctness 5 / 5. Magnitudes are not yet calibrated outside the synthetic training distribution.

Two of the five have BOUT++ ground truth available: DIII-D 145419 (γ_pred = +1.29 vs γ_BOUT = +0.094, n = 4 dominant toroidal mode, $R^2 = 0.92$ fit) and JET 79692 (γ_pred = +0.83 vs γ_BOUT = +1.81, $R^2 = 0.20$ fit so the ground truth itself is noisy). Magnitudes are off in opposite directions on these two: 14× over on DIII-D, 2× under on JET. In the OOD regime γ predictions are flagged advisory.

Discrimination scaling result. A real-equilibrium auxiliary supervised loss was added to the training and exercised in two configurations. In the first, supervision was given on the same DIII-D 145419 equilibrium at two different transport-parameter settings ($\eta = 1 \times 10^{-4}$ giving γ_BOUT = +0.094, $\eta = 1 \times 10^{-5}$ giving γ_BOUT = +0.622 — targets 6.6× apart). The model's predictions on the two cases stayed identical to within ±0.003: same-equilibrium (η, ν) variation is not enough to drive discrimination through the current γ-head. In the second configuration, supervision was given on three *different* equilibria (DIII-D 145419, SPARC 1514, SPARC 1519, all $\eta = 1 \times 10^{-4}$, γ_BOUT spanning 8×). The model now converges each supervised case to within ~10 % of its own γ_BOUT (DIII-D +0.10 vs +0.094; SPARC 1514 +0.66 vs +0.637; SPARC 1519 +0.68 vs +0.753) — a 170× larger inter-case spread than the (η, ν) result. The architecture *does* discriminate between equilibria when given diverse-equilibrium supervision; the binding constraint for cross-machine magnitude calibration is the size and diversity of the supervised equilibrium set, not γ-head capacity. Held-out OOD cases not in the supervised set still receive a deflated γ from this N = 3 fine-tune (the model has not learned the manifold beyond the supervised support), which is the expected outcome at this sample count. Phase I scope is therefore (a) acquire BOUT++ ground truth on a diverse set of equilibria (target N ≥ 8 spanning DIII-D, JET, SPARC, NSTX, MAST-U, KSTAR, ITER), and (b) train against that pooled set. All three acceptance gates remain passing throughout (Gate 1 = 0.72, Gate 2 = 14 / 15, Gate 3 = 10.9 % off PB3D — Gate 3 *improved* over the v2 canonical's 13.7 %).

Limitations stated up front

- The training distribution is **synthetic** — BOUT++ cbm18-family equilibria, not real-machine shots. Cross-machine generalization is the principal open empirical question and the reason cross-machine validation is the top engineering priority.
- The OOD detector is **binary** against G-EQDSK input by construction. The training pooled-embeddings come from BOUT++ internal meshes, so any G-EQDSK query is flagged as out-of-distribution. It cleanly separates "synthetic training" from "any real equilibrium" but does not finely rank between machines.
- The split-conformal γ confidence intervals are reported with `valid_when_ood = false` on every G-EQDSK query by design, consistent with the OOD detector behaviour above.
- Closing the rollout gate cost some cbm18a γ accuracy (4.7 % → 13.7 %, still well within the ±30 % window); SPARC classifier is unchanged at 14 / 15. Recovering the cbm18a margin and tightening the rollout ratio to a symmetric 0.9–1.1 window are Phase I scope items.

What we are looking for

A 15-minute call to scope any of:

1. **An INFUSE co-PI proposal** on integrating plasmafast with a national-lab real-time control architecture.
2. **DOE SBIR FY27 Phase I** in the FES Autonomous Plasma Control Systems and AI/ML topic areas (topics C59-21a and C59-22c in the FY25 solicitation).
3. **A DOE Milestone capability-enhancement subaward** with a pilot-plant prime, under the June 2025 program expansion that names AI/ML for fusion and plasma science as an eligible enhancement topic.

Full acceptance bench, OOD audit, and speed-benchmark reports available on request.