

Machine-Agnostic Disruption Proximity Law from Plasma Current Dynamics: An inZOR-ND Discovery

Empirical Law Candidate for Tokamak Disruption Proximity Based on I_p Features Only

Dumitru Novic · March 2026 · Machines: MAST · C-Mod (Alcator) · HL-2A · 1,384 disruptive windows · 4,260 non-disruptive windows

ABSTRACT

We propose a compact empirical indicator η for disruption proximity in tokamak plasmas, derived **solely from plasma current (I_p) dynamics** using the inZOR-ND bio-adaptive genomic discovery engine. The 3-term candidate law $\eta \approx -\text{kurtosis} + 0.80 \cdot \text{rate_ratio} + 0.65 \cdot \text{cv_ratio}$ encodes a universal structural nucleus (kurtosis, linked to MHD precursor dynamics) and two instability components (dynamic acceleration and variability amplification). Validated cross-machine on MAST (674 D / 3647 N), C-Mod (414 D / 13 N), and HL-2A (296 D / 600 N) without machine-specific retraining, the formula achieves $\text{min}(\text{sep}) = 1.071$, global ROC-AUC = 0.982, full Leave-One-Machine-Out (LOMO) PASS across all three folds, bootstrap $\text{CI}_{95} \text{min}(\text{sep}) > 0$ over 300 replicas, and coefficient sensitivity range $\text{min}(\text{sep}) \in [0.858, 1.254]$ under $\pm 20\%$ perturbation. The temporal trajectory $\eta(t)$ rises systematically toward disruption on all three machines, demonstrating a clock-like proximity behaviour. This work does not claim a universal physical law, but presents a strong empirical candidate for cross-machine I_p -based disruption proximity with transparent physical grounding.

MAIN EMPIRICAL LAW CANDIDATE (3-TERM, DUAL-INSTABILITY)

$$\eta \approx -\text{kurtosis} + 0.80 \cdot \text{rate_ratio} + 0.65 \cdot \text{cv_ratio}$$

η increases as plasma approaches disruption · Derived from I_p only · No machine-specific tuning

1.071

min(sep) cross-machine

0.982

Global ROC-AUC

3/3

LOMO folds passed

0.945

Global PR-AUC

>0

Bootstrap $\text{CI}_{95} \text{min}(\text{sep})$

3

Tokamaks validated

1. Background and Motivation

Plasma disruptions in tokamaks cause sudden termination of fusion experiments, potentially damaging the machine through induced currents, thermal loads, and halo currents. Early warning systems for disruption proximity are therefore of critical importance, particularly for next-generation devices (ITER, DEMO).

Most disruption prediction approaches rely on multi-diagnostic data (magnetic equilibrium, temperature profiles, radiation) and machine-learning classifiers trained per machine. A long-standing open question is whether any compact indicator derived from a single diagnostic (plasma current, I_p) can provide cross-machine disruption proximity information without per-machine retraining.

This work addresses that question using the **inZOR-ND** bio-adaptive discovery engine, which explores the space of candidate empirical formulas guided by cross-machine separation fitness. The result is a 3-term candidate law with transparent physical grounding.

2. Feature Definitions

All features are computed from the plasma current waveform $I_p(t)$ using two windows:

- **D** (late/disruption window): 450 samples ending at $t_{\text{disrupt}} - 50$ ms.
- **N** (reference window): full shot.

Feature	Definition	Physical role
kurtosis	Excess kurtosis of I_p in window D	MHD precursor structural change — kurtosis <i>decreases</i> before disruption as the current profile flattens
rate_ratio	$\max dI_p/dt _D / \max dI_p/dt _N$	Dynamic acceleration — late-phase surge in current rate of change
cv_ratio	$(\text{std}/\text{mean})_D / (\text{std}/\text{mean})_N$	Variability amplification — relative increase in I_p fluctuations before disruption

3. Discovery Path with inZOR-ND

The formula was not manually designed. It emerged through a systematic bio-adaptive law discovery process in five stages:

- **Stage 1 — 2D law-space:** inZOR-ND converges to $-\text{kurtosis} + \text{rate_ratio}$ as attractor. $\min(\text{sep}) = 0.68$ on C-Mod/HL-2A/MAST.
- **Stage 2 — 3D/4D extension:** A second valley emerges: $-\text{kurtosis} + \text{cv_ratio} + \text{skew}$, hinting at two machine regimes.
- **Stage 3 — C-Mod structural analysis:** On C-Mod, cv_ratio dominates over rate_ratio . The 2-term formula is the C-Mod bottleneck.

- **Stage 4 – Dual-instability 2D search:** Grid search over (a, b) in $\eta = -\text{kurtosis} + a \cdot \text{rate_ratio} + b \cdot \text{cv_ratio}$. Optimal: **a = 0.80, b = 0.65**.
- **Stage 5 – Validation:** LOMO, temporal $\eta(t)$, per-machine structural analysis, bootstrap, sensitivity. All pass.

Discovery path

Fig 1 – Discovery path: $\min(\text{sep})$ staircase from single features to 3-term law. $\min(\text{sep})$ increases monotonically: 1-term < 2-term < 3-term. The 3-term law achieves $\min(\text{sep}) = 1.071$, discovered by inZOR-ND without manual formula design.

4. Baseline / Ablation Test

To confirm the 3-term formula provides genuine value over simpler variants, we evaluate all 6 formulas on the same data, convention and machines (C-Mod, HL-2A, MAST).

Ablation

Fig 2 – Baseline / Ablation: $\min(\text{sep})$ and Global ROC-AUC across all 6 formulas. The 3-term law achieves the highest $\min(\text{sep}) = 1.071$ and balanced global ROC-AUC = 0.982.

Formula	Terms	$\min(\text{sep})$	Global ROC	Global PR
-kurtosis	1	0.581	0.965	0.897
rate_ratio	1	0.178	0.990	0.957
cv_ratio	1	0.505	0.816	0.628
-kurtosis + rate_ratio	2	0.680	0.986	0.961
-kurtosis + cv_ratio	2	0.949	0.969	0.896
-kurtosis + 0.80·rate_ratio + 0.65·cv_ratio	3	1.071	0.982	0.945

Conclusion: 1-term < 2-term < 3-term in $\min(\text{sep})$. The 3-term law resolves the C-Mod bottleneck (C-Mod sep rises from 0.68 with 2-term to 1.071 with 3-term) by combining rate_ratio (MAST-dominant) and cv_ratio (C-Mod-dominant) in a single balanced formula.

5. Leave-One-Machine-Out (LOMO) Cross-Validation

The formula coefficients are fixed (no retraining). For each fold, the formula is applied verbatim on the unseen machine. Pass criterion: $sep > 0$ and $ROC-AUC > 0.5$ on the test machine.

LOMO

Fig 3 – LOMO Validation: sep and ROC-AUC for each test (left-out) machine. All 3 folds PASS, confirming the formula generalises across tokamaks without retraining.

Train machines	Test machine	sep	ROC-AUC	n_D	n_N	Result
C-Mod + HL-2A	MAST	11.786	0.980	674	3647	PASS
MAST + HL-2A	C-Mod	1.071	0.757	414	13	PASS
MAST + C-Mod	HL-2A	9.344	0.988	296	600	PASS

Overall: PASS – cross-machine validity confirmed. C-Mod (highly imbalanced: 414 D / 13 N) is the hardest fold, yet $sep = 1.071 > 0$ and $ROC = 0.757 > 0.5$. The formula generalises.

6. Temporal Trajectory $\eta(t)$ – Disruption Clock

η is computed at sliding windows $-100, -80, -60, -40, -20$ ms before disruption. A systematic rise toward the event confirms the clock-like proximity behaviour of the indicator.

Temporal eta

Fig 4 – Temporal trajectory of $\eta(t)$ for C-Mod, HL-2A and MAST. Mean η (solid line) \pm std (shaded band). η increases monotonically in all three machines, confirming disruption clock behaviour. This property is essential for real-time proximity monitoring.

7. Per-Machine Structural Analysis

Pearson |correlation| between each feature and η reveals the dominant instability regime per machine, explaining why both $rate_ratio$ and cv_ratio are needed.

Structural

Fig 5 – Per-machine structural analysis. kurtosis is a universal nucleus. $rate_ratio$ dominates on MAST (dynamic instability); cv_ratio dominates on C-Mod and HL-2A (variability amplification). The 3-term formula captures both regimes simultaneously.

Machine	Dominant instability term	Physical interpretation
MAST	$rate_ratio$	Rapid I_p rate surges (dynamic instability, high-current spherical tokamak)

C-Mod	cv_ratio	Variability amplification (compact high-field machine, rf-heated)
HL-2A	cv_ratio (slightly dominant)	Similar variability regime to C-Mod

8. Robustness and Sensitivity Analysis

Robustness sensitivity

Fig 6 – Left: Bootstrap per-machine sep (mean \pm std, 300 replicas). Centre: Global ROC-AUC bootstrap $CI_{95} = [0.9789, 0.9845]$. Right: min(sep) distribution under $\pm 20\%$ coefficient perturbation (sensitivity). All results remain positive under perturbation.

A. Bootstrap Robustness (shot-level resampling, N = 300 replicas)

Metric	Mean	Std	CI 2.5%	CI 97.5%
min(sep)	0.9405	0.2338	0.3730	1.2108
Global ROC-AUC	0.9820	0.0014	0.9789	0.9845

PASS: $CI_{2.5\%}$ for min(sep) = 0.373 > 0 and for Global ROC = 0.979 > 0.5. Result is robust to shot-level resampling.

B. Sensitivity (coefficient perturbation $\pm 20\%$, N = 300 replicas)

Metric	Mean	Std	Min	Max
min(sep)	1.0443	0.0930	0.8583	1.2541
Global ROC-AUC	0.9818	0.0011	0.9794	0.9842

PASS: min(sep) stays in [0.858, 1.254] (always > 0) under $\pm 20\%$ perturbation. The formula is insensitive to moderate coefficient variation.

9. Physical Interpretation

The kurtosis term captures the pre-disruption flattening of the I_p distribution. In stable plasma, I_p fluctuations are small and the distribution is peaked (high kurtosis). As MHD instabilities develop, the current profile flattens and large-amplitude fluctuations appear – the distribution becomes sub-Gaussian (low kurtosis). The negative sign on kurtosis in η means the indicator rises as kurtosis falls, correctly flagging the approach to disruption.

The `rate_ratio` and `cv_ratio` terms capture the *dynamics* of this process: `rate_ratio` detects surges in $|dI_p/dt|$ in the late phase relative to the full shot, while `cv_ratio` detects amplification of relative I_p fluctuations. Together, they encode two distinct physical instability pathways: (i) rapid current collapse precursors (MAST-dominant), and (ii) turbulent variability amplification (C-Mod/HL-2A-dominant).

10. Role of inZOR-ND

The inZOR-ND bio-adaptive genomic discovery engine is the enabling technology for this result. The engine explores the law-space of candidate empirical formulas guided by a cross-machine separation fitness function, without access to domain physics.

- 2D/3D/4D law-space exploration was driven by genomic population evolution under cross-machine fitness.
- The dual-instability structure (rate_ratio vs cv_ratio, machine-specific dominance) was revealed by the search, not by manual feature engineering.
- The frozen-elite selection mechanism provided stable formula candidates for structural analysis and validation.

inZOR-ND was previously validated on power systems (PFA), TESS exoplanet prioritization, refraction emergence, and social dynamics. This application to fusion disruption proximity represents a new domain: **bio-adaptive discovery of empirical laws in plasma physics from time-series features only.**

11. Data and Reproducibility

Data used: MAST (Mega Amp Spherical Tokamak, CCFE, public), C-Mod (Alcator C-Mod, MIT / Zindi competition, public), HL-2A (Southwestern Institute of Physics, China, public research dataset). All data in HDF5 format with `ip` dataset and `IsDisrupt` attribute.

Feature convention: D = 450 samples ending at `t_disrupt - 50 ms`; N = full shot. Same across all tests.

```
python3 run_dual_instability_law_test.py      # formula + temporal  $\eta(t)$ 
python3 run_priority_tests.py                 # LOMO + per-machine structural
python3 run_robustness_sensitivity.py         # bootstrap + sensitivity
python3 run_baseline_ablation_test.py        # ablation comparison
python3 generate_publication_figures.py      # all publication figures
```

Scope and limitations: This is an empirical candidate law from 3 tokamaks. Real-time deployment, transfer to unseen machines (EAST, JET, ITER), and physics-based validation require additional diagnostics, engineering qualification, and collaboration with tokamak operators. The formula uses I_p only – no magnetic equilibrium, temperature, or density. This work is a scientific baseline and inZOR-ND capability demonstration, not a production disruption prediction system.

12. Citation

If you use this work, please cite:

```
Novic, D. (2026). Machine-Agnostic Disruption Proximity Law from Plasma Current Dynamics: An inZOR-ND Discovery. Zenodo. https://doi.org/10.5281/zenodo.\[TBD\]
```

