

Physics-Constrained Learning for Reduced-Order Thermoacoustic Stability Screening: Locating Instability Boundaries from Few Evaluations

Mário Pedro Louzeiro e Rodrigues — Independent Researcher, Lisbon, Portugal

Corresponding author: Mário Pedro Louzeiro e Rodrigues (MPLR69@gmail.com) **Preprint** — not yet peer-reviewed · August 2026 · License CC BY 4.0

Abstract

High-fidelity prediction of thermoacoustic instability in combustors — for example through reactive Large-Eddy Simulation (LES) — is accurate but expensive: each operating point can cost days to weeks of computation, so mapping a stability boundary across a design space is often prohibitive. This paper presents a physics-constrained learning approach, here designated **Constrained-by-Physics Learning (CPL)**, and demonstrates its use as a cheap, physically consistent *pre-screen* that locates instability boundaries from very few evaluations. CPL is formulated as a single projection onto the manifold of states admissible under a governing operator, and is shown to have two equivalent faces: a causal (light-cone) mask applied to an attention operator, and a physics-residual penalty added to a fitting objective. Using a reduced-order thermoacoustic model of a generic combustor (a one-dimensional acoustic network closed by a Crocco n - τ flame-transfer function), we show that CPL (i) recovers a quarter-wave acoustic mode from a handful of noisy samples where an unconstrained regression diverges, and (ii) locates the marginal-stability flame gain from four noisy evaluations with an accuracy (1.99 ± 0.12) far exceeding a physics-free cubic fit (2.04 ± 0.58), recovering the same boundary that the full reduced model obtains from a fifteen-point sweep. All models, code and data are generic and fully reproducible; no proprietary geometry or calibration is used. The method is offered as a tool to make expensive high-fidelity campaigns efficient, by identifying in advance which operating points warrant full resolution.

Keywords: thermoacoustic instability; reduced-order modelling; flame-transfer function; Rayleigh criterion; physics-informed machine learning; physics-constrained learning; large-eddy simulation; combustion instability screening.

1. Introduction

Thermoacoustic instability — the self-amplifying coupling between unsteady heat release and acoustic pressure — remains a central difficulty in the design of combustors for gas turbines, rockets and air-breathing engines [1,2]. When the unsteady heat release is in phase with the pressure oscillation, the Rayleigh integral is positive and acoustic energy grows; small changes in flame response, geometry or operating point can move a combustor across the stability boundary [2,3].

Predicting that boundary with confidence generally requires high-fidelity, time-resolved simulation of the reacting flow — most notably reactive LES, in which a measured or resolved flame-transfer function (FTF) replaces analytic closures [3,4]. Such simulations are accurate but very expensive: a single operating point can require days to weeks of computation on large clusters. Consequently, *mapping* a stability boundary — sweeping flame gain, geometry or Mach number — by brute-force high-fidelity simulation is frequently impractical.

This motivates a two-tier strategy that is standard in low-order thermoacoustics [2,5]: use a cheap reduced-order model to *screen* the design space and flag the operating points near the boundary, then spend the

expensive high-fidelity budget only where it matters. The contribution of this paper is a learning method that sharpens the cheap tier. We impose the correct physical structure on a surrogate so that it extracts the stability boundary from *very few* evaluations — each of which may, in a real campaign, be a costly simulation or experiment.

The method builds on physics-informed machine learning, in which governing equations are enforced as soft constraints in the training objective [6]. We show that this residual-penalty idea and a causal masking of an attention operator [7] are two faces of the same projection onto the physically admissible manifold, and we demonstrate the practical payoff on a reduced thermoacoustic model. Everything below is generic and reproducible; the method is independent of any particular engine, geometry or proprietary parameter.

2. Reduced-order thermoacoustic model

We use a deliberately generic, low-dimensional model of the flame–acoustic coupling in a cylindrical chamber of length L , following standard low-order thermoacoustics [2,5,8]. The chamber is treated as a one-dimensional acoustic element with a compact flame at axial fraction x_f . The speed of sound in the hot products is $c = \sqrt{\gamma RT}$.

Passive acoustic modes. For a closed-inlet / open-outlet element, the passive (quarter-wave family) eigenfrequencies are

$$f_m = (2m - 1) \frac{c}{4L}, \quad m = 1, 2, 3, \dots \quad (1)$$

Flame response. The compact flame is closed by a Crocco n - τ flame-transfer function [8],

$$F(\omega) = n e^{i\omega\tau}, \quad (2)$$

where n is the interaction index (gain) and τ the convective/chemical time delay. In a high-fidelity campaign, $F(\omega)$ would be *measured* from the resolved reactive field; here it is analytic, which is sufficient to exercise the method.

Growth rate and Rayleigh coupling. The modal growth rate follows an energy balance per cycle: the Rayleigh gain from the flame minus the acoustic losses at the boundaries,

$$\sigma \propto \underbrace{n(\gamma - 1) \cos(\omega\tau) \phi(x_f)^2}_{\text{Rayleigh gain}} - \underbrace{D}_{\text{boundary losses}}, \quad (3)$$

with $\phi(x_f) = \sin(\omega x_f L/c)$ the mode shape at the flame and D a damping coefficient from imperfect reflection. A mode is unstable when $\sigma > 0$. Crucially, near the boundary the growth rate is **linear in the flame gain** n — a classical weak-coupling result [2] — so

$$\sigma(n) \approx -D_{\text{eff}} + C_{\text{eff}} n, \quad n_{\text{crit}} = D_{\text{eff}}/C_{\text{eff}}. \quad (4)$$

This model (implemented in the accompanying `reduced_thermoacoustic_model.py`) reproduces, for a representative product-temperature envelope (2520–3060 K), three quarter-wave modes at approximately **485 / 509 / 534 Hz**, and — sweeping the flame gain n from 0.5 to 4.0 — a growth rate that crosses zero at $n_{\text{crit}} \approx 2.0$ (Table 1). It is generic: cylindrical geometry, order-unity FTF, no proprietary parameters. It proves the *method*, not any particular engine.

Envelope point	T (K)	Quarter-wave mode f_1 (Hz)
V0	2520	485
M3	2780	509
M6	3060	534

Table 1. Fundamental quarter-wave modes from the reduced model across a representative product-temperature envelope. The M3 gain sweep crosses marginal stability at $n_{\text{crit}} \approx 2.0$.

3. Constrained-by-Physics Learning (CPL)

Let $\mathcal{L}$ be the linear operator of the governing physics (here, the acoustic wave / Helmholtz operator; more generally the relevant conservation operator). A physically admissible field lies in $\ker(\mathcal{L})$. CPL is the projection of a data-driven estimate onto that kernel, and can be realised in two equivalent ways.

Face A — architecture (causal mask). In an attention operator [7], the interaction weights are

$$A = \text{softmax} \left(\frac{QK^\top}{\sqrt{d_k}} + \log M \right), \quad (5)$$

where $M_{ij} \in \{0, 1\}$ is a mask that encodes an admissibility relation between elements i and j . Taking M to be a *light-cone* (causal) mask — element i may attend to j only if j lies in the past light-cone of i , $t_i - t_j \geq 0$ and $|x_i - x_j| \leq c(t_i - t_j)$ — sets the weights of all acausal pairs to zero. Physically inadmissible couplings are removed by construction.

Face B — objective (residual penalty). Alternatively, a surrogate u_θ is fitted by minimising a data term plus a penalised physics residual [6],

$$\theta^* = \arg \min_{\theta} \underbrace{\|u_\theta - y_{\text{data}}\|^2}_{\text{data}} + \lambda \underbrace{\|\mathcal{L}[u_\theta]\|^2}_{\text{physics violation}}. \quad (6)$$

As λ grows, the solution is projected onto $\ker(\mathcal{L})$: among all fields consistent with the sparse data, CPL selects the one that best satisfies the physics.

Both faces are the same operation, $\Pi_{\ker(\mathcal{L})}$ — the mask enforces admissibility in the architecture; the penalty enforces it in the objective. The practical consequence, demonstrated next, is that CPL needs far fewer data points than an unconstrained fit, because the physics supplies the information the missing data would have.

4. Demonstrations

All demonstrations use the generic model of Section 2. Code is provided (Section 7); each figure is reproduced by running the corresponding script.

4.1 One principle, two faces

Figure 1 shows the two faces on a common footing. On the architecture side, a light-cone mask applied to a small attention operator crushes the acausal couplings (here about 70 % of the pairs) exactly to zero. On the objective side, a regression constrained by the Helmholtz operator recovers a quarter-wave acoustic mode, $p(x) = \cos(k_0 x)$, from seven sparse noisy points, driving the physics residual $\|\mathcal{L}u\|$ from $\mathcal{O}(1)$ to $\sim 10^{-7}$; the unconstrained fit, by contrast, oscillates wildly between the data points.

4.2 The physics compensates for missing data

Figure 2 quantifies the data efficiency. We measure the error in recovering the acoustic mode as a function of the number of sparse, noisy samples, averaged over repetitions, with and without the physics constraint. The unconstrained regression needs many points to reach a given accuracy and remains erratic when data are scarce; CPL reaches the same accuracy with markedly fewer points, because the imposed physics fills the gap. This is the property that matters when each “data point” is an expensive simulation or measurement.

4.3 Locating the stability boundary from few evaluations (synthetic)

We now apply CPL to the quantity of interest: the marginal-stability flame gain n_{crit} . From equation (4), the growth rate is linear in n near the boundary, so CPL imposes a linear law; a physics-free baseline fits a flexible cubic. Treating each of six evaluations as costly and noisy, and repeating 500 times, CPL recovers n_{crit} with small scatter, whereas the flexible fit — over-fitting the few noisy points — scatters widely and often crosses zero at the wrong gain (Figure 3).




Figure 1: One principle, two faces of CPL. Left: a light-cone mask on an attention operator sets acausal couplings to zero. Right: a Helmholtz-constrained regression recovers a quarter-wave mode from sparse noisy data where an unconstrained fit diverges; the physics residual collapses to $\sim 1e-7$.




Figure 2: Mode-recovery error versus number of sparse noisy samples, with and without the physics constraint. CPL reaches a given accuracy with far fewer points; the physics compensates for data that are hard to acquire.




Figure 3: Locating the instability boundary from a few costly, noisy evaluations. Imposing the physically correct near-boundary law (linear growth rate) recovers n_{crit} tightly; a physics-free flexible fit scatters and crosses zero at the wrong gain.

4.4 Locating the boundary from the reduced model’s own points

Finally, we close the loop on the reduced model of Section 2, rather than on synthetic data. The full model obtains its stability boundary from a fifteen-point gain sweep. We instead sample only **four** of those points, add measurement-scale noise as if each were an expensive run, and apply the CPL linear law. Across repetitions, CPL locates the boundary at $n_{\text{crit}} = 1.99 \pm 0.12$ (98 % of estimates within ± 0.25), against 2.04 ± 0.58 (63 %) for the physics-free cubic — the same centre, but roughly five times less scatter. CPL recovers from four points the boundary that the model computes from fifteen (Figure 4). The panel also shows the three quarter-wave modes (485 / 509 / 534 Hz) that anchor the model.




Figure 4: CPL applied to the reduced model’s own points. Left: the three quarter-wave modes across the envelope. Centre: the fifteen real sweep points, with the CPL linear law fitted to only four of them locating $n \approx 2$, versus a physics-free cubic. Right: distribution of boundary estimates over repetitions — the physics stabilises the estimate (1.99 ± 0.12 vs 2.04 ± 0.58).

5. Discussion

The demonstrations show a consistent picture: imposing the correct physical structure — a causal mask in the architecture, or a residual penalty in the objective — lets a surrogate extract the quantity of interest from far fewer evaluations than an unconstrained fit requires. For thermoacoustic screening, the quantity of interest is the stability boundary, and CPL recovers it from a handful of points because the near-boundary growth rate is, by weak-coupling theory, linear in the flame gain [2].

The intended role is explicitly that of a **pre-screen**, not a replacement for high-fidelity simulation. In a reactive-LES campaign, the analytic n - τ flame-transfer function of Section 2 would be replaced by a *measured* FTF or flame-describing function extracted from the resolved reactive field [3,4]. CPL applies unchanged, now on measured coefficients: it would identify, from a small number of resolved operating points, which regions of the design space lie near the boundary and therefore warrant the full simulation budget. This converts a weeks-per-point effort into a targeted few.

Limitations, stated plainly. The reduced model treats only longitudinal quarter-wave modes; transverse and azimuthal modes, and geometry-dependent mode structure, require three-dimensional resolution and are outside its scope. The near-boundary linear law (4) is exactly the structure CPL exploits; away from the boundary, or under strong nonlinearity (limit cycles, mode competition), a higher-order admissible law would be required, and the projection $\Pi_{\ker(\mathcal{L})}$ must be defined with the appropriate operator. The FTF here is analytic and of order unity; no claim is made about any specific combustor. These are the honest boundaries of a reduced-order screening tool.

6. Conclusion

We have presented a physics-constrained learning method (CPL) and shown that it functions as an efficient pre-screen for thermoacoustic stability. Formulated as a projection onto the manifold admissible under the governing operator, CPL has two equivalent faces — a causal mask in an attention architecture and a physics-

residual penalty in a fitting objective. On a generic reduced-order thermoacoustic model it recovers an acoustic mode from sparse noisy data, and locates the marginal-stability flame gain from four noisy evaluations with roughly five times less scatter than a physics-free fit, matching the boundary the full model obtains from a fifteen-point sweep. The method is generic, reproducible, and free of any proprietary parameter; its purpose is to make expensive high-fidelity campaigns efficient by targeting the operating points that matter. A natural next step is to couple CPL to a measured flame-transfer function from reactive LES, and to extend the admissible law beyond the linear near-boundary regime.

7. Data and Code Availability

All models, scripts and figures are generic and fully reproducible, and contain **no proprietary geometry, calibration constant, or quantitative computational-fluid-dynamics result**. The following scripts accompany this manuscript and regenerate every figure and table:

- `reduced_thermoacoustic_model.py` — the reduced-order thermoacoustic model of Section 2 (prints the modes of Table 1 and writes the M3 gain sweep to `n3_sweep.json`).
- `cpl_demonstrations.py` — reproduces Figures 1–4 (one principle/two faces; data-efficiency; stability boundary, synthetic; and stability boundary from the model’s own points). Run `reduced_thermoacoustic_model.py` first, or let the script call it automatically.

Each script runs in seconds with a standard scientific Python environment (`numpy`, `scipy`, `matplotlib`).

Acknowledgements

The author thanks Dr. Miguel Santos Pereira (SPASS Advogados, Lisbon) for legal counsel on the associated intellectual-property strategy. This methodological note arises in the context of Project STAR; it is self-contained and independent of any proprietary parameter of that project. The conceptual framework of Project STAR is described separately (Zenodo, concept DOI: 10.5281/zenodo.20607800, which always resolves to the latest version).

References

- [1] Lieuwen, T. & Yang, V. (eds.), *Combustion Instabilities in Gas Turbine Engines*, Progress in Astronautics and Aeronautics, Vol. 210, AIAA, 2005. [2] Dowling, A. P., “Nonlinear self-excited oscillations of a ducted flame,” *J. Fluid Mech.* 346:271–290, 1997. [3] Poinso, T., “Prediction and control of combustion instabilities in real engines,” *Proc. Combust. Inst.* 36(1):1–28, 2017. [4] Polifke, W., “Modeling and analysis of premixed flame dynamics by means of distributed time delays,” *Prog. Energy Combust. Sci.* 79:100845, 2020. [5] Nicoud, F., Benoit, L., Sensiau, C. & Poinso, T., “Acoustic modes in combustors with complex impedances and multidimensional active flames,” *AIAA J.* 45(2):426–441, 2007. [6] Raissi, M., Perdikaris, P. & Karniadakis, G. E., “Physics-informed neural networks,” *J. Comput. Phys.* 378:686–707, 2019. [7] Vaswani, A. et al., “Attention is all you need,” *Adv. Neural Inf. Process. Syst.* 30, 2017. [8] Crocco, L. & Cheng, S. I., *Theory of Combustion Instability in Liquid Propellant Rocket Motors*, AGARDograph 8, Butterworths, 1956.