Unsupervised Wrap-Discontinuity Repair in Wavetable Synthesis via Hybrid GA–PPO Meta-Search

A PREPRINT

Julian M. Kleber
Independent Researcher
Email: julian.m.kleber@gmail.com

3 August 2026

Code (ReelSynth engine + DenoiseOpt meta-search): https://github.com/reeldemo/denoiseopt-paper, https://github.com/reeldemo/denoise-opt-meta

ABSTRACT

Wavetable oscillators click when the stored period does not join cleanly at the wrap. We treat that click as a local seam-repair problem and score repairs with a locked residual R_{blend} (seam match to a cliff-withheld ideal sibling, plus mid-cycle identity to the cracked engine). DenoiseOpt searches compact repair operators with a hybrid of genetic search, PPO, and population-based training. We reuse those search tools; we do not propose a new NAS or RL algorithm. Primary evidence is objective R_{blend}, not listening scores (no MOS/MUSHRA; $N_{\text{listeners}}$=0).

On a five-seed holdout, a short-FitCell hybrid freeze reaches 0.9708±0.0014, below frozen Noise2Noise (0.9767±0.0015) but far above Dual Cosine (0.542±0.038; Table 7). A matched 5,000-evaluation bake-off across three search seeds ranks the hybrid first among six outer-loop families (0.9748±0.0018; Table 21). Training that hybrid to residual plateau then clears the Noise2Noise gate (R_{blend}≈0.9912 vs gate ≈0.9750; Table 5). On hard wrap cliffs, classical fades collapse while searched operators hold. Cross-domain boards reuse the same wrap protocol as residual stress tests only.

Keywords wavetable synthesis; wrap discontinuity; unsupervised repair; DenoiseOpt; seam restoration; Noise2Noise; neural architecture search; reinforcement learning; genetic algorithms; proximal policy optimization; hybrid meta-search; population-based training; audio signal processing

1 Introduction

Wavetable oscillators store one period and wrap the read pointer at the end of the table. When those endpoints disagree, each wrap injects a discontinuity that listeners hear as a period-locked click.

Wavetable synthesis in brief. A wavetable synthesizer stores a single cycle of a waveform (the wavetable) and produces continuous tone by looping a read pointer through that table. Pitch is set by how fast the pointer advances. Timbre comes from the stored shape (and, in multi-table banks, from morphing across tables). Tiling the stored period yields a sustained note under playback. If the endpoints disagree ($x[0] \neq x[L-1]$), every wrap is a discontinuity at the wrap join (the seam), heard as a click locked to the fundamental. Figure 1 shows one stored period, tiled playback with wrap marks, and a zoomed mismatched wrap. The same geometry appears whenever a short loop is tiled for listening or scoring.

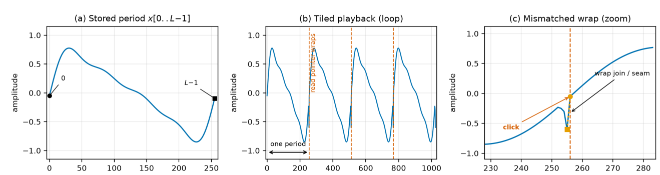

Figure 1: Wavetable synthesis primer. (a) One stored period $x[0..L-1]$ (L=256). (b) The same period tiled for continuous playback. Dashed lines mark read-pointer wraps. (c) Zoom at a mismatched wrap join (seam), where the endpoint cliff produces an audible click.

Terminology. Key names used below. A **seam** (wrap join) is the local end/start junction where endpoint mismatch produces the period-locked click: the seam is local, while the signal is the whole waveform. **Repair operator** Θ means a searchable local edit at that seam (legacy: bake cell), and a tiny network inside Θ is a **residual operator**. **FitCell** continuously fits one candidate Θ under residual loss (Appendix A). Scores R and R_{blend} are unsupervised: prolonged R is debug-only, and R_{blend} drives primary selection together with the latency-aware objective $J = R_{\text{blend}} - \lambda \cdot \text{latency\_norm}$ (Section 3.3). **MoE** soft-gates bake experts, and **PBT** mutates elites in the outer loop. The **outer loop** (meta-search) runs GA plus PPO over repair-operator graphs rather than end-to-

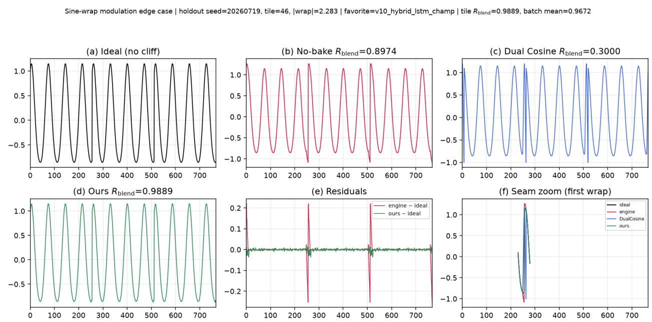

Figure 2: Hard sine-wrap holdout tile (seed `20260719`). (a) Ideal sibling $r^\star$. (b) No-bake cracked engine. (c) Dual Cosine fade. (d) Ours (searched bake). (e)–(f) Residual and seam zoom. Accompanying JSON scores use locked R_{blend} (α=0.7; tile 46: no-bake 0.897, Dual Cosine 0.300, Ours 0.989; Table 7).

end training of one large network. A **holdout** freezes the evaluation set or seed for reported scores.

Unsupervised, as used here. Search never trains on paired studio-clean audio. The ideal sibling $r^\star$ is a procedural cliff-withheld twin of the cracked engine (self-supervised geometry for scoring), not a human-labeled clean recording. Operators are ranked under R_{blend} without requiring labeled clean targets at search time.

We treat the task as *cycle-local seam restoration*: edit a neighborhood of the wrap join (the seam), and optionally a small residual operator over the period, so multi-period playback matches an ideal reference that never opens the wrap cliff. Agreement uses a discontinuity-weighted blend R_{blend} (seam fidelity vs the ideal, plus mid-cycle identity vs the cracked engine) under a latency-aware objective J (Section 3.3).

Four roles. Keep these distinct (Figure 2 and Section 3.4):

1. **Ideal sibling** $r^\star$: cliff withheld. Scoring target only. Not a method.
2. **No-bake**: unrepaired cracked engine x, scored vs $r^\star$. A do-nothing reference.
3. **Dual Cosine**: classical raised-cosine end fade. Classical baseline.
4. **Ours (DenoiseOpt)**: searched repair operator from the hybrid outer loop.

Corrupt-to-corrupt **Noise2Noise** is a fifth comparator: the primary neural gate under the locked R_{blend} protocol (Section 5.2). It is not shown in Figure 2, which illustrates roles on one hard sine-wrap tile.

Classical virtual-analog tools already address wrap energy locally (BLIT/BLEP [1], PolyBLEP [3], BLAMP [2]). We score cycle-local bake versions of those operators beside Dual Cosine (Section 5.15). Label-free restoration such as Noise2Noise shows that useful reconstructors can be ranked without paired clean targets when the evaluation geometry is explicit [4,5]. On the search side, we *reuse* known NAS, aging evolution, PBT, CMA-ES, and PPO tools [16,19,21–23]: GA for diversity, PPO for discrete architecture edits, PBT-style exploit–mutate on elites. The engineering choice is to combine them as an outer loop over compact seam-repair graphs; that combination is not claimed as an algorithmic contribution.

Key findings. Under the locked v10.1 R_{blend} protocol with multi-seed evaluation (primary story: Table 7):

- On the primary holdout, the short-FitCell hybrid champion sits just below frozen Noise2Noise—the strongest neural baseline for this locked protocol—on every seed (0.9708 ± 0.0014 vs 0.9767 ± 0.0015) and well above Dual Cosine (0.542 ± 0.038).
- On twenty multi-family boards, Ours and Noise2Noise are statistically indistinguishable on board-mean deltas (0.9337 ± 0.0028 vs 0.9308 ± 0.0045; mean win fraction 0.47; sign-test p=1.0).
- Matched-5k outer-loop comparison ranks hybrid first among six search families (0.9748 ± 0.0018, three search seeds); a hybrid-only FitCell-to-plateau follow-up then clears the Noise2Noise gate on seed `1902771841` ($R_{\text{blend}} \approx 0.9912$; Section 5.3).
- On hard wrap cliffs, classical fades collapse while searched operators and Noise2Noise hold (Section 5.11).
- On Factory+FX and AKWF corpora under R_{blend}, Noise2Noise leads Ours; both still beat Dual Cosine.

We do not claim a new NAS/RL algorithm. Existing NAS, aging evolution, PBT, CMA-ES, and PPO primitives are reused [16, 19, 21, 22]. Any novelty is in the wrap-discontinuity problem framing, the compact repair-operator vocabulary, and the unsupervised sibling/R_{blend} evaluation—not in inventing search algorithms. The engineering story is applying a hybrid GA+PPO(+PBT) outer loop to those operators, then selecting hybrid from a matched six-way bake-off before deepening FitCell. Compact searchable repair operators remain competitive when classical fades fail and need no paired clean targets at search time. Whole-curve prolonged R near 0.991 is a superseded debug score, not a current holdout result under R_{blend}.

Transfer datasets. Transfer boards apply the *same* wrap-discontinuity protocol to other short periodic signals (Section 4.2): open a cliff at the period join, score against a no-cliff ideal sibling. They are mechanistic stress tests of wrap geometry under that protocol, not evidence of general domain diagnosis. Boards: **CWRU**, **MFPT**, and **Paderborn** K001 (one shaft revolution as a period), **MIT-BIH** and a **PTB-XL** 500 Hz subset (one heartbeat as a period), plus **synthetic CNC** and **synthetic PMU/AC** pilots when real downloads are walled (Table 15).

Scope. Cycle-local wavetable seam repair for DSP / synthesis venues (DAFx / AES / arXiv cs.SD). Formal wrap-closure proofs, search-convergence theorems, and speech/music enhancement SOTA are out of scope; caveats live in Section 9.

1.1 Contributions

1. **Problem framing.** Wrap discontinuity as cycle-local seam restoration under tiled evaluation, with explicit R_{seam}, R_{body}, and primary locked R_{blend} (α=0.7) plus latency-aware J.

2. **Ideal-sibling protocol.** Unsupervised scoring via a procedural cliff-withheld twin plus mid-cycle engine identity, without paired studio-clean targets at search time.

3. **Compact repair-operator space.** Searchable graphs over classical seam ops (DualCosine, FIR, VA residuals) and tiny residual nets (MLP/FIR/U-Net-lite/LSTM-lite/MoE), not end-to-end Demucs-scale stacks (Section 3.4, Appendix A).

4. **Applied hybrid search + selection evidence.** DenoiseOpt applies GA–PPO(+PBT) to that operator space; a matched six-way 5k bake-off under locked $R_{\text{blend}}+J$ ranks hybrid first, motivating hybrid-only FitCell-to-plateau (Tables 21, 5).

5. **Open multi-seed evaluation.** Frozen Noise2Noise gate, five-seed holdout, multi-family and wavetable-native boards, hard-cliff strata, wrap-protocol transfer stress tests, and released code/artifacts (Section 8).

What is *not* claimed as a contribution: a new foundational NAS/RL algorithm, formal wrap-closure theorems, or perceptual MOS/MUSHRA superiority.

2 Related Work

We group prior work by mechanism. **Used** citations ground methods we build on or compare against under our protocol. **Screened** citations mark prior work we reviewed and ruled out (wrong domain, cost, or objective). We cite those papers for positioning only. They are not building blocks of DenoiseOpt.

Wavetable / virtual-analog foundations and discontinuity control. Foundational wavetable and virtual-analog (VA) oscillator work frames alias-free classic waveforms and wrap/corner discontinuities via BLIT/BLEP [1], PolyBLEP [3], and BLAMP [2]. We treat those OA papers as the wavetable/VA lineage for this seam task and score cycle-local approximations beside Dual Cosine (Section 5.15). DDSP frames DSP modules as composable blocks [6]. We reuse that framing for small bake/FIR/MLP seam operators at cycle scale, rather than end-to-end vocoders.

Label-free restoration. Noise2Noise learns restorers from corrupted observations alone [4,5]. Mixture-invariant training ranks unsupervised systems without stem labels [29]. We train a corrupt-to-corrupt seam CNN on the same generator as the primary neural baseline (Section 5). Full Demucs / DiffWave stacks are screened: domain and bake-budget mismatch [34, 37].

Compact waveform operators. Wave-U-Net, Conv-TasNet, dual-path, and attention audio models supply block priors we shrink to bake width [8, 11–15, 30]. They supply block priors rather than serving as frozen full-scale separators inside every meta trial.

Search and hybrids. Neural architecture search (NAS) taxonomies, reinforcement-learning (RL) controllers, progressive and latency-aware NAS, aging evolution, population-based training (PBT), CMA-ES, and evolution-guided policy gradients motivate applying a hybrid outer loop to this DSP task [16–19, 21–24, 27, 28]. Mixture-of-experts (MoE) soft gates motivate routing among small heterogeneous operators under one residual score [25]. Matched outer-loop bake-offs (Random, CMA-ES, REINFORCE, aging evolution, TPE vs hybrid) sit in Appendix B. We do not claim a new foundational NAS/RL algorithm.

Used vs screened (short). Used anchors include VA discontinuity control [1–3], unsupervised restoration [4, 5, 29], DDSP [6], compact waveform operators as above, and NAS/RL/PPO/PBT/CMA-ES/ERL/MoE as above. Screened as contrast only: MAML [32], DARTS [33], Demucs-scale enhancers [34, 36], DiffWave [37], and SEGAN-style full GAN loops [38]. Claims stay inside cycle-local seam restoration under a declared compute budget.

3 Methods

3.1 Notation

Throughout, $x \in \mathbb{R}^L$ is the cracked single-period wavetable (engine input with open-wrap cliff). Let $\mathscr{H}$ be the space of repair-operator configurations (discrete operator graph plus continuous bake parameters). The repair operator (bake operator Θ) is

$$\Theta : \mathbb{R}^L \times \mathscr{H} \to \mathbb{R}^L, \qquad y = \Theta(x;h), \quad h \in \mathscr{H}. \tag{1}$$

When continuous parameters alone are emphasized we write θ for the continuous slice of h and $y = \Theta(x;\theta)$. $r^\star \in \mathbb{R}^L$ is the procedural ideal sibling defined below. Tiling length N (typically N=16) forms $y_{\text{tiled}} = \text{Tile}(y,N)$ and $r^\star_{\text{tiled}} = \text{Tile}(r^\star,N)$. Seam width W (code: `SEAM_W`, typically 8) is the neighborhood edited by classical fade and polish ops. $R \in [0,1]$ is the prolonged residual score (1 = best). These symbols follow a different convention from Noise2Noise clean/noisy pairs.

3.2 Problem setup: cracked period and ideal sibling

Problem. Given a cracked period x, choose a bake configuration h so that the repaired cycle $y = \Theta(x;h)$ matches a no-cliff ideal under prolonged tiling.

Ideal sibling and mid-cycle engine identity (explicit). Let $G(\text{seed},\text{cliff})$ denote the procedural generator used by the CUDA FitCell batch maker (FitCell = inner continuous-parameter fit of one candidate Θ; Algorithm 5): one draw

Table 1: Core symbols used in Methods.

Symbol	Meaning
L	Cycle length (samples per period)
x	Cracked engine wavetable period (cliff on)
$y = \Theta(x;h)$	Baked cycle after seam ops / residual operator
Θ	Repair operator (bake operator) $\mathbb{R}^L \times \mathscr{H} \to \mathbb{R}^L$
$h \in \mathscr{H}$	Repair-operator configuration (ops + continuous θ)
$r^\star$	Ideal sibling (same seed, cliff withheld)
G	Procedural period generator (cliff on/off)
N	Prolong tiling count
W	Seam neighborhood width
R	Prolonged residual score in $[0,1]$ (1 = best)

of frequency, phase, and mid-cycle content, with an optional open-wrap cliff of width W at both ends. For any seed,

$$r^\star = G(\text{seed}, \text{cliff=off}), \qquad x = G(\text{seed}, \text{cliff=on}). \quad (2)$$

Thus $r^\star$ (ideal sibling) and x (cracked engine) are *paired twins of one procedural draw*: same mid-cycle waveform content; they differ only by whether the wrap cliff (and engine-only seam-boosted noise) is applied.

- $r^\star$ is a **scoring target only**. It is never emitted as a method output and is not a human-labeled clean stem.
- x is the **operator input**. Repair Θ edits x toward agreement with $r^\star$ at the seam and identity to x in the mid-cycle body.
- Locked R_{blend} (Section 3.3) combines $R_{\text{seam}} = R(r^\star, y; \mathscr{S})$ with $R_{\text{body}} = R(x, y; \mathscr{B})$ at α=0.7: seam fidelity to the cliff-withheld twin, plus mid-cycle identity to the cracked engine (so the bake cannot freely rewrite the period interior).

For an arbitrary cracked period from this family, the sibling is obtained by replaying the same seed with the cliff withheld (that is the pair returned by `make_batch`). Here *unsupervised* means: no human-labeled clean audio and no studio-clean targets in the search loop; $r^\star$ is a procedural self-supervised scoring twin, not a perceptual ground truth [4, 5].

3.3 Blended residual score and latency-aware objective

Primary metric. After tiling each cycle N times, let y be the repaired output, $r^\star$ the ideal sibling, and x the cracked engine. Define the wrap-neighborhood mask $\mathscr{S}$ of width W=8 at each period head/tail and the complementary mid-cycle body mask $\mathscr{B}$. Component scores use the unit clamp of relative RMS:

$$R(a,b;\mathscr{M}) = \operatorname{clamp}\left(1 - \frac{\operatorname{RMS}\big((a-b)\odot \mathbf{1}_{\mathscr{M}}\big)}{\operatorname{RMS}\big(a \odot \mathbf{1}_{\mathscr{M}}\big) + \varepsilon}, 0, 1\right). \quad (3)$$

Then $R_{\text{seam}} = R(r^\star, y; \mathscr{S})$ (seam fidelity vs the ideal) and $R_{\text{body}} = R(x, y; \mathscr{B})$ (mid-cycle identity vs the engine). The primary reported score is

$$R_{\text{blend}} = \alpha R_{\text{seam}} + (1-\alpha) R_{\text{body}}, \quad \alpha = 0.7. \quad (4)$$

Whole-curve prolonged R remains a debug / superseded key only.

Roles of the three quantities.

Quantity	Role	Used for
R_{blend}	Primary quality metric	Tables, figures, reporting
J	Latency-aware search objective	Champion selection
N2N gate	Frozen neural benchmark	Holdout pass/fail

Search objective. Candidates maximize latency-aware

$$J = R_{\text{blend}} - \lambda \cdot \text{latency\_norm}, \quad \text{latency\_norm} = \frac{\log(1+t_{\text{ms}})}{\log(1+50)}, \quad \lambda = 0.02. \quad (5)$$

FitCell minimizes $1 - R_{\text{blend}}$. Selection and champion updates use J. A champion must also strictly beat the frozen corrupt$\to$corrupt Noise2Noise holdout R_{blend} (primary baseline).

Every scored row applies some Θ to x and is evaluated under (4). The ideal sibling is the seam scoring target. Body identity uses the engine.

Remark 1 (Noise2Noise primary baseline). `SeamN2N` (≈53.5k) is the primary neural comparator and champ gate. Search may discover `n2n_unet` operators at the same scale. TinyUNet1D `unet` is a different, smaller block [4, 5].

3.4 Seam bake operator and classical baselines

A **repair operator** $\Theta(\cdot;h)$ composes classical seam ops (DualCosine, polish, soft seam, FIR blends) with optional tiny residual operators (MLP/FIR/U-Net-lite/attention-lite/LSTM-lite) and optional mixture-of-experts (MoE) soft gates over experts (Appendix A, Algorithms 2–4). Bake parameters sit in a bounded box; architecture strings are discrete graphs over those ops. Separate classical controls apply cycle-local BLIT/BLEP, PolyBLEP, and BLAMP residual corrections at the wrap on the cracked engine tile: they share the DualCosine / FIR interface and are not a full oscillator rewrite. Roles in plain language: BLIT/BLEP and PolyBLEP correct a *value* step at wrap, BLAMP corrects a *slope*/corner jump, and DualCosine fades endpoints without a BLEP residual. Section 5.15 scores them.

Four roles (keep separate). Four objects must stay distinct (Figure 2). The **ideal sibling** $r^\star$ withholds the cliff and is a scoring target only (panel (a)). **No-bake** is the unrepaired passthrough $\Theta_{\text{no-bake}}(x) = x$: the cracked engine scored against $r^\star$ (panel (b)). It leaves the engine unchanged and must not be confused with $r^\star$, residual-net skips, or learned identity maps. **DualCosine** is the classical raised-cosine end fade on x (panel (c)). It appears both as a classical board row and, separately, as a PPO advantage baseline (Remark 2). **Ours** is the searched bake $y = \Theta(x;h^\star)$ from the hybrid outer loop (panel (d)).

On mild cliffs, residual mass sits far below ideal RMS under the unit clamp of (4), so no-bake R_{blend} can look high even when the wrap click remains audible. Treat that number as a do-nothing passthrough reference rather than a perceptual wrap-quality score. Released JSON may still use the legacy key `identity` for no-bake. Manuscript prose uses *no-bake*.

Classical (non-AI) comparison board. Learned / searched methods are always ranked against the full classical board on absolute R (same $r^\star$): no-bake, DualCosine / cosine fade, `seam_fir3`, classic quadratic, hann/crossfade/soft fades, detrend/ensemble, poly seam, and cycle-local BLIT/BLEP / PolyBLEP / BLAMP where scored. When tables show ΔR vs DualCosine, that is a reporting gap vs one classical bake, not the definition of reward.

Remark 2 (DualCosine centering is PPO advantage only). The outer objective remains (5): maximize absolute $R(y, r^\star)$. In the PPO branch only, we optionally form a centered advantage

$$\rho = R - R_{\text{DualCosine}}, \tag{6}$$

so early-search advantages are roughly zero-mean. This does *not* change selection, FitCell loss, champion ranking, or the matched 5k meta-approach objective. All of those still maximize absolute R. DualCosine centering is soft relative shaping for the actor–critic, not a hard absolute ranking target.

3.5 Search pipeline

Figure 3 summarizes the residual-scored hybrid loop. Each outer iteration proposes a bake configuration h, runs FitCell (Appendix A, Algorithm 5) to fit continuous parameters under loss $1 - R_{\text{blend}}$, evaluates R_{blend} (and J) vs $r^\star$, and updates the population / PPO policy.

Main-text algorithm summary (before appendix listings).

1. **FitCell (inner):** Adam on $1 - R_{\text{blend}}$ over i.i.d. G batches until relative residual change $< \varepsilon$ for patience P steps, or hard cap $T_{\max}$ (Algorithm 5).

2. **GA branch:** tournament parents $\rightarrow$ crossover/mutate op graph $\rightarrow$ FitCell $\rightarrow$ insert by J (Algorithm 6).

3. **PPO branch:** actor proposes discrete architecture edits; advantage uses optional DualCosine centering; clipped surrogate update (Algorithm 7); selection still maximizes absolute R_{blend}/J.

4. **PBT branch:** copy hyperparameters from elites into bottom quantile; multiplicative mutate; brief refit (Algorithm 8).

5. **Hybrid rotation:** each outer step runs one or more of the above; champion is the best J cell that also clears Dual Cosine and, when required, the N2N gate (Algorithm 9).

Full line-by-line pseudocode is in Algorithms 6–9 (Appendix A).

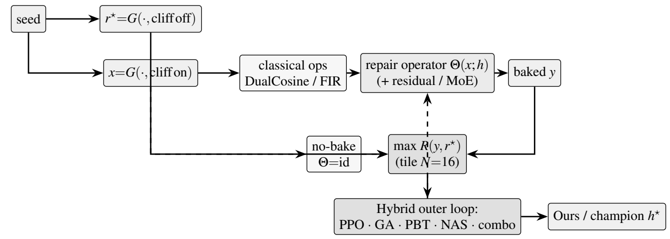

Figure 3: DenoiseOpt overview (before component deep-dives). Same-seed generator G yields ideal sibling $r^\star$ (cliff off) and cracked engine x (cliff on). Repair operator $\Theta(x;h)$ may include DualCosine/FIR and optional residual/MoE experts. No-bake is the identity control $x \mapsto x$. Outer objective is maximize absolute prolonged $R(y, r^\star)$. DualCosine centering (not shown) is PPO advantage shaping only. Dashed feedback: architecture proposals update the operator. Pseudocode in Appendix A.

3.6 Search space and inner fit

Each genome encodes: ordered op list (fade, polish, FIR, pin, residual write), depth, residual operator type (none, MLP, FIR, U-Net-lite, attn-lite, LSTM-lite), optional MoE expert count $K \in \{2, 3, 4\}$, and continuous θ in the CUDA runner box bounds. The searchable bake vocabulary includes a width-capped 1–2 layer LSTM block over the seam window or full cycle (distinct from the fixed supervised seq-LSTM SOTA baseline). The live hybrid tag is `PPO+GA+PBT+NAS+depth+MoE` with LSTM enabled in the block graph.

FitCell draws sibling batches via G, applies $\Theta(\cdot; h)$, and steps Adam on $1 - R_{\text{blend}}$ until residual plateau or a searchable hard cap $T_{\max}$ (Algorithm 5; default patience / ε co-tuned online). The short-cap D1 FitCell ($\approx$16–48 steps) is retained only for the six-way matched bake-off; the hybrid-only follow-up uses the converge protocol (Section 5.3).

3.7 Hybrid outer loop

Each iteration rotates among branches: PPO, GA, PBT, NAS, and combo. Selection and FitCell are always by absolute residual R (maximize toward the ideal sibling, loss $1 - R$).

Near-ceiling scores need reward shaping. On the sine+cliff generator, no-bake already sits near $R \approx 0.97$, and strong classical / searched bakes live in a narrow band toward 1 (often 0.99–0.991). Absolute gaps of order 10^{-3}–10^{-2} are real but tiny as raw PPO/GA signals. Without shaping, advantages collapse. Hyperparameters (clip, entropy, mutation rate, population size, fit steps, learning rate) then decide whether the outer loop can climb the last percent. We therefore treat reward shaping and search hyperparameters as first-class tunables alongside bake architecture. In the present hybrid, PPO uses a searchable shaping mode among $\{$`abs_r`, `vs_dualcosine`, `vs_nobake`, `neglog_gap`$\}$ (PBT-mutated, default $\rho = R - R_{\text{DualCosine}}$ per Remark 2). PBT mutates fit / PPO hyperparameters and reward mode in-population. Plateau adaptation retunes branch mix and depth/MoE bias when champions stall. The matched 5k meta-approach comparison (Section 3.8) selects the outer-loop family under a shared short FitCell budget; the hybrid-only converge campaign then spends GPU on that winner.

3.8 Meta-learning approach comparison

Beyond the hybrid outer loop, we compare five additional outer-loop families under a matched 5,000-evaluation budget, three search seeds (`1902771841`, `2026072701`, `2026072702`), and the same FitCell / $R_{\text{blend}}+J$ protocol (LSTM and xLSTM included in the searchable bake vocabulary). Table and figure legends use the short display names below. Artifact folders keep stable code keys (see the nomenclature):

- **Random NAS** (`random`): independent uniform draws over bake graphs and fit hyperparameters (control).
- **Cont. CMA-ES** (`cmaes`): diagonal covariance-matrix adaptation over a continuous encoding of depth, width, operator kind, block logits, and fit hyperparameters [24].
- **Arch REINFORCE** (`reinforce`): vanilla policy gradient over discrete architecture-edit actions (no clipped surrogate). Contrast to PPO [17, 19].
- **Aging evolution** (`aging_evo`): regularized evolution with oldest-member replacement after tournament parent selection [20, 21], distinct from the hybrid's non-aging tournament GA branch.
- **TPE Bayes NAS** (`tpe`): lightweight tree-structured Parzen estimator over discrete bake categoricals in the spirit of practical Bayesian optimization [26].

We also run **Ours (hybrid GA–PPO)** (`hybrid_lstm`): the multi-branch hybrid loop for 5,000 evaluations with LSTM/xLSTM in the block vocabulary and online reward-mode / fit–PPO hyperparameter co-tuning, as the proposed matched-budget arm. Results appear in Appendix B (Table 21, Figures 18–19). Healed-sample overlays appear in Figure 7.

Remark 3 (Trial complexity). Per outer iteration the dominant cost is $O(T_{\text{fit}} \cdot C_{\text{fwd}})$ for fit steps times one forward bake. Bake-width caps (tiny residual operators) keep C_{fwd} far below full-band Demucs/DiffWave forwards screened in Related Work.

3.9 Hyperparameters

Table 2 lists the CUDA search defaults used for the 5k-gate freeze reported in Results. Under the near-ceiling residual regime (Section 3.7), these outer-loop and fit hyperparameters are co-tuned with architecture. PBT and plateau adaptation continue to mutate several of them online. Companion listings: Algorithms 1–9 (Appendix A).

Table 2: Hybrid search hyperparameters (RTX 3090 DenoiseOpt CUDA search-runner defaults). Co-tuned with reward shaping under near-ceiling R (Section 3.7).

Parameter	Value
Population size n	12
GA tournament k	3
GA crossover fraction p_c	0.5 (else clone elite parent)
GA mutation	≥ 1 mutate after inherit (+50% second mutate)
PPO clip ε	0.2 (default, searchable $\{0.1, 0.2, 0.25, 0.3\}$)
Adam learning rate (fit)	3×10^{-3} (default, PBT-searchable)
Adam learning rate (policy)	3×10^{-4}
Entropy coefficient (PPO)	0.02 (default, searchable ≈ 0.005–0.06)
PBT exploit / mutate	elite fraction 0.25, multi-step arch+hp mutate
MoE routing	`sequential` or `moe_parallel` when enabled
Plateau boredom threshold	1000 iters without champ improve
Cycle length L / tiles N / seam W	256 / 16 / 8
Fit steps / batch (default)	24 / 48
Algo tag	`PPO+GA+PBT+NAS+depth+MoE`

4 Experiments

Evaluation protocol (locked). All primary tables use the same locked residual. Given a cracked engine tile $x \in \mathbb{R}^L$ and an ideal sibling $r^\star$ (cliff withheld), prolong each cycle N=16 times, form wrap neighborhoods of width W=8 at every period head/tail, and compute unit residuals

$$R_{\text{seam}} = \operatorname{clamp}\big(1 - \text{rms}_{\text{seam}}(y, r^\star)/\text{rms}_{\text{seam}}(x, r^\star), 0, 1\big),$$

$$R_{\text{body}} = \operatorname{clamp}\big(1 - \text{rms}_{\text{body}}(y, x)/\text{rms}_{\text{ref}}, 0, 1\big),$$

then

$$R_{\text{blend}} = \alpha R_{\text{seam}} + (1-\alpha) R_{\text{body}}, \qquad \alpha{=}0.7.$$

Body residual is scored against the engine x (not $r^\star$) so mid-cycle identity is rewarded. No per-table rescaling: reported scores are that clamp in $[0, 1]$. Search maximizes $J = R_{\text{blend}} - \lambda \cdot \text{latency\_norm}$ with λ=0.02 and $\text{latency\_norm} = \log(1+t_{\text{ms}})/\log(1+50)$. The champion gate is to strictly beat frozen Noise2Noise corrupt→corrupt holdout R_{blend}. Secondary keys (SNR, SDR, wrap-jump, seam-local edge RMSE) are reported where tabulated; whole-curve prolonged R is debug-only and never used for selection. PESQ/STOI/MUSHRA are out of scope on non-speech cycles ($N_{\text{listeners}}$=0).

Which numbers are multi-seed. Holdout / multi-family means are mean±std over five evaluation seeds (`20260719–20260723`). Matched-5k outer-loop bake-off means are mean±std over three search seeds

(`1902771841`, `2026072701`, `2026072702`). The FitCell-to-plateau follow-up (Section 5.3) is a single-seed freeze (`1902771841`), not a three-seed mean. N2N train seed is `424242`. Transfer Table 15 is one domain-trained hybrid per board (search seed `1902771841`), not multi-seed averages. We report the hybrid_lstm freeze plus holdout / multi-family / Factory+FX / AKWF / transfer benches. Venue: DSP / synthesis artifact repair (DAFx / AES / arXiv cs.SD).

Statistical reporting (what we compute). Unless a table caption says otherwise, multi-seed cells are arithmetic mean$\pm$sample std over the named seed list. Paired multi-family comparisons use a two-sided sign test on per-seed board-mean deltas (reported p values; p=1.0 means no sign advantage). We do *not* run ANOVA across outer-loop families beyond mean$\pm$std in Table 21, and we do not invent MOS confidence intervals. Listening: $N_{\text{listeners}}$=0 (Table 14); no preference rates.

Experimental setup (datasets, compute).

- **Primary procedural datasets.** Generator G: L=256 two-harmonic sine+cliff; seam width W=8; cliff amplitude $\pm\mathscr{U}(0.08, 0.43)$; FitCell/search draws i.i.d.; frozen holdout metrics draw n=4096 (seed `20260719`) plus five-seed score batches (`20260719`–`20260723`). Multi-family SOTA: 10 families $\times$ 2 seeds $\times$ batch 64 (n=1280). Factory+FX / AKWF wavetable-native boards: n=1280 each (ReelSynth exports; AKWF CC0).
- **Transfer stress boards.** CWRU / MFPT / Paderborn K001 / MIT-BIH / PTB-XL subset / synth CNC / synth PMU: n=256 periods each under the same wrap protocol (Table 3; public sources + procedural pilots).
- **FitCell / outer loop.** Adam on $1 - R_{\text{blend}}$; D1 short cap $\approx$16–48 steps; converge campaign searchable patience/ε/$T_{\max} \in \{512, \ldots, 2048\}$, batch $\approx$96, two sequential proposals/iter (Section 3.5). Matched 5k evals: Random, CMA-ES, REINFORCE, Aging, TPE, hybrid (3 search seeds). Hybrid-only FitCell-to-plateau: 750 outer iters, seed `1902771841` ($\approx$46.9 h wall).
- **Baselines.** no-bake, DualCosine, FIR/endpoint classical board, SeamN2N ($\approx$53.5k, 4000 Adam steps, train seed `424242`), VA residuals where scored.
- **Compute.** Single NVIDIA GeForce RTX 3090 (24 GB); host desktop Windows; Python 3.12, PyTorch 2.6+CUDA 12.4; no multi-GPU. Historical prolonged-R 5k-style freeze $\approx$7.9 h (Table 23); D1 hybrid mean wall $\approx$8.4 h/seed.
- **Artifacts.** Frozen JSON matrices and figure sources ship in the paper repository (Section 8).

4.1 Dataset

Creation. The paper-facing evaluation corpus is a frozen draw from the same synthetic generator used by the CUDA search runner (`make_batch`). Cycles are length L=256. A two-harmonic sine is drawn with frequency $\sim \mathscr{U}(1,4)$ and phase $\sim \mathscr{U}(0, 2\pi)$, then an open-wrap seam cliff of amplitude $\pm\mathscr{U}(0.08, 0.43)$ is applied over `SEAM_W`=8 samples at both ends. Seam-boosted Gaussian noise (scale 0.02 at the ends, 0.15$\times$ in the mid-cycle) forms the engine input. The ideal withholds the cliff and serves only as scoring target $r^\star$, never as a method output, while no-bake leaves the cracked engine unchanged and scores that unrepaired x against the same $r^\star$. Prolonged residual R tiles each cycle N=16 times before the RMS ratio. Holdout seed `20260719` is distinct from hybrid search seed `1902771841`. The search campaign draws fresh i.i.d. batches from this generator and never trains on the frozen holdout tensors: there is no supervised clean/noisy training split of studio audio.

Multi-family diversity ($\geq$ 20 waveforms). Beyond the single sine+cliff holdout, Phase 3a scores methods on **20** generative draws: ten family variants (`sine_cliff`, `harmonic_fft`, `am_fm`, `nonlinear`, `combo`, `triple_mix`, `extreme_overlay`, `open_wrap_bias`, `soft_noise`, `wide_seam`) $\times$ two seeds (`20260719` and `20260720`), batch 64 each (128 cycles per family, 1,280 total). These Python generative variants stress harmonic / AM–FM / nonlinear / overlay / open-wrap behavior of `make_batch` and remain the primary SOTA matrix. Separately, `export_sound_bench_tiles` exports **20** Rust `sound_bench` engine/ideal pairs (10 families $\times$ 2 seeds, L=256, byte-aligned with `generate_sound` / `generate_sound_ideal`) scored under the same residual geometry as a secondary transfer block (Table 13). Every method sees the same tensors per waveform. **Primary classical comparison** is absolute R_{blend} against the non-AI board (no-bake, DualCosine, `seam_fir3`, poly, fades, VA residuals, Section 3.4). The search objective is maximize R toward the ideal sibling (best R=1). DualCosine is one classical comparator and a PPO centering constant. It is not the reward target.

Dataset statistics. Table 3 and Figures 4–5 inventory every evaluation board used in this paper. The primary procedural holdout is a frozen 4,096-cycle metrics draw (L=256, N=16 tiles, seed `20260719`) plus a matched score batch of 64 cycles (five consecutive seeds for classical mean$\pm$std). Multi-family SOTA scores 20 waveforms (10 families $\times$ 2 seeds, batch 64, 1,280 cycles). Wavetable-native boards use ReelSynth Factory+FX exports (n=1,280: 640 dry + 640 FX) and AKWF CC0 cycles (n=1,280) so Plot 4 corpus bars are comparable to multi-family (1,280). Sci/eng transfer tiles 1,536 periods across CWRU, MFPT (n=256), MIT-BIH, PTB-XL subset, and synthetic CNC/PMU. Search never trains on the frozen holdout. N2N/seq baselines use disjoint train seeds `424242`/`424243`. Artifacts: dataset artifacts.

Table 3: Dataset statistics (evaluation boards). L=256 unless noted. Hours are raw source duration on disk where measurable, not tiled period count. Plot 4 compares ≈1,280-height boards only (multi-family / Factory+FX / AKWF / transfer).

Board	Size	Notes
Holdout metrics	4,096 cycles	seed `20260719`, N=16 tiles
Score batch	64 cycles	classical / favorite tables
Multi-family SOTA	1,280 cycles	10 families × 2 seeds × 64
Factory+FX export	1,280 periods	640 dry + 640 FX mid-tile extracts
AKWF OA	1,280 periods	CC0 single-cycle WAVs
Meta hybrid_lstm	1,200 iters	$R_{\text{blend}}+J$, seed `1902771841`
CWRU	256 periods	4 MAT files, DE @12 kHz
MFPT	256 periods	shaft-rate windows
MIT-BIH	256 periods	10 records @360 Hz
PTB-XL subset	256 periods	94 `*_hr` @500 Hz
synth CNC / PMU	256+256	procedural pilots

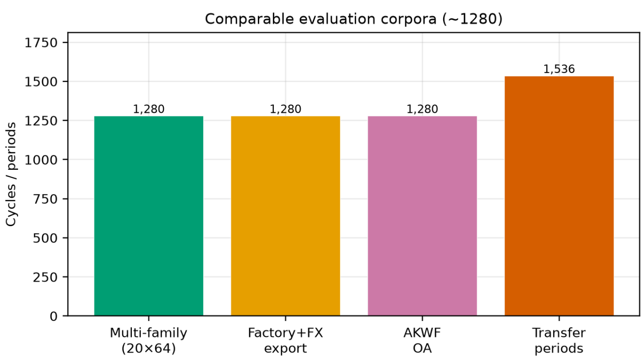

Figure 4: Corpus sizes across evaluation boards (cycles or periods). Holdout metrics draw, matched score batch, multi-family SOTA, factory/AKWF wavetable-native sets, and summed transfer periods.

Dataset metrics. A metrics draw of 4,096 cycles under the holdout seed yields mean ideal RMS 0.721 ± 0.016, mean engine residual RMS 0.021 ± 0.013, and mean absolute wrap jump $|x_0 - x_{L-1}|$ of 0.913 ± 0.634. No-bake (passthrough) on that draw scores mean $R \approx 0.971$. Method tables use a matched score batch of 64 cycles from the same seed (the canonical evaluation tensors), five consecutive seeds for mean±std on the classical table, and the 20-waveform matrix for SOTA secondary metrics. Figure 6 summarizes the 4,096-cycle metrics draw: ideal RMS clusters near 0.72, wrap jumps span nearly three orders of magnitude (median ≈ 0.83), and no-bake R is high on easy tiles but falls below 0.91 on the hardest decile of wrap jumps. Figure 2 is the documented sine-wrap edge case: a high-$|wrap|$ holdout tile where cliff modulation is visible, scored with DualCosine and the top neural favorite under the same residual R (tile $R \approx 0.9917$, favorite batch mean $R \approx 0.991$).

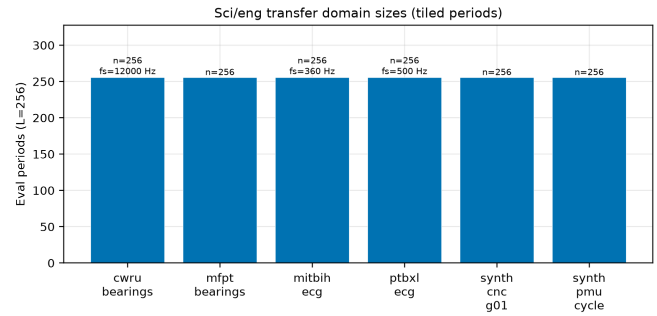

Figure 5: Sci/eng transfer domain sizes: tiled eval periods at L=256 (with source f_s where applicable).

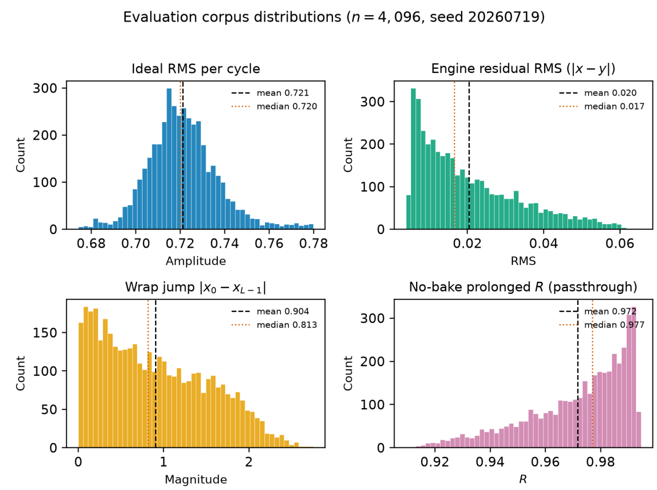

Figure 6: Holdout distribution statistics (seed `20260719`, n=4,096 cycles). Distinguish: ideal sibling $r^\star$ (cliff withheld), no-bake = unrepaired cracked engine (not $r^\star$), and DualCosine = one classical bake. Top left: ideal RMS per cycle. Top right: engine residual RMS before any bake. Bottom left: wrap jump magnitude $|x_0 - x_{L-1}|$ (cliff hardness proxy). Bottom right: no-bake prolonged R (passthrough, $x \mapsto x$). Dashed and dotted vertical lines mark mean and median. Hard-cliff reporting (Section 5) uses the top 10% wrap-jump stratum where no-bake R drops.

Hybrid search campaign (5,000-iteration freeze). A CUDA search tagged `PPO+GA+PBT+NAS+depth+MoE` runs with a large declared iteration budget (order 10^5). At the reported freeze (run `gpu-rl-arch-20260719T083019Z`, ≈ 6,922 clean iters / ≈ 7.9 h on RTX 3090) the champion residual is $R \approx 0.99093$ against DualCosine baseline ≈ 0.81658 ($\Delta R \approx +0.174$). Branch-best residuals at the freeze: GA ≈ 0.99016, PBT ≈ 0.99012, PPO ≈ 0.98924, combo ≈ 0.98854, NAS ≈ 0.98677. Isolated short-budget re-runs (GA-only, PPO-only, GA+PPO, full, 150 iters, same search seed, plateau adapt off) are reported beside those freezes in

Table 22. Figures in Section 5 regenerate from the archived search results.

Inference, classical, and learned baselines. High-R fitted operators ($R \geq 0.98$) are timed on CUDA (batch 64, prolonged residual). Ten non-learnable bake denoisers share the frozen canonical holdout against the neural favorite (Table 9). Cycle-local BLIT/BLEP, PolyBLEP, and BLAMP seam repairs are scored on the same holdout / cliff / multi-family protocol (Table 12 and the VA results) and illustrated on the intro high-wrap tile (Figure 11). Fixed-arch MLP-on-R and CNN/U-Net-lite-on-R baselines (no outer NAS) are trained on $1-R$ with the same generator and scored in the SOTA matrix. Phase B adds a primary corrupt→corrupt Noise2Noise seam CNN (train seed `424242`, disjoint from holdout), a secondary sibling-supervised ceiling on the same architecture, and lightweight LSTM / depthwise 1D-CNN seq restorers (train seed `424243`). No search-champion weights initialize those baselines. Family-conditional hardness uses the separate Rust procedural 100,000-cycle bench reported in Results.

Cliff strata and wavetable-native realism. A 4,096-tile holdout draw (seed `20260719`) is stratified by engine wrap-jump percentiles (top 25% / top 10%), with cutoffs in the cliff-strata data. Hard-cliff is the operative regime for claims about DualCosine collapse vs favorite retention. Edge RMSE is the locked required seam-local secondary. Click energy is optional. Wavetable-native transfer (Phase F1) scores true ReelSynth-exported factory periods (primary) and external AKWF CC0 WAVs (secondary) under the same close-seam→open-wrap protocol. Procedural stand-ins are tertiary smoke only. Speech/music LibriSpeech/MUSDB corpora remain off.

Audible demos and vibrato spectrograms. Downloadable clips (no-bake / DualCosine / Ours, 44.1 kHz mono, A4, 3 s): hear-sample WAVs (n=5 tiles × 3 methods = 15 WAVs). Rebuild script: `export_meta_hear_samples.py`. Slow-vibrato spectrogram protocol (`bench_vibrato_spectrogram.py`, artifacts: vibrato folder) scores the same bake operators under modulated-pitch playback. Figures appear in Section 5.18. These assets support informal audition only. **Listening statistics:** $N_{\text{listeners}}$=0; no MOS, MUSHRA, forced-choice A/B rates, or significance tests are reported (and none are invented). Objective companions are cycle-local / prolonged R on the same tiles. The planned listening protocol (if run later) is in Section 5.22.

4.2 Cross-domain wrap transfer protocol

The primary claim of this paper is wavetable seam repair. Transfer boards answer a narrower mechanistic question: does the *same* wrap-discontinuity protocol (open a cliff at a period join, score against a no-cliff ideal sibling, fit a compact repair operator) still make sense on other short periodic signals? They are stress tests of that geometry. They are *not* claims that an audio-trained bake transfers as bearing-fault or ECG diagnosis SOTA, and they are not claims that fatigue cracks or grid phase events "are" wavetable seams. Reported transfer evidence is only locked R_{blend} on the seven wrap boards in Table 15; we do not claim cross-domain generalization beyond that residual protocol.

Shared transfer recipe. For each domain we cut fixed-length periods (L=256) and build an ideal sibling (smooth / template / cubic path) plus a cracked engine (same content with an open-wrap cliff, sometimes with a deliberately worse interpolant). We then score classical baselines and a domain-trained DenoiseOpt / Noise2Noise stack under locked R_{blend} as defined in Section 3.3 (`hybrid_lstm` FitCell, 150–250 outer iterations where searched, short `fit_steps` = 40 on transfer boards). Classical baselines on every board: no-bake (passthrough), DualCosine (raised-cosine end fade), `seam_fir3`, and endpoint-pin (shared DualCosine/FIR interface; Section 3.4). SeamN2N: corrupt→corrupt CNN, ≈53.5k params, 4000 Adam steps, train seed `424242`. Domain-trained Noise2Noise (corrupt→corrupt, 4000 Adam steps) is scored on the same seven boards and reported beside classical rows in Table 15. Artifacts: signal_heal_transfer.

Seam analogy by domain family. Bearing revolutions (CWRU / MFPT / Paderborn). One shaft revolution is the period. The seam is the join between the last and first angle samples after resampling to L=256. A wrap cliff models a bad change-of-time / interpolant at that join, not a physics model of fatigue crack growth. Ideal siblings use a smoother angular path (cubic) than the cracked engine (linear + cliff).

ECG beats (MIT-BIH / PTB-XL subset). One heartbeat (R–R window) is the period. The seam is the join between beat end and beat start when the window is tiled. A wrap cliff is an artificial endpoint mismatch on that window. The ideal is a mild endpoint-equalized beat template. Clinical arrhythmia restore metrics are out of scope.

Synthetic CNC / PMU. Stand-ins for closed machine-tool contours and one AC / PMU voltage cycle when real KIT / IEEE DataPort downloads are walled. The seam is again a period-join cliff under the same residual protocol. These rows are procedural pilots, not industrial toolpath or live-grid claims.

What each dataset is (pointers). CWRU / MFPT / Paderborn K001: public vibration, n=256 revolutions each (Paderborn also hosts a separate Al Firdausi 4-class fault classifier, labeled apart from wrap-R_{blend}). MIT-BIH and PTB-XL 500 Hz lead-I subset: PhysioNet beat windows, n=256. Synthetic CNC G01 and synthetic PMU/AC: procedural, n=256.

Scope note. Transfer tables compare classical rows, domain-trained Noise2Noise, and Ours under this residual protocol. Author ECG Cycle-GAN and BeatDiff carry OOD wrap probes, while Paderborn Al Firdausi is a trained fault classifier (plus labeled arch-reuse wrap bake). Formal listening remains out of scope (Section 9 and Table 20). Results appear in Section 5.21.

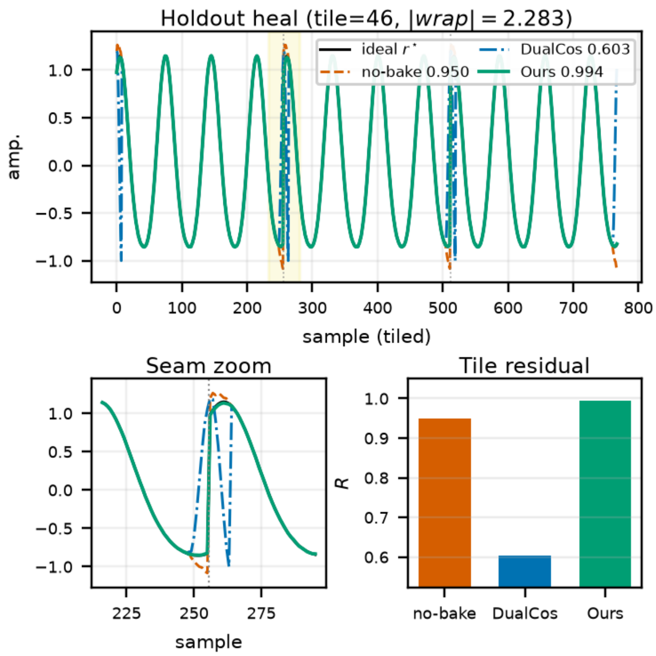

Figure 7: Holdout wrap-seam heal for the FitCell-to-plateau hybrid champion ($R_{\text{blend}}\approx 0.9912$, seed `1902771841`) vs no-bake, Dual Cosine, and ideal sibling $r^{\star}$.

5 Results

This section reports objective wrap-repair scores under locked R_{blend} (Sections 5.1–5.2). Higher is better. Unless noted, multi-seed cells are mean±std. Appendix B holds the matched outer-loop bake-off, ablations, HP probes, and transfer latency. Limitations are collected in Section 9. The evaluation protocol is stated in Section 4.

How to read the main numbers.

1. **Holdout vs Noise2Noise and Dual Cosine** (Table 7): short-FitCell hybrid is close to, but below, Noise2Noise, and far above Dual Cosine.
2. **Which outer loop to keep** (Table 21): hybrid wins a matched 5k bake-off, so later compute goes to hybrid alone.
3. **Longer FitCell** (Table 5): FitCell-to-plateau hybrid training then clears the Noise2Noise residual gate.
4. **Where classical fails** (Section 5.11): hard wrap cliffs collapse Dual Cosine; searched operators and Noise2Noise hold.

5.1 A short hybrid search under R_{blend}

We first ask whether a hybrid GA+PPO search can beat frozen Noise2Noise under R_{blend} and J. A 1,200-iteration run (seed `1902771841`, ≈1.65 h) selects on J (5) with λ=0.02. The champion reaches $R_{\text{blend}}\approx 0.9695$ ($J\approx 0.9603$, $t\approx 5.08$ ms). Dual Cosine on the same runner scores ≈0.504. The frozen Noise2Noise holdout gate is ≈0.9750. So this short search clears Dual Cosine but not Noise2Noise. Artifacts: hybrid search data.

5.2 Noise2Noise vs Ours on four boards

Table 7 is the main comparison under locked R_{blend}. Dual Cosine is the classical fade on every board. Noise2Noise is the primary neural gate for this protocol (not a DenoiseOpt variant).

Holdout. Over five evaluation seeds (`20260719`–`20260723`), Noise2Noise leads the short-FitCell hybrid on every seed: 0.9767±0.0015 vs 0.9708±0.0014. Dual Cosine sits at 0.542±0.038. The hybrid is therefore competitive with Noise2Noise and far better than a classical fade, but it does not win the holdout gate under short FitCell.

Multi-family board. On twenty generative waveforms, board means are close: Ours 0.9337±0.0028, Noise2Noise 0.9308±0.0045 (mean win fraction 0.47; sign test on board-mean deltas p=1.0). That is no significant mean difference. Both beat Dual Cosine.

Factory and AKWF. On wavetable-native corpora (n=1,280 each), Noise2Noise leads both Factory+FX (0.771 vs Ours 0.699; Dual Cosine 0.589) and AKWF (0.731 vs Ours 0.679; Dual Cosine 0.424).

Table 4: Parameter count and inference latency under the locked holdout story. Favorite Θ is the v10.1 hybrid champion (search seed `1902771841`). SeamN2N is the frozen corrupt→corrupt gate (train seed `424242`). Holdout R_{blend} is mean±std over five evaluation seeds. Latency is CUDA ms/batch at B=64 (Table 6).

Method	Params	Holdout R_{blend}	ms/B=64
Favorite Θ (v10.1 hybrid)	≈121k	0.9708±0.0014	8.20
SeamN2N (corrupt→corrupt)	≈53.5k	0.9767±0.0015	1.12
Dual Cosine	0	0.542±0.038	—

Table 4 shows size and speed. The hybrid favorite is larger and about 7.4× slower per holdout batch than SeamN2N. That is partly by design: selection maximizes $J = R_{\text{blend}} - \lambda \cdot \text{latency\_norm}$, so a slower cell can win when residual gains outweigh the λ penalty. We do not claim architectural superiority over Noise2Noise on this freeze. Search seeds, the N2N train seed `424242`, and holdout seeds remain disjoint.

Seed protocol. Holdout and multi-family rows in Table 7 use five evaluation seeds. The matched 5k outer-loop bake-off uses three search seeds (Table 21: hybrid 0.9748±0.0018, Aging 0.9635±0.0049).

5.3 Hybrid-only FitCell-to-convergence

The matched bake-off (Table 21) already ranked hybrid first under a short FitCell budget (≈16–48 steps). We therefore stop spending matched budget on Random, CMA-ES,

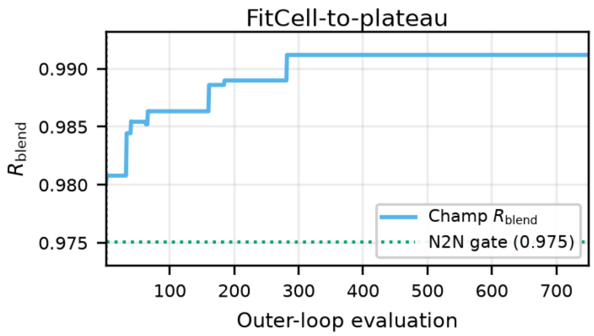

Figure 8: Hybrid FitCell-to-plateau learning curve (seed `1902771841`, 750 evaluations). Champion R_{blend} vs. outer-loop evaluation with the frozen Noise2Noise gate (y-axis zoomed to the high-R regime; Dual Cosine sits far below). Final champion R_{blend}≈0.9912.

Table 5: Hybrid FitCell-to-plateau under locked $R_{\text{blend}}+J$ (seed `1902771841`, 750/750 evaluations). Gate: frozen Noise2Noise R_{blend}≈0.9750.

Setting	Iters	R_{blend}	t ms	vs N2N gate
FitCell-to-plateau	750/750	0.9912	8.70	clears
Short FitCell (ref.)	1,200	≈0.9695	≈5.08	below
Frozen N2N gate	—	≈0.9750	—	—

TPE, Aging, and REINFORCE. The follow-up is hybrid-only FitCell-to-plateau: every candidate trains until residual plateau (or a searchable hard cap $T_{\max} \in \{512, \dots, 2048\}$), with searchable patience, relative ε, batch size, and λ (artifacts `meta_approach_compare_v14_converge/`).

The outer budget is 750 evaluations with two sequential proposals per iteration on one GPU. On search seed `1902771841`, the champion reaches R_{blend}≈0.9912 (J≈0.9854, t≈8.70 ms) after ≈46.9 h wall and clears the frozen Noise2Noise gate (≈0.9750). Figure 8 shows the learning curve; Table 5 summarizes the freeze.

5.4 Ablations and hyperparameter checks

Two narrow checks sit in Appendix B; they do not replace Table 21.

Ablations. We either (A) read branch tags inside one recorded hybrid history, or (B) restart a short 150-iteration search with GA-only, PPO-only, GA+PPO, or full hybrid (Table 22; seed `1902771841`). Type B does not collapse short-budget R_{blend} when GA or PPO is dropped on that probe, but those scores stay well below the matched-5k hybrid mean.

Hyperparameters. A one-at-a-time ±50% probe at 500 iterations moves champion R by at most ≈0.012 (Table 24). Defaults are not claimed globally optimal.

5.5 Ours vs Noise2Noise inference latency

Table 6 compares CUDA forward time across boards and batch sizes. At the default batch B=64, Ours is about 7× slower than Noise2Noise on every board (holdout 8.20 vs 1.12 ms; multi-family mean 7.50±0.56 vs 1.06±0.07 ms).

Table 6: Inference latency: Ours vs Noise2Noise on CUDA (warmup 8, timed 30). Ratio = Ours/N2N ms/batch.

Board / batch	Ours	N2N	Ratio
Holdout (B64)	8.20	1.12	7.4×
Multi-family (B64)	7.50±0.56	1.06±0.07	7.1×
Factory+FX (B64)	7.31	1.06	6.9×
AKWF (B64)	7.50	1.11	6.7×
Holdout (B1)	8.52	1.30	6.6×
Holdout (B8)	7.71	1.07	7.2×
Holdout (B256)	7.62	1.52	5.0×

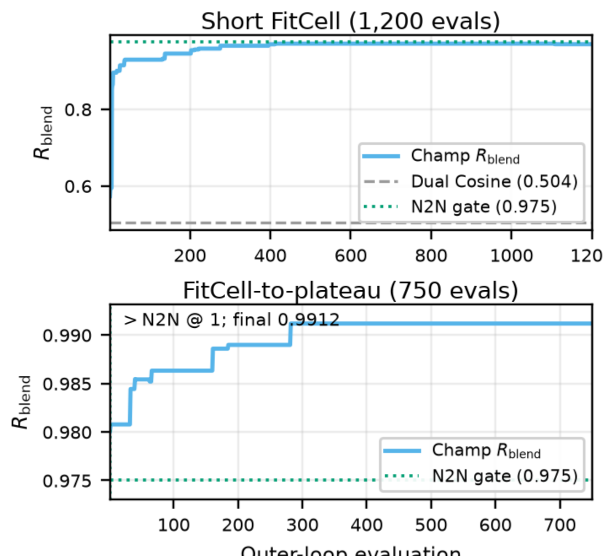

Figure 9: Hybrid search learning curves under R_{blend} (seed `1902771841`). Top: short FitCell (1,200 evaluations; champ. ≈0.9695; Dual Cosine + N2N gate). Bottom: FitCell-to-plateau (750 evaluations; champ. ≈0.9912; y-axis zoomed to the high-R regime).

A holdout batch sweep ($B \in \{1, 8, 64, 256\}$) keeps the ratio near 5–7×. Larger batches help both methods, but Noise2Noise stays faster. This matches the parameter gap (≈121k vs ≈53.5k) and the J latency term in (5). Artifacts: `v13_ours_vs_n2n_latency/latency_compare.json`.

5.6 Search efficiency (learning curves)

Figure 9 plots champion score against outer-loop evaluations on seed `1902771841`. The top panel is the short FitCell hybrid under locked $R_{\text{blend}}+J$ (1,200 evaluations): raw R_{blend} clears Dual Cosine immediately and plateaus below the Noise2Noise gate (≈0.975). The bottom panel is the FitCell-to-plateau follow-up (750 evaluations): the same outer loop continues until residual plateau and clears the Noise2Noise gate (final R_{blend}≈0.9912; Section 5.3). The matched multi-seed R_{blend} bake-off is in Table 21. Live hybrid champions stay compact (≈24k±2k params across short-FitCell seeds), while mid-search graphs can swing from ≈15k to ≈370k.

Superseded whole-curve freezes. Whole-curve prolonged-R freezes (R≈0.991 vs DualCosine ≈0.825) are not current

Table 7: Ours vs Noise2Noise and Dual Cosine under R_{blend} (α=0.7). Holdout / multi-family: mean±std over five seeds (`20260719`–`20260723`). Δ = Ours − baseline.

Board	Ours	N2N	DualCos.	Δ_{N2N}	Δ_{DC}
Holdout	0.971±0.001	0.977±0.001	0.542±0.038	−0.006	+0.429
Multi-family	0.934±0.003	0.931±0.005	—	+0.003	—
Factory+FX	0.699	0.771	0.589	−0.072	+0.110
AKWF OA	0.679	0.731	0.424	−0.051	+0.255

results. Selection and published comparison under v10.1 use $R_{\text{blend}}+J$ with Noise2Noise as the primary neural baseline.

5.7 Which families stay hard

On the fitted DenoiseOpt 100,000-cycle bench, denoise score D is lowest for `nonlinear` ($D \approx 0.39$) and `combo` ($D \approx 0.40$), then `triple_mix` / `extreme_overlay` ($D \approx 0.51$–0.52), while `harmonic_fft` and `am_fm` sit near $D \approx 0.69$–0.70. Lit-combo champion family stress (prolonged residual) similarly ranks `open_wrap_bias` ($R \approx 0.83$) and `extreme_overlay` / `combo` ($R \approx 0.88$–0.89) below `harmonic_fft` ($R \approx 0.95$). Hard sounds recur.

Superseded 5k-gate freeze. At the reported 5k-gate search freeze ($\approx$ 6,922 clean iterations, $\approx$ 7.9 h on RTX 3090) the champion residual was $R \approx 0.99093$ with DualCosine ≈ 0.81658 on the CUDA search runner ($\Delta R \approx +0.174$). That freeze optimized whole-curve prolonged R. It is superseded by R_{blend} for v10.1 selection. Monitoring plots below come from that older freeze and should be read as search-trace artifacts, not as locked-protocol scores.

5.8 Inference latency at matched score

Across the prior high-prolonged-R fitted operators, we retimed CUDA forward passes under locked R_{blend} (batch 64). Live R_{blend} spreads more than the old prolonged band: the max-R_{blend} model is `champion_iter_000265` ($R_{\text{blend}}\approx 0.9769$, $\approx$ 6.53 ms/batch, $\approx$ 94k params). The compact historically favorite `iter_000235` remains lowest-latency among the top-5 distinct topologies ($R_{\text{blend}}\approx 0.9749$, $\approx$ 3.42 ms/batch, $\approx$ 37k params). Selection rule: maximize live R_{blend}, then minimize latency inside $R \geq R_{\max} - 5 \cdot 10^{-4}$.

5.9 Top-5 architectures

Table 8 ranks five distinct fitted bake graphs from the CUDA inference bench (inference data, batch 64), preferring live R_{blend} first and latency on ties, and skipping near-duplicate copies of the same topology signature.

Table 9 scores the classical (non-AI) bake set (no-bake, DualCosine, and `seam_fir3`) plus the locked v10.1 hybrid favorite on the frozen corpus of Section 4.1 (CUDA, batch 64, seed `20260719`) under R_{blend}. DualCosine reaches $R_{\text{blend}} \approx 0.541$ ($\approx$ 1.66 ms/batch). Active classical `seam_fir3` reaches 0.884 / $\approx$ 1.11 ms/batch (0.888 ± 0.009 over five seeds). No-bake (passthrough) scores 0.914 (0.917 ± 0.009). The hybrid favorite reaches 0.970 (0.971 ± 0.001 over five seeds) at $\approx$ 5.99 ms/batch: $\Delta R_{\text{blend}} \approx +0.056$ vs no-bake and

Table 8: Top-5 distinct fitted architectures on the CUDA inference bench (batch 64) under R_{blend} (α=0.7). R_{saved} is the checkpoint hint (often prolonged R); R_{live} is measured R_{blend}. *Max-R_{blend} (also sole member of the $R \geq R_{\max} - 5 \cdot 10^{-4}$ band here).

Tag	R_{saved}	R_{live}	ms/batch	params
`iter_000265`*	0.9906	0.9769	6.53	93 566
`iter_000140`	0.9906	0.9762	7.16	104 484
`iter_000133`	0.9908	0.9752	7.16	104 484
`iter_000229`	0.9909	0.9751	7.56	100 500
`iter_000235`	0.9910	0.9749	3.42	37 489

*Favorite under maximize-R_{blend} then latency. Compact `iter_000235` is still the latency floor in this top-5.

$\approx +0.429$ vs DualCosine. Noise2Noise on the multi-family matrix is Appendix Table 17. The four-board gate is Table 7.

Table 9: Canonical holdout under R_{blend} (α=0.7; seed `20260719`). Mean±std over five seeds. No-bake = unrepaired engine. Favorite = locked v10.1 hybrid champ.

Method	R_{blend}	R_{blend} (5-seed)	ms/batch
neural favorite	0.9697	0.9708 ± 0.0012	5.99
no-bake (passthrough)	0.9142	0.9166 ± 0.0090	1.06
`seam_fir3`	0.8842	0.8879 ± 0.0087	1.11
classic quadratic	0.5742	0.5727 ± 0.0318	2.07
DualCosine	0.5409	0.5420 ± 0.0342	1.66
cosine fade	0.5409	0.5420 ± 0.0342	1.87
crossfade	0.5310	0.5273 ± 0.0311	1.86
hann blend	0.5197	0.5214 ± 0.0337	1.05
soft wide fade	0.4828	0.4803 ± 0.0258	2.85
detrend	0.3437	0.3453 ± 0.0399	0.96
ensemble detrend+DC+FIR	0.3126	0.3100 ± 0.0421	2.00

5.10 Multi-family comparison

Does the locked hybrid favorite beat the classical non-AI board (no-bake, DualCosine, FIR, fades), fixed MLP/CNN baselines, *and* Noise2Noise across twenty generative families under R_{blend}? Appendix Table 17 reports mean±std over the 20-waveform multi-family set (Section 4.1) from the SOTA matrix. Appendix Figures 17–16 plot the same ranking. Primary reading is absolute R_{blend} rank: Ours 0.934, corrupt→corrupt N2N 0.931, seq CNN 0.920, no-bake 0.893, `seam_fir3` 0.868, DualCosine 0.514 (aligned with Table 7 multifamily). Fixed MLP-on-R and CNN/U-Net-lite-on-R (no outer NAS) sit between DualCosine and the neural rows. Paired Wilcoxon signed-rank (normal approx.) of favorite vs DualCosine on the 20 waveforms: W=0, $p\approx 8.9\times 10^{-5}$, mean ΔR_{blend}=+0.417 (bootstrap 95% CI $[0.398, 0.435]$). Vs no-bake the multifamily mean gap is $\approx +0.038$. Vs corrupt→corrupt N2N the mean-of-means gap is $\approx +0.003$ (mean win fraction 0.47 across five seeds).

5.11 Hard cliffs expose classical fade failure

Mild cliffs can leave no-bake R_{blend} looking high even when a wrap click remains. Appendix Table 19 therefore strat-

Table 10: Required seam-local secondary on top-10% wrap-jump tiles (n=410): edge RMSE (lower better). Frozen in the cliff-strata data.

Method	top10 edge RMSE
seq LSTM	0.011
N2N sibling-sup.	0.017
N2N corrupt→corrupt	0.020
neural favorite	0.027
seq CNN1D	0.035
no-bake	0.095
`seam_fir3`	0.164
DualCosine	0.990

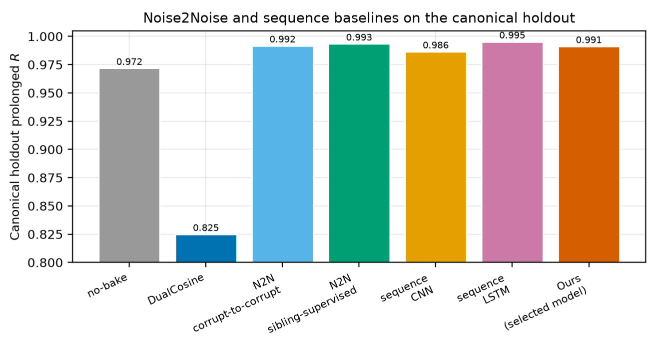

Figure 10: Canonical holdout R_{blend}: N2N and seq baselines vs Ours (locked hybrid). Train seeds `424242`/`424243` are disjoint from holdout `20260719`.

ifies a 4,096-tile holdout (seed `20260719`) by wrap-jump percentiles (p75≈1.387, p90≈1.832). On hard cliffs, Dual Cosine collapses (0.525 → 0.305), while the hybrid favorite holds (0.971 → 0.978) and Noise2Noise stays slightly ahead (0.976 / 0.983 on all / top-10%). The claim is tiled residual repair, not forcing wrap-jump to zero. Edge RMSE on the top-10% subset drops from no-bake 0.095 to favorite 0.027 (N2N 0.020, Dual Cosine 0.990; Table 10).

5.12 Noise2Noise and sequence baselines

How does DenoiseOpt compare to true Noise2Noise-style and lightweight seq restorers on the same generator under R_{blend}? Primary corrupt→corrupt N2N (train seed `424242`) reaches holdout $R_{\text{blend}} \approx 0.975$ (Table 7), above the hybrid favorite (0.970). On the 4,096-tile cliff draw, sibling-supervised N2N reaches 0.981 and seq LSTM 0.986 (Appendix Table 19). Holdout tiles never enter training. Classical DualCosine / `seam_fir3` remain in Appendix Table 19. Figure 10 plots the four-board R_{blend} comparison.

5.13 Factory exports and AKWF cycles

Does the wrap protocol transfer beyond synthetic sine cliffs onto true exports and third-party OA cycles? Table 11 scores ReelSynth-exported factory periods (primary realism) and Adventure Kid Waveforms (AKWF) CC0 single-cycle WAVs (secondary) after close-seam idealization and the same open-wrap cliff, under R_{blend}. Procedural stand-ins are demoted to

Table 11: Wavetable-native wrap protocol under R_{blend} (α=0.7). Close-seam ideal → open-wrap cliff (seed `20260719`). Primary = ReelSynth export. Secondary = AKWF CC0 files.

Method	Export	AKWF OA
`seam_fir3`	0.863	0.725
poly seam (d=3)	0.846	0.696
no-bake	0.845	0.702
endpoint-pin (mean)	0.792	0.745
N2N corrupt→corrupt	0.771	0.731
seq CNN1D	0.706	0.677
neural favorite	0.699	0.679
DualCosine	0.589	0.424
seq LSTM	0.628	0.579

tertiary smoke and no longer carry the primary realism claim. On exported factory periods the favorite (R_{blend}≈0.699) trails no-bake / FIR / poly / N2N while still beating Dual Cosine (0.589). On AKWF OA cycles the favorite (0.679) trails no-bake / N2N and beats Dual Cosine (0.424). Export gaps are listed in Section 9.

5.14 Polynomial baseline and endpoint pinning

Does a cheap classical poly fitter close the residual, and does endpoint pinning show the wrap-jump vs R_{blend} tradeoff? A seam-local degree-3 polynomial fitter (no learning) reaches canonical $R_{\text{blend}} \approx 0.917$ (slightly above no-bake 0.914) with wrap-jump essentially unchanged (polynomial results). SSM seq models are deferred. The LSTM remains the supervised seq ceiling. An endpoint-pin control (pinning results) drives wrap-jump to 0 on all / hard subsets while dropping R_{blend} (all 0.733, hard 0.597). Methods that zero wrap-jump are not the same objective as R_{blend} (Limitations).

5.15 Classical virtual-analog seam repairs

Do the cited VA antialias tools [1–3] beat DualCosine on R_{blend} when applied as cycle-local seam residuals?

What each operator changes. Figure 11 shows the family on one high-wrap sine+cliff tile (same seed and tile as Figure 2). **BLIT/BLEP** (Stilson/Smith lineage) treats the wrap as a value step. It subtracts a band-limited step residual (here a raised-cosine BLEP-family lobe scaled by the measured wrap jump) so the tiled click spectrum is softened near the seam. **PolyBLEP** (Nam et al.) is a cheap polynomial surrogate of that step residual with fractional-delay support near phase wrap. It edits only a few samples on each side of the discontinuity. **BLAMP** (Esqueda et al.) corrects a slope or corner jump (discontinuity in the first derivative), which matters for triangles, rectification, and hard clipping. On a pure value cliff the slope jump is small, so BLAMP is nearly a no-bake passthrough. **DualCosine** (our classical bake baseline) fades both ends of the stored period with raised-cosine windows. It does not insert a BLEP-style residual. It trades wrap energy

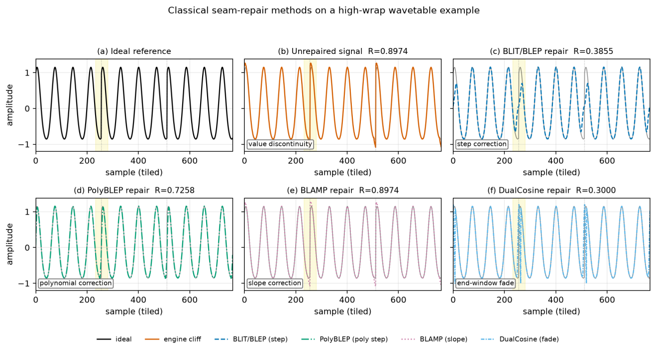

Figure 11: Classical VA seam techniques on one cracked sine+cliff tile (seed `20260719`, tile 46). (a) ideal; (b) cracked engine; (c) BLIT/BLEP; (d) PolyBLEP; (e) BLAMP; (f) DualCosine. Batch means: Table 12.

Table 12: Classical VA seam repairs under R_{blend} (α=0.7; seed `20260719`). Holdout R_{blend}, seam-local edge RMSE, cliff top-10% R_{blend}, multi-family mean.

Method	R	edge	top10	multi
BLAMP	0.914	0.080	0.922	0.893
no-bake	0.914	0.080	0.922	0.893
`seam_fir3`	0.884	0.112	0.869	0.868
PolyBLEP	0.795	0.205	0.690	0.784
BLIT/BLEP	0.632	0.365	0.351	0.612
DualCosine	0.541	0.506	0.305	0.514

Latency (CUDA): PolyBLEP ≈0.87 ms/batch, DualCosine ≈1.46, BLIT/BLEP ≈71, BLAMP ≈20.

for a wide end-fade distortion that R_{blend} penalizes on hard cliffs.

Scores. Table 12 reports batch means from the VA results (seed `20260719`, batch 64) under R_{blend}. Code is in the ReelSynth repository: PolyBLEP matches `osc::va::poly_blep`, BLIT/BLEP uses a raised-cosine step residual, and BLAMP applies a polyBLAMP lobe to the wrap slope jump. PolyBLEP reaches mean $R_{\text{blend}} \approx 0.795$ ($\Delta \approx +0.254$ vs DualCosine) but trails no-bake / `seam_fir3`. BLIT/BLEP nearly zeros wrap-jump (0.004), yet mean R_{blend} falls to 0.632 and hard cliffs collapse to 0.351. BLAMP stays near no-bake on this value-cliff geometry (mean 0.914). On the annotated severe tile in Figure 11, tile R_{blend} is lower still (PolyBLEP ≈ 0.726, BLIT/BLEP ≈ 0.385, Dual Cosine 0.300, BLAMP ≈ engine). We score these classical tools, not only cite them: they beat DualCosine on mean R_{blend} in places, and none matches the locked hybrid favorite on holdout R_{blend}.

5.16 When search helps vs classical methods

Per-cycle ΔR on export, AKWF, Rust `sound_bench`, and sine+cliff holdout (transfer results): favorite win-rate vs Dual Cosine is high on sine cliffs (1.00) and AKWF (0.88), and moderate on export (0.64) / Rust (0.75). Win-rate vs no-bake is high only on sine cliffs (0.92) and collapses on export (0.24) / Rust (0.20). Search value is highest on hard cliffs and when compact bake graphs are needed without engine→ideal labels at search time. Export / Rust gaps vs no-bake are listed in Section 9.

Table 13: Rust `sound_bench` transfer (n=20) under R_{blend} (α=0.7). Δ is vs Dual Cosine.

Method	R_{blend}	Δ vs DC
no-bake	0.846 ± 0.171	$+0.422$
`seam_fir3`	0.751 ± 0.157	$+0.327$
neural favorite	0.506 ± 0.236	$+0.082$
classic quadratic	0.474 ± 0.186	$+0.050$
Dual Cosine / cosine fade	0.424 ± 0.187	0

5.17 Transfer to Rust engine tiles

Table 13 scores the Python-trained favorite on 20 Rust-exported tiles under R_{blend}. The favorite beats Dual Cosine on mean ΔR_{blend}+0.082, but the signed-rank test is not significant at α=0.05 ($p\approx$0.30), and absolute score trails no-bake and `seam_fir3`. The primary SOTA claim remains the Python multi-family matrix.

5.18 Vibrato spectrograms

Do wrap-seam differences stay visible under slow vibrato, beyond static tiled periods? Appendix Figure 13 renders the same holdout cracked period, DualCosine bake, and refit **Ours (hybrid GA–PPO)** champion under shared sinusoidal vibrato (±3% depth at 5 Hz around 220 Hz). Panels (a–c) show magnitude spectrograms; (d–e) show DualCosine−no-bake and Ours−no-bake differences; (g) reports cycle-local prolonged R and a modulated-playback residual under the same pitch envelope. On the annotated high-wrap tile, Ours retains high cycle R and high modulated R, while DualCosine remains the weakest classical comparator. This is objective spectrogram / residual evidence only—not a formal listening study. Protocol: `bench_vibrato_spectrogram.py`; artifacts: vibrato folder.

5.19 Downloadable audio samples

As an audible companion to Figure 7, we release five fixed-pitch (A4, 3 s, 44.1 kHz mono) wavetable clips for no-bake, DualCosine, and Ours-healed periods from the matched-suite holdout (hear-sample WAVs, rebuild: `export_meta_hear_samples.py`). Appendix Figure 14 shows short waveform excerpts from those WAVs with per-clip absolute R.

Listening statistics (null study). Table 14 is the complete statistical statement for human listening in this paper. Assumptions: no IRB-recruited listeners; no forced-choice protocol executed; WAV pack is for optional informal inspection of the same tiles already scored by R_{blend}. Therefore preference mean, std, CI, and p-values are undefined (not omitted by accident). Subjective quality is *out of scope*; objective residual is the claim.

Table 14: Listening / subjective evaluation statistics for this manuscript (honest null study).

Quantity	Value
$N_{\text{listeners}}$	0
MOS / MUSHRA trials	0
Forced-choice A/B trials	0
Preference mean±std	n/a (undefined)
Significance tests on preference	none
Released informal WAVs	15 (5 tiles × 3 methods)
Primary quality metric	locked R_{blend}

Table 15: Wrap-heal transfer under locked R_{blend} (higher better; wrap-protocol stress test). Bearing boards (n=256 each). DenoiseOpt: `hybrid_lstm`, seed `1902771841`.

Method	CWRU	MFPT	Paderborn
Ours	0.960	0.943	0.937
SeamN2N	0.866	0.866	0.842
no-bake	0.845	0.866	0.838
endpoint-pin	0.505	0.545	0.567
`seam_fir3`	0.448	0.573	0.560
DualCosine	0.461	0.483	0.471

5.20 Factory and Adventure Kid waveform examples

Beyond the synthetic sine+cliff narrative, Appendix Figure 15 shows a compact gallery of true ReelSynth-exported factory periods (one mid-morph frame per bank) and Adventure Kid Waveforms (AKWF) CC0 single-cycle files already used by the wrap protocol. Scored means remain in Table 11 and the wavetable results. This panel is diversity evidence only; no new NAS was run for the gallery.

5.21 Transfer to bearings, ECG, and synthetic cycles

Setup reminder. Table 15 reports the primary metric **locked** R_{blend} (higher better; α=0.7 seam/body blend from Section 3.3) after training/searching DenoiseOpt (`hybrid_lstm`, seed `1902771841`) on each transfer board from Section 4.2 (n=256 periods, L=256). No PESQ/STOI/MOS appears on these boards: periods are non-speech wrap tiles, and the claim is wrap-protocol residual repair only. Columns are: CWRU/MFPT/Paderborn bearing revolutions, MIT-BIH/PTB-XL ECG beats, and synthetic CNC/PMU pilots. Numbers come from the released results_table.json (`primary_metric=r_blend`). Domain-trained Noise2Noise-style SeamN2N (corrupt→corrupt, 4000 Adam steps, train seed `424242`, holdout seed `20260719`) is scored on the same seven boards under the *same* R_{blend}. Author ECG Cycle-GAN and BeatDiff remain OOD wrap probes on ECG only (clinical restore ≠ wrap residual; status notes in Appendix B, Table 20). These rows do *not* transfer an audio-trained bake as fault/ECG diagnosis SOTA; they only stress the wrap geometry. Artifacts: signal_heal_transfer.

Takeaways. Under R_{blend}, Ours leads every classical-board column in Tables 15–16, including Dual Cosine and SeamN2N on every domain. SeamN2N beats Dual Cosine everywhere and sits near no-bake on several bearing / ECG columns, but trails Ours on ECG and synth CNC. On MIT-BIH and PTB-XL, endpoint-pin remains a strong classical ECG row above SeamN2N. Ours still leads both. Author Operational Cycle-GAN and BeatDiff collapse under OOD wrap scoring relative to Ours / SeamN2N / no-bake. We do not treat clinical morph metrics as wrap residual. Synthetic CNC shows the largest classical collapse with Ours near 0.991. Synth PMU is the strongest SeamN2N row (0.9589) but still below Ours (0.9919). Paderborn wrap board shows Ours 0.9369 vs SeamN2N 0.8421 and DualCosine 0.4710. Separately, we train Al Firdausi CNN as a 4-class fault classifier and do not rename classification as wrap residual (Appendix Table 20). Figure 12 summarizes the same boards as grouped bars. Latency timings in Appendix B remain wavetable-checkpoint zero-shot. Primary claim remains cycle-local wavetable seam restoration.

Table 16: Wrap-heal transfer under locked R_{blend} (continued): ECG and synthetic boards. BeatDiff / Cycle-GAN are author OOD wrap probes on ECG only.

Method	MIT-BIH	PTB-XL	CNC	PMU
Ours	0.975	0.945	0.991	0.992
SeamN2N	0.750	0.754	0.553	0.959
no-bake	0.750	0.762	0.450	0.936
endpoint-pin	0.865	0.870	0.410	0.944
`seam_fir3`	0.766	0.760	0.534	0.912
DualCosine	0.414	0.512	0.494	0.827
BeatDiff (OOD)	0.491	0.493	—	—
Cycle-GAN (OOD)	0.106	0.211	—	—

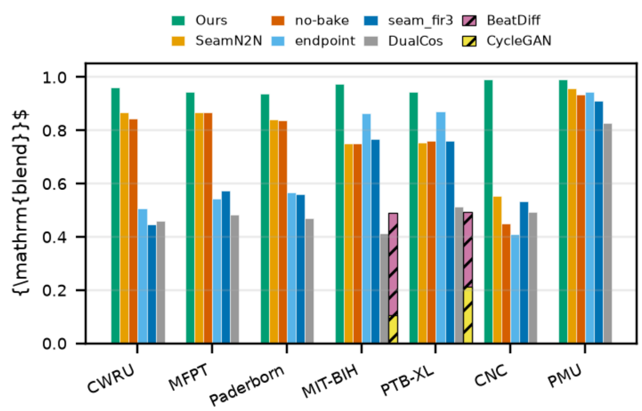

Figure 12: Grouped bars for Tables 15–16: R_{blend} by board and method (higher better). Core rows (Ours, SeamN2N, no-bake, endpoint-pin, `seam_fir3`, DualCosine) appear on all seven boards. Hatched bars are author OOD wrap probes on ECG (BeatDiff, Cycle-GAN) and the Paderborn Al Firdausi architecture-reuse wrap bake (not the published fault classifier, see Table 20). Missing cells are omitted rather than zero-filled.

5.22 Listening protocol (no formal listening study)

Downloadable WAVs (hear-sample WAVs) and vibrato spectrograms are released for informal file inspection (Sections 5.18–5.19). They are regenerated from the same public procedural / holdout tiles used in the residual tables—not private studio stems. Formal MOS / MUSHRA / forced-choice A/B was not run ($N_{\text{listeners}}=0$), so we report no preference rates or significance tests. Until a powered listening study runs, we make *no* perceptual quality claim; residual R_{blend} is the sole primary evidence. See Section 9.

6 Discussion

Main story under R_{blend}. Under the locked holdout protocol, Noise2Noise still leads the short-FitCell hybrid freeze (0.9767±0.0015 vs 0.9708±0.0014 over five seeds; Dual Cosine 0.542±0.038). The operative claim is narrower: compact searchable repair operators beat classical fades, stay competitive with Noise2Noise on multi-family means without a significant win (sign-test p=1.0), and hold on hard cliffs where Dual Cosine collapses (Table 7, Section 5.11). Favorite Θ on the v10.1 freeze is larger than SeamN2N (≈121k vs ≈53.5k) and still trails N2N on holdout R_{blend} (Table 4). Whole-curve numbers near R≈0.991 are superseded debug freezes, not current holdout results.

Which families stay hard. Across large procedural benches, wrap error is uneven. Families such as `nonlinear`, `combo`, `triple_mix`, and `extreme_overlay` dominate the residual budget; `harmonic_fft` and `am_fm` are easier. Mean scores hide that structure.

Search near the ceiling. Under near-ceiling residuals, further gains are small in absolute units and sensitive to reward shaping and branch weights. Learning curves (Figure 9) show Dual Cosine cleared early under short FitCell, then FitCell-to-plateau clearing the Noise2Noise gate (Figure 8). Matched-5k outer-loop bake-offs under locked $R_{\text{blend}}+J$ rank hybrid first across three search seeds (Appendix B, Table 21), which is why the follow-up spends GPU on hybrid alone (Section 5.3). Training candidates to residual plateau (rather than a 16–48-step FitCell cap) is what lets seed `1902771841` clear the Noise2Noise gate (R_{blend}≈0.9912). Future search work can target family-specialist graphs and cliff-conditioned FitCell schedules under the same outer loop.

Evidence beyond static tiles. Vibrato spectrograms and downloadable hear WAVs are informal file checks only (Sections 5.18–5.19; Table 14: $N_{\text{listeners}}=0$). They do *not* support perceptual claims of any kind—not preference, not "sounds better," not MOS/MUSHRA proxies. Sci/eng wrap transfer reports locked R_{blend} on seven boards as a wrap-protocol stress test (Section 5.21, Table 15), not audio→fault/ECG diagnosis SOTA. Scope limits are collected in Section 9.

7 Tooling and AI use

AI-assisted workflow. The following tools assisted preparation. The author checked claims, numbers, and citations against code and frozen JSON artifacts by hand (evidence alignment only—not formal verification, not a trusted computing base / TCB, and not machine-checked proofs).

- **Cursor** assisted software development (search runners, baselines, transfer benches, and figure scripts in the ReelSynth repository).
- **SciSpace** assisted literature discovery and paper triage.
- **NotebookLM** assisted citation checks (bibliography entries cross-checked against open-access sources and the manuscript text).
- **aixiv.science** assisted manuscript quality review (clarity, structure, and reviewer-style feedback used to guide revisions).

AI did not invent experimental numbers. Fitting, residual scoring, baselines, and hybrid GPU search are conventional Rust / PyTorch with deterministic logs.

Literature and reproducibility artifacts. Open literature was also gathered with local OpenAlex / Crossref / arXiv providers. Every bibliography entry has a downloadable OA PDF (literature PDFs). Audible demos: hear-sample WAVs.

8 Broader impact and reproducibility

Broader impact. This work targets cycle-local wavetable seam repair for virtual-analog / software-instrument workflows. Misuse potential is low: the method does not enhance speech for surveillance, does not remove watermarks from protected media, and operates only on short periodic tiles. Training, search, and reported tables use procedural sine+cliff generators and public corpora (Factory/AKWF exports; CWRU/MFPT/Paderborn/MIT-BIH/PTB-XL for wrap stress tests). No private studio stems enter the scored matrices. Released hear-pack WAVs are public regenerations of those holdout tiles for informal inspection, not private stems and not a listening study. The main resource cost is GPU time on a consumer RTX 3090 (order of hours for the 5k-gate freeze, Table 23; tens of hours for FitCell-to-convergence). Do not over-claim speech enhancement, mastering quality, or perceptual superiority from residual metrics alone.

Reproducibility. Holdout seeds: `20260719`–`20260723`. Search seeds: `1902771841`, `2026072701`, `2026072702`. Noise2Noise train seed: `424242`. Paper sources, frozen figure JSONs, and build scripts live at `https://github.com/reeldemo/denoiseopt-paper`; experimental search runners remain at `https://github.com/reeldemo/reelsynth`. CUDA benches need Python 3.12, PyTorch 2.6+CUDA 12.4, and NumPy/matplotlib. PESQ/STOI are omitted for non-speech cycles (domain mismatch), and MUSHRA was not run.

8.1 Funding

This research received no external funding.

8.2 Author contributions (CRediT)

Julian M. Kleber: Conceptualization, Methodology, Software, Investigation, Formal analysis, Data curation, Writing (original draft), Writing (review and editing), Visualization.

8.3 Data availability

Frozen evaluation artifacts (JSON matrices, search histories, fitted champions) are published at `https://github.com/reeldemo/denoiseopt-paper` under `figures/`. The PDF rebuilds with `build.ps1` in that repository. No private studio recordings were used.

8.4 Competing interests

The author declares no competing financial or non-financial interests. ReelSynth and denoise-opt-meta are MIT-licensed open-source projects maintained by the author; no commercial sponsor funded the reported experiments or manuscript. There is no conflict of interest with the Noise2Noise, VA, or NAS literature cited as baselines.

9 Limitations

These constraints bound claim strength; they are not optional honesty footnotes.

- **No formal listening study.** R_{blend} and J use a procedural no-cliff sibling for seams and an engine-identity prior for the mid-cycle body. They do not measure listening quality. No MOS, MUSHRA, or forced-choice A/B was run ($N_{\text{listeners}}=0$). Released WAVs and vibrato spectrograms are for informal file inspection only (Section 5.22); they are *not* a substitute for perceptual scores. Why listening was deferred: primary claims are objective cycle-local residual under a locked protocol for a DSP venue; a powered listener study needs IRB/recruitment budget outside this campaign. Generalization to “sounds better” is therefore *not* claimed. PESQ/STOI are excluded (non-speech tiles). We do not invent THD+N or MOS tables.

- **Metric lock and seed labeling.** Primary holdout claims use locked R_{blend} (Table 7, Section 3.3) as mean$\pm$std over five evaluation seeds. Matched-5k outer-loop means use three search seeds (Table 21). FitCell-to-plateau $R_{\text{blend}}\approx 0.9912$ is a single-seed freeze (Table 5), not a multi-seed mean. Whole-curve prolonged R near 0.991 is a superseded debug score and must not overwrite holdout or converge readings.

- **Synthetic / proxy boards.** Primary search and holdout use procedural sine+cliff generators, not private studio stems. Synthetic CNC/PMU transfer rows are procedural pilots when KIT / IEEE DataPort downloads are login-walled; they do *not* validate industrial toolpath or live-grid claims.

- **Domain Noise2Noise vs transfer depth.** Domain-trained Noise2Noise was scored on the seven wrap-protocol stress-test boards (Table 15). It trails Ours on every column under R_{blend}. Latency timings in Appendix B still use the wavetable checkpoint zero-shot. Cycle-GAN / BeatDiff / real KIT–PMU remain incomplete where noted (Table 20).

- **Domain shift on exports.** On Factory+FX / Rust tiles the meta favorite can trail no-bake or FIR while still beating Dual Cosine. We do not claim universal factory superiority.

- **Search geometry.** Search scores Python generator draws. The Python multi-family matrix is primary. Rust export is secondary. Per-trial width is capped. Demucs-/diffusion-scale stacks were screened out.

- **No wrap-closure theorems / no formal TCB.** Endpoint-pin can zero wrap-jump while residual falls. We claim tiled residual repair toward $r^{\star}$, not endpoint closure or search-convergence proofs. Tooling checks are manual evidence alignment against frozen JSON (Section 7), not a trusted computing base or machine-checked verification.

- **Transfer scope.** Sci/eng wrap transfer is a classical-board stress test of the wrap protocol, not deep CWRU/ECG SOTA (Table 20).

- **Hyperparameter probe.** The $\pm 50\%$ one-at-a-time probe (Appendix B) is not a full re-search of Table 2.

9.1 Future work

A follow-up can treat which sounds fail as the primary object: stratified residual heatmaps over procedural families, cliff-size conditioning, and family-specialist operators under the same outer loop. A powered formal listening study (MOS or forced-choice on high-wrap tiles) remains the clearest missing perceptual check and is required before any “sounds better” claim.

10 Conclusion

DenoiseOpt repairs wavetable wrap clicks with compact operators searched by hybrid GA+PPO(+PBT) under locked R_{blend} and latency-aware J. We reuse existing search tools; we do not introduce a new NAS algorithm. Noise2Noise is the neural gate. Dual Cosine is the classical fade.

On five holdout seeds, short-FitCell hybrid scores 0.971 ± 0.001: below Noise2Noise (0.977 ± 0.001), far above Dual Cosine (0.542 ± 0.038). Multi-family board means are statistically indistinguishable from Noise2Noise. Hard cliffs collapse classical fades while searched operators hold. A matched 5k bake-off ranks hybrid first (0.9748 ± 0.0018), so compute then goes to hybrid-only FitCell-to-plateau. Seed `1902771841` clears the Noise2Noise gate ($R_{\text{blend}}\approx 0.9912$). Transfer boards are wrap-protocol residual stress tests only. No formal listening study was run; the claim is objective residual.

Paper repository: `https://github.com/reeldemo/denoiseopt-paper`.

Acknowledgments

Authors. This preprint has one author, Julian M. Kleber. There are no co-authors.

Author contributions (CRediT). Julian M. Kleber performed all roles: Conceptualization; Methodology; Software; Investigation; Formal analysis; Data curation; Writing—original draft; Writing—review and editing; Visualization; Resources (local RTX 3090 workstation).

Named stress boards. We thank the maintainers of the wrap-protocol stress boards cited in Tables 15 and 3: CWRU bearing vibrations, MFPT bearing, Paderborn K001, MIT-BIH ECG, the PTB-XL ECG subset, Adventure Kid Waveforms (AKWF) CC0 single-cycle files, and ReelSynth Factory+FX export periods. Synthetic CNC and synthetic PMU rows are procedural pilots authored for this work when KIT and IEEE DataPort downloads were login-walled.

Software and literature. Thanks also to the Rust/Cargo, PyTorch, and egui communities used by ReelSynth. Open literature was gathered with local OpenAlex / Crossref / arXiv providers (Section 7). Any remaining errors are the author's.

References

[1] T. Stilson and J. O. Smith. Alias-free digital synthesis of classic analog waveforms. In *Proc. ICMC*, 1996. `https://ccrma.stanford.edu/~stilti/papers/blit.pdf`.

[2] F. Esqueda, V. Välimäki, and S. Bilbao. Rounding corners with BLAMP. In *Proc. DAFx*, 2016. `https://www.dafx.de/paper-archive/2016/dafxpapers/18-DAFx-16_paper_33-PN.pdf`.

[3] J. Nam, V. Välimäki, J. S. Abel, and J. O. Smith. Efficient antialiasing oscillator algorithms using low-order fractional delay filters. In *Proc. DAFx*, 2009. `https://ccrma.stanford.edu/~juhan/pubs/jnam-dafx09.pdf`.

[4] J. Lehtinen et al. Noise2Noise: Learning image restoration without clean data. In *ICML*, 2018. arXiv:1803.04189.

[5] M. M. Kashyap, A. Tambwekar, K. Manohara, and S. Natarajan. Speech denoising without clean training data: A Noise2Noise approach. In *INTERSPEECH*, 2021. `https://www.isca-archive.org/interspeech_2021/kashyap21_interspeech.pdf`.

[6] J. Engel, L. Hantrakul, C. Gu, and A. Roberts. DDSP: Differentiable digital signal processing. In *ICLR*, 2020. arXiv:2001.04643.

[7] H. Purwins et al. Deep learning for audio signal processing. *IEEE J. Sel. Topics Signal Process.*, 2019. arXiv:1905.00078.

[8] D. Stoller, S. Ewert, and S. Dixon. Wave-U-Net: A multi-scale neural network for end-to-end audio source separation. In *ISMIR*, 2018. arXiv:1806.03185.

[9] C. Macartney and T. Weyde. Improved speech enhancement with the Wave-U-Net. arXiv:1811.11307, 2018.

[10] A. Jansson et al. Singing voice separation with deep U-Net convolutional networks. In *ISMIR*, 2017. `https://openaccess.city.ac.uk/id/eprint/19289/1/7bb8d1600fba70dd79408775cd0c37a4ff62.pdf`.

[11] Y. Luo and N. Mesgarani. Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation. *IEEE/ACM TASLP*, 2019. arXiv:1809.07454.

[12] Y. Luo, Z. Chen, and T. Yoshioka. Dual-path RNN: Efficient long sequence modeling for time-domain speech separation. In *ICASSP*, 2020. arXiv:1910.06379.

[13] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. arXiv:1505.04597.

[14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. arXiv:1512.03385.

[15] A. Vaswani et al. Attention is all you need. In *NeurIPS*, 2017. arXiv:1706.03762.

[16] T. Elsken, J. H. Metzen, and F. Hutter. Neural architecture search: A survey. *JMLR*, 2019. arXiv:1808.05377.

[17] B. Zoph and Q. V. Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017. arXiv:1611.01578.

[18] H. Pham et al. Efficient neural architecture search via parameter sharing. In *ICML*, 2018. arXiv:1802.03268.

[19] J. Schulman et al. Proximal policy optimization algorithms. arXiv:1707.06347, 2017.

[20] E. Real et al. Large-scale evolution of image classifiers. In *ICML*, 2017. arXiv:1703.01041.

[21] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le. Regularized evolution for image classifier architecture search. In *AAAI*, 2019. arXiv:1802.01548.

[22] M. Jaderberg et al. Population based training of neural networks. arXiv:1711.09846, 2017.

[23] S. Khadka and K. Tumer. Evolution-guided policy gradient in reinforcement learning. In *NeurIPS*, 2018. arXiv:1805.07917.

[24] N. Hansen. The CMA evolution strategy: A tutorial. arXiv:1604.00772, 2016.

[25] N. Shazeer et al. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017. arXiv:1701.06538.

[26] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *NeurIPS*, 2012. arXiv:1206.2944.

[27] M. Tan et al. MnasNet: Platform-aware neural architecture search for mobile. In *CVPR*, 2019. arXiv:1807.11626.

[28] C. Liu et al. Progressive neural architecture search. In *ECCV*, 2018. arXiv:1712.00559.

[29] S. Wisdom et al. Unsupervised sound separation using mixture invariant training. In *NeurIPS*, 2020. arXiv:2006.12701.

[30] C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong. Attention is all you need in speech separation. In *ICASSP*, 2021. arXiv:2010.13154.

[31] Y. Gong, Y.-A. Chung, and J. Glass. AST: Audio spectrogram transformer. In *INTERSPEECH*, 2021. arXiv:2104.01778.

[32] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. arXiv:1703.03400. *(Screened.)*

[33] H. Liu, K. Simonyan, and Y. Yang. DARTS: Differentiable architecture search. In *ICLR*, 2019. arXiv:1806.09055. *(Screened.)*

[34] A. Défossez et al. Demucs: Deep extractor for music sources. arXiv:1911.13254, 2019. *(Screened.)*

[35] A. Défossez, G. Synnaeve, and Y. Adi. Real time speech enhancement in the waveform domain. In *INTERSPEECH*, 2020. *(Screened.)* arXiv:2006.12847.

[36] A. Défossez. Hybrid spectrogram and waveform source separation. arXiv:2111.03600, 2021. *(Screened.)*

[37] Z. Kong et al. DiffWave: A versatile diffusion model for audio synthesis. In *ICLR*, 2021. arXiv:2009.00713. *(Screened.)*

[38] S. Pascual, A. Bonafonte, and J. Serrà. SEGAN: Speech enhancement generative adversarial network. In *INTERSPEECH*, 2017. *(Screened.)* arXiv:1703.09452.

A Algorithms

Companion pseudocode for the residual score, classical controls, FitCell, and hybrid outer-loop branches. Canonical narrative lives in Section 3. These listings mirror the pseudocode notes and the CUDA search runner (`overnight_gpu_rl_arch.py`).

How the pieces connect. Algorithm 1 is the debug whole-curve score; primary selection uses the blended form in Section 3.3 (seam vs $r^\star$, body vs engine). Algorithms 2–3 are classical / passthrough controls scored under the same residual. Algorithm 4 is an optional soft gate inside a candidate repair graph. Algorithm 5 (**FitCell**) is the inner trainer: Adam on $1 - R_{\text{blend}}$ until plateau or hard cap $T_{\max}$ (converge campaign) or a short step budget (D1). Algorithms 6–8 are the outer-loop branches that propose new graphs / hyperparameters; the hybrid runner interleaves them each iteration and keeps elites by J. Ablations that drop GA or PPO (Table 22, right column) still call Algorithm 5 but disable the corresponding outer branch.

Algorithm 1 ResidualScore

Require: Ideal cycle $r^\star$, baked cycle y, periods N

$I \leftarrow \text{Tile}(r^\star, N)$
$Y \leftarrow \text{Tile}(y, N)$
$r_{\text{rms}} \leftarrow \text{RMS}(Y - I)$
$i_{\text{rms}} \leftarrow \max(\text{RMS}(I), \varepsilon)$
return $\text{clamp}(1 - r_{\text{rms}}/i_{\text{rms}}, 0, 1)$

Algorithm 2 DualCosine baseline

Require: Engine cycle x

Fade both ends of x with raised-cosine windows of width W
return $\text{ResidualScore}(r^\star, x_{\text{faded}}, N)$

Algorithm 3 No-bake (passthrough) control

Require: Engine cycle x (open-wrap cliff present)

return $\text{ResidualScore}(r^\star, x, N)$ ▷ no operator: $x \mapsto x$

Algorithm 4 MoE soft gating over bake experts

Require: Cycle x, expert operators $\{E_k\}_{k=1}^K$, gate logits $g \in \mathbb{R}^K$

$w \leftarrow \text{softmax}(g)$
$y \leftarrow \sum_{k=1}^K w_k E_k(x)$
return y

Algorithm 5 FitCell (inner loop; converge protocol)

Require: Architecture cfg; hard cap $T_{\max}$; patience P; relative ε; batch B

cell $\leftarrow$ SeamCell(cfg)
opt $\leftarrow$ Adam(cell.params)
prev $\leftarrow$ null; stagnant $\leftarrow 0$
for $t = 1, \ldots, T_{\max}$ **do**
  ideal, eng $\leftarrow$ MakeBatch(B)
  out $\leftarrow$ ApplyOps(eng, cell, cfg.ops)
  $R \leftarrow$ mean ResidualScoreBlend(ideal, eng, out)
  loss $\leftarrow 1 - R$
  if finite(loss) **then**
    AdamStep(opt, loss)
  end if
  if prev $\neq$ null and $|\text{prev} - R| / \max(|\text{prev}|, \varepsilon_{\text{num}}) < \varepsilon$ **then**
    stagnant $\leftarrow$ stagnant $+1$
    if stagnant $\geq P$ **then**
      break ▷ converged
    end if
  else
    stagnant $\leftarrow 0$
  end if
  prev $\leftarrow R$
end for
return cell, R, converged, t

Algorithm 6 GA tournament step

Require: Population P, tournament size k, crossover rate p_c, mutation rate p_m

Sample k genomes; parents $\leftarrow$ top-2 by R
child $\leftarrow$ Crossover(parents) with prob. p_c, else clone elite
Mutate(child) with rate p_m (ops, depth, operator type, θ)
Fit and evaluate child; insert into P (replace worst if full)
return updated P

Algorithm 7 PPO architecture update

Require: Actor–critic π_ϕ, clip ε, experience buffer of mutate actions

Sample action a editing ops / hyperparameters from π_ϕ
Fit proposed cfg; observe centered advantage $\rho = R - R_{\text{DualCosine}}$ ▷ objective remains maximize R
Advantage $\hat{A}$ from critic; ratio $r_\phi = \pi_\phi(a)/\pi_{\phi_{\text{old}}}(a)$
Update ϕ with clipped surrogate $\min(r_\phi \hat{A}, \text{clip}(r_\phi, 1-\varepsilon, 1+\varepsilon)\hat{A})$
return updated π_ϕ

Algorithm 8 PBT exploit–mutate

Require: Population P, exploit quantile q, mutate scale σ

Rank P by R; for each bottom member, copy hyperparameters from a top-q elite
Perturb copied hyperparameters by multiplicative noise of scale σ
Refit briefly; write back into P
return updated P

Algorithm 9 HybridMetaSearch

Require: Iteration budget T_{out}, population size n

$P \leftarrow \text{InitPopulation}(n)$
$\pi \leftarrow \text{ActorCritic}()$
champ $\leftarrow$ null
boredom $\leftarrow 0$
for $it = 1, \dots, T_{\text{out}}$ **do**
  branch $\leftarrow \text{Rotate}(\text{ppo}, \text{ga}, \text{pbt}, \text{nas}, \text{combo})$
  cfg, hp $\leftarrow$ Propose(branch, P, π)
  cell, $R_{\text{fit}} \leftarrow$ FitCell(cfg, hp.fit_steps)
  $R_{\text{eval}} \leftarrow$ EvalCell(cell)
  $R \leftarrow \frac{1}{2}R_{\text{fit}} + \frac{1}{2}R_{\text{eval}}$ (+ optional depth/MoE bonus)
  UpdatePopulation(P, cfg, R)
  PPOUpdate(π, centered $\rho = R - R_{\text{DualCosine}}$)
  if $R >$ champ.R **then**
    champ $\leftarrow$ (cfg, cell, R)
    boredom $\leftarrow 0$
  else
    boredom $\leftarrow$ boredom $+1$
  end if
  if boredom ≥ 1000 **then**
    PlateauAdapt
    boredom $\leftarrow 0$
  end if
end for
return champ

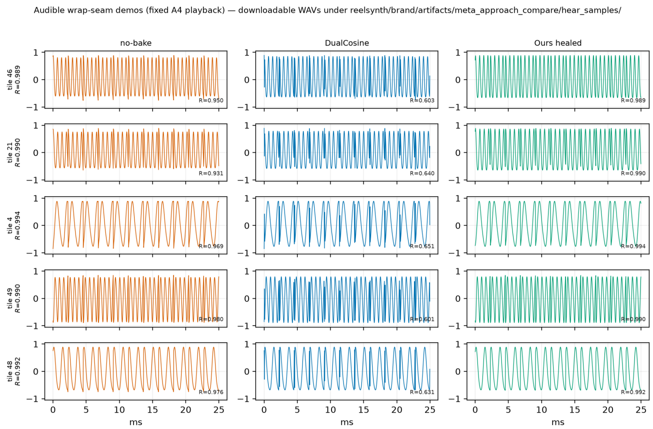

Figure 14: Hear-sample panel from the online WAV pack (hear-sample WAVs, five holdout tiles, no-bake / DualCosine / Ours). Waveforms show the first ≈ 25 ms of each clip. R annotations are cycle-local prolonged residuals vs the ideal sibling. Not a formal listening study.

B Supplementary tables and figures

Material moved here for the conference-length reading path: vibrato / hear / diversity panels, multi-family / cliff boards, matched outer-loop bake-off, branch/compute freezes, hyper-parameter probe, transfer latency / status notes, and search-monitor plots. Primary claims remain in Sections 5–5.21.

B.1 Vibrato, hear-sample, and diversity panels

Figures 13–15 support Sections 5.18–5.20. They are large multi-panel artifacts, so they sit here rather than opening empty main-text columns.

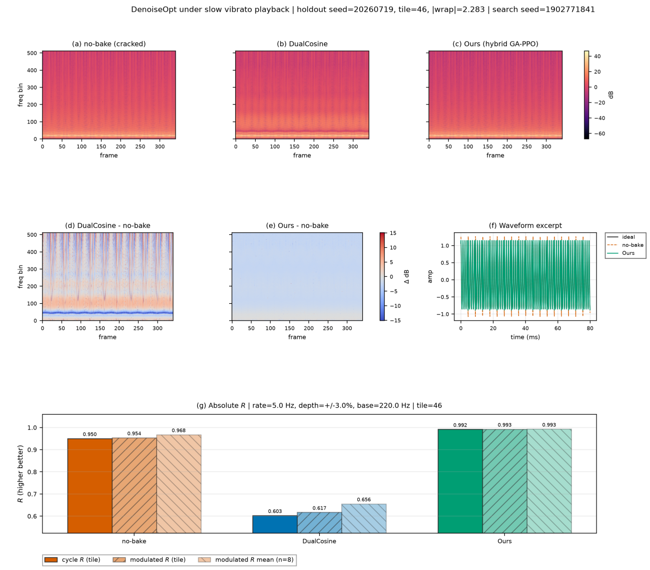

Figure 13: Slow-vibrato spectrogram eval on a high-wrap holdout tile (seed `20260719`). (a–c) Magnitude spectrograms for no-bake, DualCosine, and Ours. (d–e) Spectrogram differences vs no-bake. (f) Short waveform excerpt. (g) Absolute R: cycle-local prolonged residual, modulated-playback residual on the same tile, and mean modulated R over top-wrap tiles. Accessibility: Okabe–Ito hues. Difference maps use diverging coolwarm. Formal human A/B scores are out of scope.

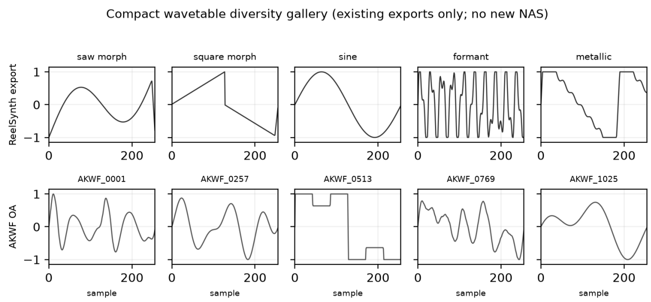

Figure 15: Compact wavetable diversity gallery from existing exports (one mid-morph dry frame per unique ReelSynth factory bank, top, and matched AKWF CC0 cycles, bottom). No new NAS for this panel. Quantitative wrap scores: Table 11.

B.2 Multi-family ranking plots

Figure 16 ranks methods on the multi-family board; Figure 17 breaks the same board down by generative family.

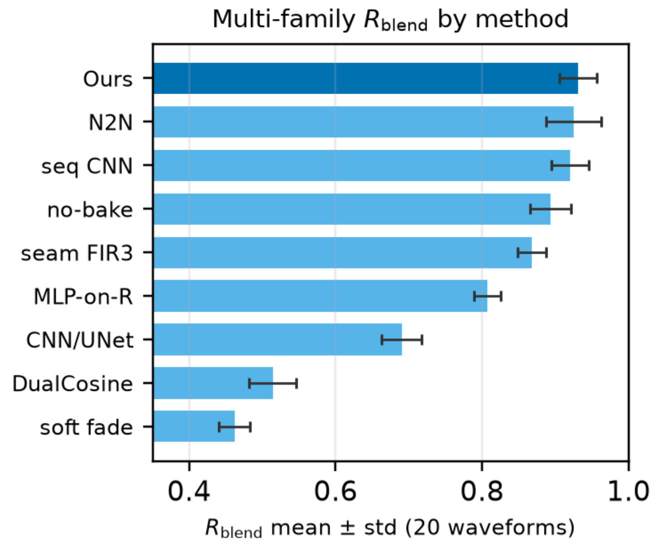

Figure 16: Multi-family mean R_{blend} by method (same 20-waveform board as Table 17).

Table 19: Cliff strata on 4,096 holdout tiles under R_{blend} (α=0.7). Strata by engine wrap-jump (n: all 4096, top25 1024, top10 410). $\pm$ is tile std on the all stratum.

Method	all	top25%	top10%
seq LSTM	0.986 ± 0.009	0.989	0.990
N2N sibling-sup.	0.981 ± 0.013	0.985	0.985
N2N corrupt→corrupt	0.976 ± 0.015	0.981	0.983
neural favorite	0.971 ± 0.020	0.978	0.978
seq CNN1D	0.960 ± 0.028	0.969	0.970
no-bake	0.913 ± 0.077	0.929	0.922
`seam_fir3`	0.886 ± 0.067	0.870	0.869
DualCosine	0.525 ± 0.239	0.318	0.305

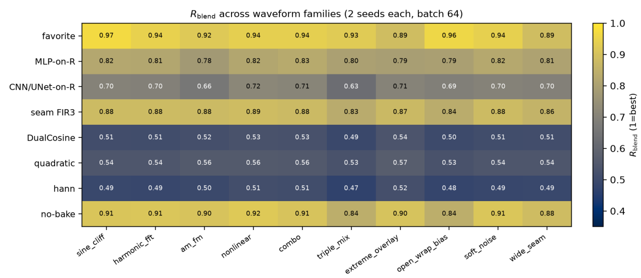

Figure 17: Colorblind-safe (`cividis`) heatmap of mean R_{blend} for selected methods $\times$ generative families (2 seeds each). Higher is better.

B.3 Multi-family and cliff boards

Full multi-family ranking (Table 17) and hard-cliff strata (Table 19) for the Results narrative in Sections 5.10–5.11.

Table 17: Multi-family board (20 waveforms) under R_{blend} (α=0.7). Core rank columns; SNR/SDR/jump in Table 18. ΔR and ms vs Dual Cosine. N2N/seq seeds `424242`/`424243`.

Method	R_{blend}	ΔR	ms
neural favorite	0.931 ± 0.026	+0.417	4.50
N2N corrupt→corrupt	0.925 ± 0.038	+0.410	0.74
N2N sibling-sup.	0.924 ± 0.035	+0.410	0.74
seq CNN1D	0.920 ± 0.026	+0.405	0.53
no-bake	0.893 ± 0.028	+0.379	0.00
seq LSTM	0.891 ± 0.075	+0.376	1.53
`seam_fir3`	0.868 ± 0.019	+0.353	0.17
MLP-on-R	0.807 ± 0.018	+0.292	1.76
CNN/U-Net-on-R	0.690 ± 0.027	+0.175	3.08
DualCosine	0.514 ± 0.032	0	0.98
soft wide fade	0.462 ± 0.021	−0.053	1.87

Table 18: Multi-family secondary metrics (same board as Table 17).

Method	SNR (dB)	SDR (dB)	wrap-jump
neural favorite	34.6 ± 3.2	34.7 ± 3.2	0.89 ± 0.14
N2N corrupt→corrupt	33.8 ± 4.0	33.9 ± 4.0	0.89 ± 0.15
N2N sibling-sup.	33.9 ± 4.1	34.0 ± 4.1	0.91 ± 0.15
seq CNN1D	32.2 ± 3.5	32.4 ± 3.5	0.92 ± 0.15
no-bake	30.9 ± 1.9	31.0 ± 1.9	0.91 ± 0.15
seq LSTM	32.2 ± 5.8	32.4 ± 5.8	0.86 ± 0.15
`seam_fir3`	28.1 ± 1.2	28.1 ± 1.2	0.46 ± 0.08
DualCosine	16.9 ± 1.2	16.8 ± 1.2	0.32 ± 0.05

B.4 Transfer wrap-protocol status notes

Table 20: Author-method and data-cut status under this wrap-protocol stress test (no invented scores).

Scope	Note
SeamN2N (ours, N2N-style)	Scored on all seven Table 15 boards under R_{blend}. Not Lehtinen code
Cycle-GAN (ECG, author)	OOD wrap probe in Table 15 (MIT-BIH 0.1063, PTB-XL 0.2105). Clinical restore $\neq$ wrap residual
BeatDiff (Bedin NeurIPS 2024)	OOD wrap probe in Table 15 (MIT-BIH 0.4907, PTB-XL 0.4933). Drive prior (not HF). Clinical morph prior $\neq$ wrap residual
Paderborn KAt deep (Al Firdausi CNN)	Trained CNN_1D_2L from scratch (not frozen `model.pth`) on bearings K001/KA04/KI04/KB23. File-holdout seed `20260726`, train seed `42`, 20 epochs. 4-class holdout acc 0.8590 (n=2560 segs, NORMAL/IR/OR/OR+IR). Classifier wrap residual N/A. Arch-reuse wrap bake listed in Table 15[†]. Wrap board: Ours 0.9369 / SeamN2N 0.8421 / DualCosine 0.4710
Expanded PTB-XL	records500 lead-I expanded (∼2k + `_hr` records, board n=256, Ours 0.9451). Multi-lead still out of scope
Real KIT CNC / IEEE PMU	Synthetic pilots only (download walls)
Formal MOS / MUSHRA	Not collected

B.5 Matched-budget outer-loop comparison

Under a 5,000 evaluation budget with LSTM and xLSTM in the searchable bake vocabulary, Table 21 compares Random NAS, Cont. CMA-ES, Arch REINFORCE, Aging evolution, TPE Bayes NAS, and our hybrid GA+PPO approach (Ours). The table reports results from a matched Deferred D1 re-search under locked $R_{\text{blend}}+J$ (mean $\pm$ std over three search seeds; artifacts `meta_approach_compare_v13_rblend/`). Our hybrid GA+PPO leads with an R_{blend} of 0.9748$\pm$0.0018, taking approximately 8.4 hours on average. Aging evolution ranks second at 0.9635$\pm$0.0049, while Arch REINFORCE shows the highest variance with 0.9599$\pm$0.0116. Only seed `1902771841` of our hybrid approach clears the frozen Noise2Noise gate under the short-FitCell D1 protocol, achieving an R_{blend} of approximately 0.9766 compared to the gate's approximate value of 0.9750. Despite mid-search graphs ranging from about 15k to 370k parameters, our hybrid approach maintains compact models (≈24k$\pm$2k params across seeds). Since our hybrid method dominated this matched bake-off, the subsequent campaign focuses solely on running hybrid-only FitCell-to-convergence (Section 5.3, Table 5). Note that Noise2Noise is not an outer-loop method in this context (see Tables 7 and 17). Learning-curve and bar figures use trajectories from seed `1902771841` under the same locked objective.

Table 21: Matched 5,000-evaluation outer-loop comparison under locked $R_{\text{blend}}+J$ (Deferred D1 complete). Mean$\pm$std over three search seeds (`1902771841`, `2026072701`, `2026072702`). ΔR uses seed-`1902771841` Dual Cosine (≈0.504). Wall-h is mean CUDA wall per approach. LSTM?/xLSTM?: whether any seed's champion included that block.

Method	R_{blend}	ΔR vs DC	h	LSTM	xLSTM
Random NAS	0.9367$\pm$0.0032	+0.4327	4.7	no	no
Cont. CMA-ES	0.9565$\pm$0.0016	+0.4525	1.8	no	no
Arch REINFORCE	0.9599$\pm$0.0116	+0.4559	5.8	yes	no
Aging evolution	0.9635$\pm$0.0049	+0.4595	9.0	yes	no
TPE Bayes NAS	0.9451$\pm$0.0019	+0.4411	3.8	no	no
Ours (hybrid GA–PPO)	0.9748$\pm$0.0018	+0.4708	8.4	no	no

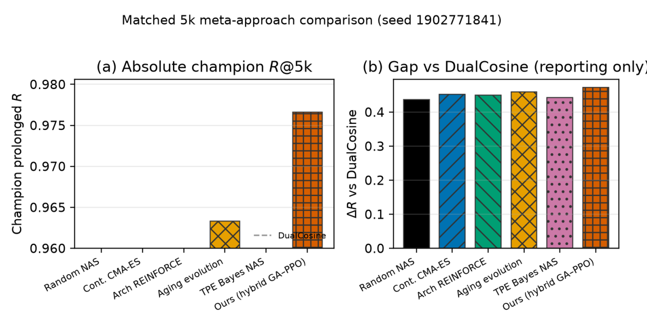

Figure 18: Matched 5,000-evaluation search methods under locked $R_{\text{blend}}+J$ (seed `1902771841`). Table 21 reports mean$\pm$std across three search seeds. Dashed line: Dual Cosine on the same runner.

Table 22: Ablations (interpret narrowly). Left: best prolonged R by branch label in one hybrid history (final iteration ≈6,922; superseded search metric; not separate fits). Right: isolated 150-iteration `ablate-*` final champions re-scored under R_{blend} (α=0.7; seed `1902771841`). Full R_{blend} branch re-search remains deferred; use Table 21 for matched outer-loop comparison.

Config	Prolonged R freeze	150-it R_{blend}
Full hybrid	0.99093	0.92710
GA-only	0.99016	0.94168
PBT (search branch)	0.99012	n/a
PPO-only	0.98924	0.94405
GA+PPO	n/a	0.94478
Combined branch	0.98854	n/a
NAS branch	0.98677	n/a
Dual Cosine (runner)	0.81658	≈0.539

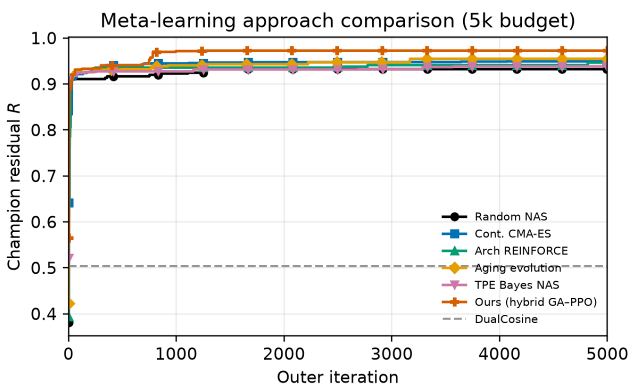

Figure 19: Champion R_{blend} over a matched 5,000-evaluation budget (Deferred D1 complete; seed `1902771841`). Dashed line: Dual Cosine.

B.6 Ablations and compute

Table 22 answers two narrow questions, not a full factorial re-search of Table 21. **Left column:** within one recorded hybrid history under the superseded prolonged-R search metric, what was the best freeze tagged by outer-loop branch (PPO/GA/PBT/...) at the 5k gate? These are branch labels in a single run, not separately fitted champions, so they mainly show that no single branch monopolizes the trajectory. **Right column:** isolated 150-iteration `ablate-*` restarts (seed `1902771841`) whose final cells are re-scored under locked R_{blend}. They probe whether dropping GA or PPO collapses residual under a short budget; they are *not* matched-5k competitors and must not overwrite Table 21.

How to read Table 22. Left-column values are close to each other (≈0.987–0.991 prolonged R): in that single history, branch tags do not identify a unique winner, so we do not treat "GA-only freeze" as proof that GA alone suffices under locked R_{blend}. Right-column short restarts under R_{blend} stay far below the matched-5k hybrid mean (0.9748) and the converge seed (0.9912) because the budget is only 150 iterations with short FitCell—they are sensitivity checks, not leaderboard rows. The matched comparison that drives method choice remains Table 21. Table 23 summarizes campaign compute (historical prolonged-R overnight).

Table 25: Transfer-domain inference latency (ms/batch). Batch=64, CUDA.

Method	CWRU	MFPT	MIT-BIH	PTB-XL	CNC	PMU
no-bake	0.26	0.26	0.27	0.27	0.28	0.26
Dual Cosine	1.10	1.11	1.12	1.14	1.15	1.14
Ours (hybrid)	3.92	6.61	2.87	1.67	1.69	1.67
N2N (WT to domain)	0.95	0.99	0.98	0.97	0.98	0.97

Table 23: Compute budget for the reported 5,000-evaluation search run (RTX 3090, search seed `1902771841`; prolonged-R objective). Final champion R below is the historical prolonged freeze, not R_{blend}.

Item	Value
GPU	NVIDIA GeForce RTX 3090
Elapsed wall time	≈7.92 h
Clean iterations / arch evals	≈6,922
Population size	12
Final champion prolonged R	0.99093
Dual Cosine baseline (runner, prolonged)	0.81658
Peak VRAM (replay)	≈3.29 GiB

B.7 Hyperparameter sensitivity

One-at-a-time ±50% probe of Table 2 knobs at 500 outer iterations (seed `1902771841`). Not a full 5k re-search. Default champion $R \approx 0.9888$. The largest $|\Delta R|$ among OAT arms is ≈ 0.0123 (50% lower learning rate). All 11/11 arms finished under this budget.

Table 24: One-at-a-time ±50% hyperparameter sensitivity probe (500 outer iterations, seed `1902771841`, sine+cliff FitCell). Champion absolute R is compared with the Table 2 default. This is sensitivity evidence only, not a full 5,000-iteration re-search for each parameter.

Parameter change	Champ. R	ΔR	Iter.
Default	0.98882	+0.00000	500
Population size −50%	0.98136	-0.00746	500
Population size +50%	0.98073	-0.00809	500
PPO clip −50%	0.98159	-0.00723	500
PPO clip +50%	0.99041	+0.00159	500
Learning rate −50%	0.97654	-0.01228	500
Learning rate +50%	0.98138	-0.00744	500
GA mutation rate −50%	0.98334	-0.00548	500
GA mutation rate +50%	0.99002	+0.00120	500
PPO entropy −50%	0.98912	+0.00030	500
PPO entropy +50%	0.99116	+0.00234	500

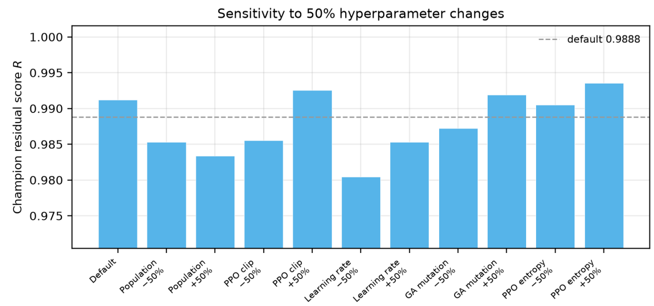

Figure 20: One-at-a-time ±50% HP sensitivity of hybrid champion R at 500 iterations. Blue bar: Table 2 default.

B.8 Transfer inference latency

Table 25 and Figure 21 pair with the prolonged-R transfer board (Table 15). N2N latency uses the wavetable checkpoint zero-shot. Domain-trained N2N quality scores are in Table 15.

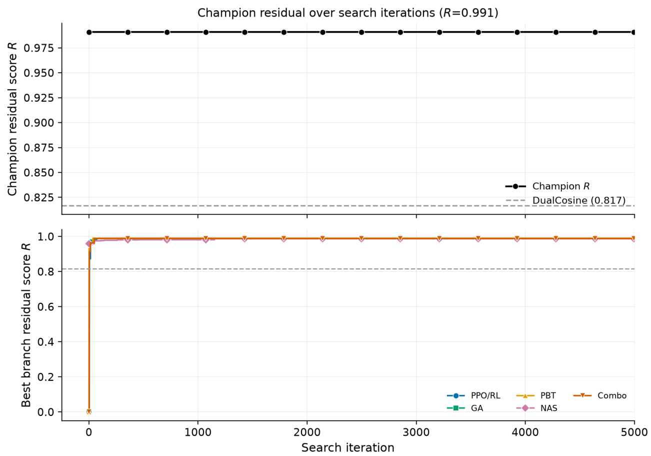

Figure 24: Search progress: champion R (top) and branch bests (bottom).

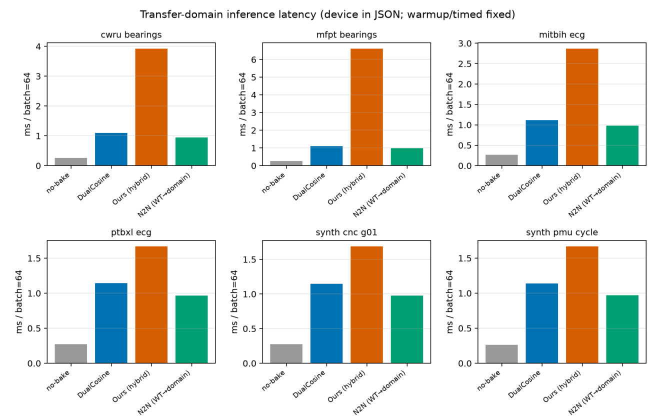

Figure 21: Transfer-domain inference latency (ms/batch). N2N is wavetable-trained zero-shot when available.

B.9 Search-monitor figures

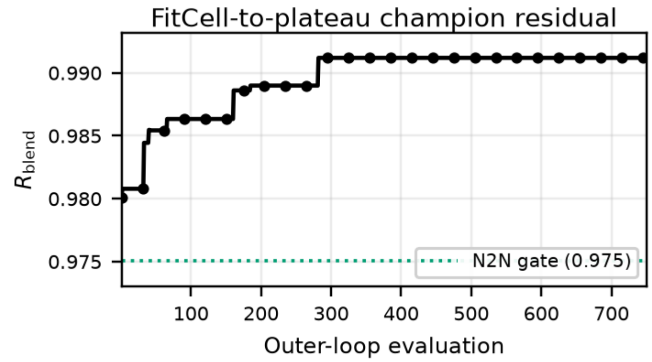

Figure 22: FitCell-to-plateau champion R_{blend} vs. outer-loop evaluation (seed `1902771841`; y-axis zoomed to the high-R regime with Noise2Noise gate marked).

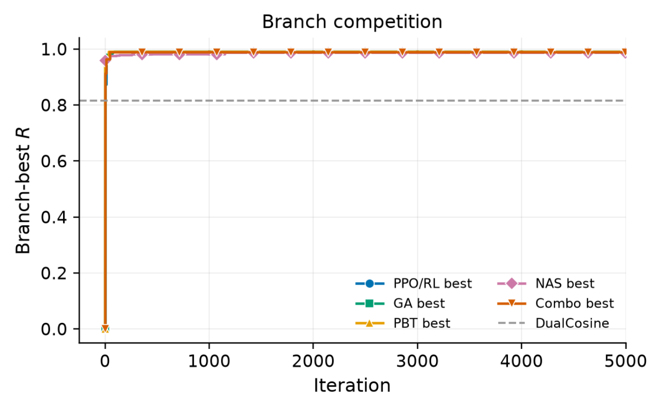

Figure 23: Best residual by search branch (older overnight monitor freeze; search-trace artifact).

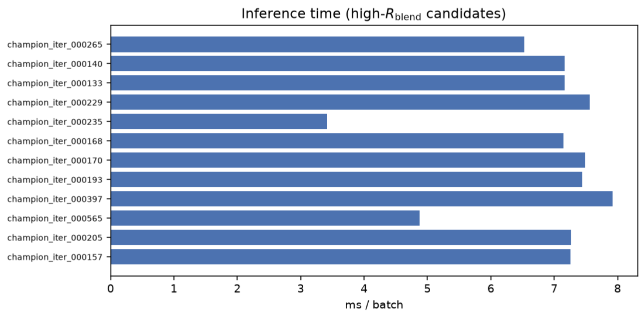

Figure 28: Inference latency for high-R fitted models (CUDA, batch 64).

Figure 25: Per-trial residual by branch with champion trajectory overlaid.

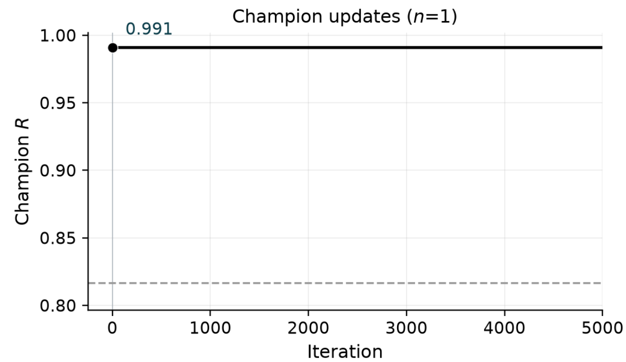

Figure 26: Champion updates (older overnight monitor freeze; search-trace artifact).

B.10 Inference and classical diagnostic plots

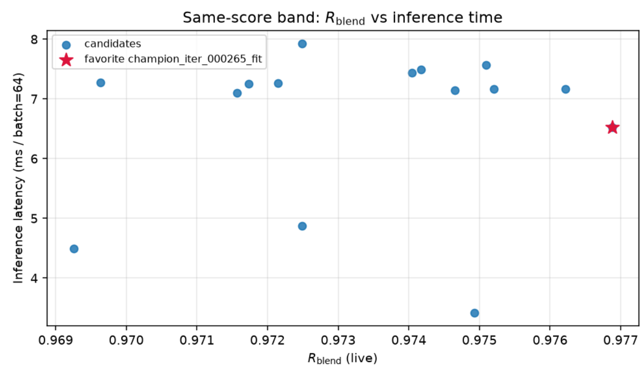

Figure 27: High-R fitted operators: residual vs CUDA latency. Star: favorite.

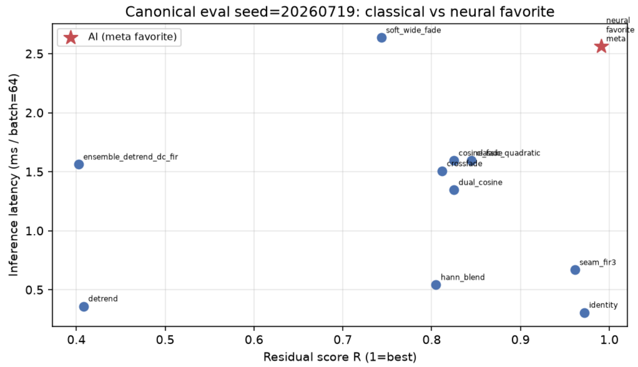

Figure 29: Canonical holdout: classical methods and favorite in the R vs ms/batch plane.

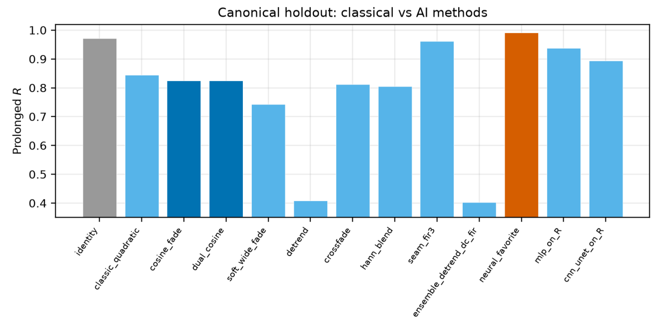

Figure 30: Canonical holdout ranking on R_{blend} (left) and latency (right).

C Transfer artifact locations

Full cross-domain wrap-heal transfer protocol, classical-board tables, and deep-model status rows are in Sections 4.2–5.21 of the main text. Algorithms for the residual score and hybrid loop remain in Appendix A. Raw JSON and figures: signal_heal_transfer.