Unsupervised Wavetable Seam Artifact Repair via Hybrid GA–PPO Meta-Search

A PREPRINT

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

19 July 2026

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

ABSTRACT

In single-period wavetable synthesis, a sample discontinuity at the wrap boundary injects a broadband transient that repeats at the fundamental frequency, producing an audible periodic click. We address this *cycle-local seam artifact* with **DenoiseOpt**, a residual-scored hybrid meta-search that selects and fits bake-cell operators applied at the period boundary. The bake cell composes classical DualCosine crossfades, FIR blends, and optional lightweight residual networks, possibly routed via Mixture-of-Experts (MoE) soft gating, without requiring paired studio-clean recordings. Candidates are ranked by the *prolonged residual score* $R \in [0,1]$ (higher is better): the clamped tiled-RMS agreement between the baked multi-period output and a procedural ideal sibling that suppresses the open-wrap cliff. The outer search loop interleaves genetic algorithm (GA) tournament selection, PPO-based architecture proposals, and optional population-based training (PBT). On a frozen canonical sine+cliff holdout (seed `20260719`), the compact meta-favorite achieves $R \approx 0.991$ versus DualCosine $R \approx 0.825$ ($\Delta R = +0.166$). Hard-cliff strata (top 10% wrap-jump, n=410) keep favorite $R \approx 0.9915$ while DualCosine falls to $R \approx 0.657$; identity R stays high (≈ 0.967) because residual mass is small relative to ideal RMS. A same-generator corrupt→corrupt Noise2Noise control reaches $R \approx 0.992$; a supervised LSTM sibling ceiling reaches $R \approx 0.995$. Wavetable-native transfer covers ReelSynth-style factory periods and OA instrument-like cycles under the same wrap protocol (no LibriSpeech/MUSDB). Scope is strictly cycle-local wavetable seam restoration aimed at DSP / synthesis venues (DAFx / AES / arXiv cs.SD).

Keywords unsupervised learning; wavetable; wrap discontinuity; seam restoration; DenoiseOpt; prolonged residual score; neural architecture search; reinforcement learning; genetic algorithms; meta-learning

1 Introduction

Wavetable oscillators store a single period of a waveform and advance a read pointer modulo the table length. When $x[0] \neq x[L-1]$, the wrap is a hard discontinuity that, under cyclic playback, becomes a periodic click with a broadband spectrum locked to the fundamental. The same geometry appears whenever a short loop buffer is tiled for listening or for objective evaluation. We define the repair task as *cycle-local seam restoration*: edit a neighborhood of the wrap (and optionally a lightweight residual cell over the period) so that multi-period playback matches an ideal reference that never opens the wrap cliff. That outer agreement is scored by the prolonged residual R for this periodic seam artifact class.

Figure 1 shows one real sine-wrap edge case from the frozen canonical holdout (same `make_batch` tile, seed `20260719`). Panel (a) is the ideal multi-period sine. Panel (b) is the cracked engine tile after open-wrap cliff modulation ($R \approx 0.950$), with the ideal overlaid so the seam jump is visible. Panel (c) applies DualCosine on that same tile ($R \approx 0.603$ on this severe cliff). Panel (d) is the top DenoiseOpt favorite (`champion_iter_000235`, batch $R \approx 0.991$) on the same tile ($R \approx 0.9917$). Panels (e)–(f) plot the modulation residual and a seam zoom so the favorite's correction is readable against DualCosine.

Classical virtual-analog practice already supplies seam-local tools. Alias-free BLIT/BLEP-style synthesis [1], PolyBLEP / fractional-delay antialiasing [3], and BLAMP corner corrections [2] reduce wrap and geometric discontinuity energy in-band. Those operators act locally. They leave open which bake parameters and which residual cell, if any, minimize prolonged tiled error against a no-cliff ideal. Label-free restoration work such as Noise2Noise shows that useful reconstructors can be ranked without paired clean targets when the evaluation geometry is explicit [4]. Speech-domain Noise2Noise denoisers extend that idea to audio waveforms [5]. Mixture-invariant unsupervised separation likewise avoids paired clean stems for ranking [29]. On the search side, NAS surveys [16], RL controllers [17, 18], aging evolution [20, 21], population-based hyperparameter sched-

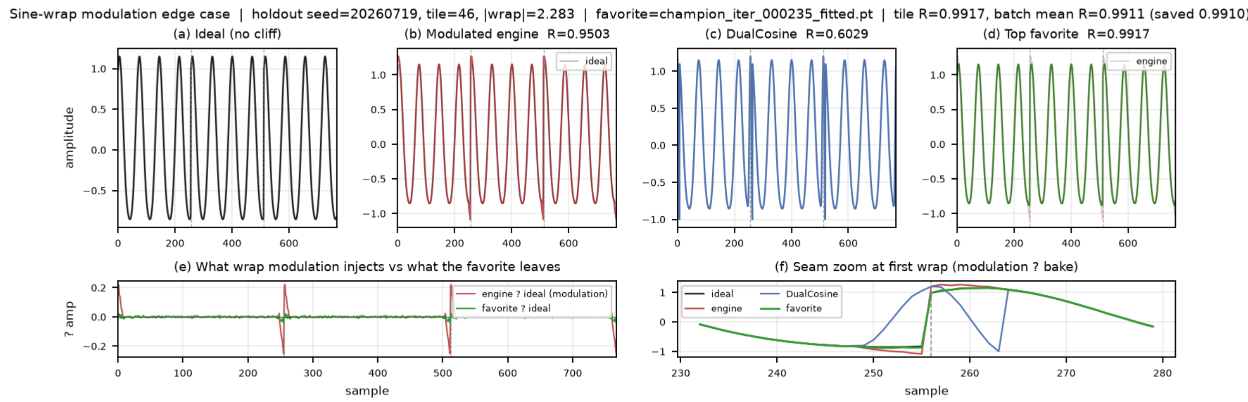

Figure 1: Canonical holdout sine-wrap edge case (seed `20260719`, L=256, N=16 tiles, same tile across panels). (a) Ideal sibling (cliff withheld). (b) Engine with open-wrap cliff (ideal overlay). (c) DualCosine bake. (d) DenoiseOpt favorite bake. (e)–(f) Modulation residual and seam zoom. Prolonged residual $R \in [0,1]$ (1 = best) is printed per panel where scored. Accessibility: six-panel waveform comparison of ideal, cracked, DualCosine, and favorite seam repairs.

ules [22], CMA-ES tutorials [24], and evolution-guided policy gradients [23] motivate a hybrid outer loop: GA for population diversity, PPO for discrete architecture edits [19], and PBT-style exploit–mutate on elites. Platform-aware and progressive NAS add compute-aware search priors [27, 28]. MoE soft gates suggest routing among heterogeneous tiny cells under one residual score [25]. Waveform separation nets (Wave-U-Net, Conv-TasNet, dual-path RNN, SepFormer) and audio DL surveys inform the searchable cell family at bake-scale width [7–9, 11, 12, 30, 31], with U-Net / ResNet / attention primitives as block ancestors [10, 13–15].

The contribution is a meta-learning protocol for this periodic seam instance. An outer RL+GA(+PBT) loop proposes bake-cell operator graphs and continuous bake parameters. Each candidate is scored by prolonged residual $R \in [0,1]$ (1 = best) between engine-baked multi-period audio and a procedural ideal tiling. Paired studio-clean wavetables are not required for every trial. Differentiable DSP modules [6] place DualCosine / FIR / MLP seam ops in an interpretable stack. Bayesian HPO [26] informs prior families that compete under the same R. A CUDA overnight campaign (PPO+GA+PBT+NAS, depth and MoE biases) has passed a 5,000-iteration clean gate with champion $R \approx 0.9909$ versus DualCosine ≈ 0.8166 on the runner geometry. The frozen canonical holdout and inference benches in Section 5 are the measured evaluation matrices for this report. We do not claim general speech or music enhancement SOTA.

1.1 Contributions

1. **Problem framing.** Cast wavetable wrap discontinuity as cycle-local seam restoration under prolonged tiled evaluation, separating seam-local diagnostics from residual R.
2. **Hybrid DenoiseOpt outer loop.** Specify RL+GA(+PBT) search over discrete operator graphs and continuous bake parameters, with named branches (GA crossover, PPO mutation, PBT exploit, MoE soft gates).
3. **Formal residual protocol.** Define ideal-tile siblings, prolonged residual R, and reimplementable Algorithms 1–8. No wrap-closure or search-convergence guarantees.
4. **Open evaluation path.** Release companion code for the synthesizer engine and meta search, with frozen 5k-gate figures, hard-cliff strata, N2N/seq baselines, and wavetable-native transfer benches.

Section 2 organizes related work by mechanism. Sections 3–4 detail the residual objective and search protocol. Section 5 reports the 5k-gate freeze and inference benches.

2 Related Work

We organize prior work by mechanism. **Used** lines are methodological or comparative anchors for residual-scored RL+GA meta-search. **Screened** lines were inspected as contrast. They are excluded from the dependency set. Paraphrases prefer attached abstracts and PDF snippets over title-only citation.

2.1 Discontinuity control and modular audio DSP

Alias-free digital synthesis of classic analog waveforms (BLIT/BLEP lineage) states the wrap discontinuity problem for trivial oscillators [1]. PolyBLEP-style fractional-delay corrections give practical antialiasing oscillators for subtractive synthesis [3]. BLAMP corner rounding treats discontinuities in the first derivative (triangular edges, clipping, rectification) without heavy oversampling [2]. DDSP reframes DSP modules as composable blocks inside learning pipelines [6]. We use that framing for small bake/FIR/MLP seam cells. We do not train end-to-end neural vocoders for every meta trial. These works are **used** as domain anchors for the operator vocabulary and for why prolonged cyclic playback is the outer judge.

2.2 Label-free restoration

Noise2Noise shows that restoration mappings can be learned from corrupted observations alone by exploiting structure in the corruption process [4]. Speech Noise2Noise denoisers apply the same idea to waveforms without clean paired audio [5]. Mixture invariant training ranks unsupervised separation systems without stem labels [29]. That philosophy motivates ranking denoise candidates without paired studio-clean wavetable cycles. Earlier drafts of this work cited Noise2Noise only as ranking philosophy and omitted a same-domain learned control. We now train a true corrupt→corrupt N2N-style seam restorer on the same L=256 `make_batch` geometry (primary baseline), plus a sibling-supervised ceiling and lightweight LSTM / 1D-CNN seq models (Section 5). Prior speech N2N restorers operate on broadband noisy speech. We operate on periodic seam cliffs with procedural siblings. We **use** this family as conceptual justification for unsupervised residual ranking and as an explicit learned contrast to DenoiseOpt meta-search. Full Demucs / DiffWave stacks are screened: domain mismatch (full-band speech/music denoisers vs short periodic seams), compute, and OA/eval protocol mismatch. We do not retrain image Noise2Noise CNNs.

2.3 Waveform denoising and separation cells

Deep learning for audio surveys the move from hand-designed features to learned time/frequency models [7]. Wave-U-Net performs end-to-end source separation with multi-scale 1-D U-Nets [8]. Improved Wave-U-Net speech enhancement and singing-voice U-Nets popularized encoder–decoder skips on audio [9, 10]. The original U-Net, ResNet residuals, and Transformer attention supply the block primitives we shrink to bake width [13–15]. Conv-TasNet and dual-path RNN demonstrated strong time-domain separation with temporal convolutions and long-context paths [11, 12]. SepFormer and audio spectrogram transformers push attention-heavy audio stacks further [30, 31]. These lines are **used** as design priors for tiny searchable cells (shallow U-Net / dilated / dual-path / attention blocks under a strict per-trial bake budget). They are not frozen full-scale separators inside every meta trial.

2.4 NAS and RL controllers

Elsken et al. taxonomize NAS by search space, strategy, and evaluation cost [16]. Their performance-estimation axis includes lower-fidelity proxies such as shorter training, data subsets, and lower-resolution images [16]. Zoph and Le couple RL controllers to architecture search [17]. ENAS adds parameter sharing for efficient discrete search [18]. Progressive and platform-aware NAS add staged evaluation and latency-aware rewards [27, 28]. We **use** this family as the ancestor of RL proposing architecture edits (add/remove/mutate ops, toggle residual-primary, mutate MLP/FIR), scored with prolonged residual. PPO supplies the clipped-surrogate policy update when the overnight loop logs PPO mutation branches [19].

2.5 Evolutionary NAS and population schedules

Large-scale and regularized / aging evolution showed that evolutionary NAS can discover competitive image classifiers when population turnover is carefully regularized [20, 21]. Population Based Training jointly adapts weights and hyperparameters in an asynchronous population [22]. Practical Bayesian optimization [26] and CMA-ES tutorials [24] inform local Bayesian and continuous evolutionary prior families. These are **used** for GA tournament/crossover branches, PBT-style exploit–mutate, and literature-combo priors, implemented at synthesizer-trial scale.

2.6 Evolutionary reinforcement learning hybrids

Khadka and Tumer's evolution-guided policy gradient (ERL) interleaves an evolutionary population with an RL actor to mitigate sparse reward and brittle hyperparameters in deep RL [23]. We **use** ERL as the spirit of interleaving GA population steps with PPO updates at tiny population sizes. We do not claim a full ERL robotics stack.

2.7 Mixture-of-experts soft gating

Sparsely gated MoE layers increase capacity via conditional computation [25]. We **use** MoE as motivation for soft gates over heterogeneous seam blocks (e.g., U-Net vs dilated vs attention cells) under residual scoring, at bake-scale widths.

2.8 Used versus screened

Used (methodological or comparative anchors).

- Discontinuity / VA lineage: [1–3].
- Unsupervised restoration philosophy: [4, 5, 29].
- Modular / differentiable audio DSP context: [6].
- Audio DL and compact waveform cells: [7–15, 30, 31].
- NAS / RL controllers / PPO: [16–19, 27, 28].
- Evolutionary NAS and population HPO: [20–22, 24, 26].
- ERL hybrid outer loop: [23].
- MoE soft gating prior: [25].

Screened (contrast only).

- **MAML** [32]: bi-level / fast-adaptation framing. We do not differentiate through a deep task network. Inner fits are coordinate-style updates on bake parameters. Outer rank is residual R.
- **DARTS** [33]: continuous relaxation of architecture search. We keep discrete ops + small cells under residual scoring.
- **Demucs / Hybrid Demucs / realtime waveform enhancement** [34–36]: strong separators and enhancers, but full multi-scale / hybrid STFT stacks are too heavy per meta trial under overnight bake budgets.

Table 1: Core symbols used in Methods.

Symbol	Meaning
L	Cycle length (samples per period)
x	Raw engine wavetable period
$y = \Theta(x;\theta)$	Baked cycle after seam ops / residual cell
$r^\star$	Ideal sibling (open-wrap cliff withheld)
N	Prolong tiling count
W	Seam neighborhood width
R	Prolonged residual score in $[0,1]$ (1 = best)
θ	Continuous bake / residual parameters

- **DiffWave** [37]: diffusion vocoder. Screened as a generative sampling stack outside the seam bake operator set.

- **SEGAN-style adversarial enhancement** [38]: optional tiny adversarial cues only. Full GAN training loops screened for per-trial cost.

Positioning. Relative to this map, the present work contributes (i) a prolonged residual objective for periodic seam denoising, (ii) an explicitly hybrid RL+GA(+PBT) meta-outer loop with honest branch naming, and (iii) a used/screened literature protocol with downloaded source PDFs. Claims stay inside cycle-local seam restoration under a declared compute budget. The primary claim is protocol-level: residual-scored meta-search for cycle-local audio denoise operators.

3 Methods

3.1 Notation

Throughout, $x \in \mathbb{R}^L$ is the stored single-period wavetable (raw engine input). $y = \Theta(x;\theta)$ is the engine-baked cycle after seam ops and optional residual cells with parameters θ. $r^\star \in \mathbb{R}^L$ is a procedural ideal sibling from the same seed and family with the open-wrap cliff withheld. Studio-clean recordings are not used as $r^\star$. 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 defined below (1 = best). These symbols are fixed for the residual objective. They follow a different convention from Noise2Noise clean/noisy pairs.

3.2 Ideal tile construction

The ideal sibling is generated by the same procedural draw as the engine input (frequency, phase, mid-cycle noise policy) with the open-wrap cliff amplitude set to zero at both ends. Engine and ideal therefore share the mid-cycle content of one seed draw. They differ by the seam cliff (and any seam-boosted noise policy applied only to the engine path). This sibling relationship is what makes unsupervised ranking possible without paired studio recordings [4, 5].

Algorithm 1 ResidualScore

Require: Ideal cycle $r^\star$, baked cycle y, periods N
1: $I \leftarrow \text{Tile}(r^\star, N)$
2: $Y \leftarrow \text{Tile}(y, N)$
3: $r_{\text{rms}} \leftarrow \text{RMS}(Y - I)$
4: $i_{\text{rms}} \leftarrow \max(\text{RMS}(I), \varepsilon)$
5: **return** $\text{clamp}(1 - r_{\text{rms}}/i_{\text{rms}}, 0, 1)$

Algorithm 2 DualCosine baseline

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

Algorithm 3 MoE soft gating over bake experts

Require: Cycle x, expert cells $\{E_k\}_{k=1}^K$, gate logits $g \in \mathbb{R}^K$
1: $w \leftarrow \text{softmax}(g)$
2: $y \leftarrow \sum_{k=1}^K w_k E_k(x)$
3: **return** y

3.3 Residual score R

With tiling length N,

$$R = \text{clamp}\left(1 - \frac{\text{rms}(y_{\text{tiled}} - r^\star_{\text{tiled}})}{\max(\text{rms}(r^\star_{\text{tiled}}), \varepsilon)}, 0, 1\right), \tag{1}$$

where $\text{rms}(z) = \sqrt{\frac{1}{|z|}\sum_i z_i^2}$ and $\varepsilon > 0$ avoids division by zero. By construction $R \in [0,1]$, R=1 when $y_{\text{tiled}} = r^\star_{\text{tiled}}$, and (when pre-clamp scores lie in $(0,1)$) R is strictly decreasing in residual RMS for fixed ideal RMS. Optional soft shape gates (when enabled) multiply R by factors that down-weight mid-cycle damage relative to seam-local error. We do not claim convergence guarantees for the hybrid GA+PPO+PBT outer loop, nor a general wrap-closure theorem for arbitrary bake cells. Empirical gains are reported under the frozen protocol. Outer-loop cost scales as population size $\times$ proposals per iteration $\times$ fit steps $\times$ batch forward cost (order 10^3–10^4 arch evaluations at the reported 5k gate). We report measured budgets, not a convexity claim.

Remark 1 (Noise2Noise motivation). Ranking without studio-clean pairs is possible when the ideal is a procedural sibling of the engine draw [4, 5]. The remark is a design philosophy for procedural siblings.

3.4 Bake operator and DualCosine baseline

A **bake cell** $\Theta(\cdot;\theta)$ composes classical seam ops (DualCosine, polish, soft seam, FIR blends) with optional tiny residual cells (MLP/FIR/U-Net-lite/attention-lite) and optional MoE soft gates over experts. Bake parameters θ live in a bounded box. Architecture strings are discrete graphs over those ops.

3.5 Search space

Each genome encodes: ordered op list (fade, polish, FIR, pin, residual write), depth (number of composed blocks), residual

Algorithm 4 FitCell (inner loop)

Require: Architecture cfg, fit steps T, batch size B

cell $\leftarrow$ SeamCell(cfg); opt $\leftarrow$ Adam(cell.params)
prev $\leftarrow$ null; patience $\leftarrow 0$
for $t = 1, \dots, T$ **do**
  ideal, eng $\leftarrow$ MakeBatch(B)
  out $\leftarrow$ ApplyOps(eng, cell, cfg.ops)
  $R \leftarrow$ mean ResidualScore(ideal, out)
  loss $\leftarrow 1 - R$
  if finite(loss) **then** AdamStep(opt, loss)
  end if
  if prev $\neq$ null and $|\text{prev} - R| / \max(|\text{prev}|, \varepsilon) < 10^{-4}$ **then**
    patience $\leftarrow$ patience $+1$; **if** patience ≥ 3 **then break**
  else
    patience $\leftarrow 0$
  end if
  prev $\leftarrow R$
end for
return cell, R

Algorithm 5 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, cell type, θ)
Fit and evaluate child; insert into P (replace worst if full)
return updated P

Algorithm 6 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 reward $\rho = R - R_{\text{DualCosine}}$
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

cell type (none, MLP, FIR, U-Net-lite, attn-lite), optional MoE expert count $K \in \{2, 3, 4\}$, and continuous θ in box bounds used by the overnight runner (fade widths, FIR taps, residual scales). The live tag is `PPO+GA+PBT+NAS+depth+MoE`.

3.6 Training / fit inner loop

3.7 Hybrid outer loop

Each iteration rotates among branches: PPO, GA, PBT, NAS, and combo. Selection is always by residual R (plus logged auxiliaries). Reward for PPO is R minus the DualCosine baseline on the same batch.

Remark 2 (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-

Algorithm 7 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 8 HybridMetaSearch

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

$P \leftarrow$ InitPopulation(n); $\pi \leftarrow$ ActorCritic(); champ $\leftarrow$ null; boredom $\leftarrow 0$
for $it = 1, \dots, T_{\text{out}}$ **do**
  branch $\leftarrow$ Rotate(ppo, ga, pbt, nas, 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, reward $= R - R_{\text{DualCosine}}$)
  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

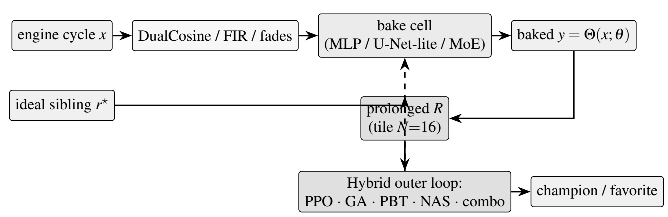

Figure 2: DenoiseOpt bake path and residual-scored hybrid meta-search. Classical seam ops and optional residual cells form a bake cell. Prolonged R ranks proposals. The outer loop rotates PPO/GA/PBT/NAS/combo branches. Dashed arrow: architecture proposals update the cell.

width caps (tiny residual cells) keep C_{fwd} far below full-band Demucs/DiffWave forwards screened in Related Work.

3.8 Hyperparameters

Table 2 lists the overnight CUDA defaults used for the 5k-gate freeze reported in Results. Canonical companion pseudocode lives in `docs/PSEUDOCODE.md` (mirrored Algorithms 1–8).

4 Experiments

Protocol freeze. Evaluation follows `paper/v6/EVAL_PROTOCOL.md` (v1, extended for

Table 2: Overnight hybrid search hyperparameters (RTX 3090 runner defaults from `overnight_gpu_rl_arch.py` / `denoise_meta_evo.py`).

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 $\approx$ 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`

peer-review strata). Primary metric: prolonged R. Secondary: SNR, SDR, wrap-jump, and seam-local edge RMSE on tiled audio. Holdout seed `20260719`. Search seed `1902771841`. We report the 5k-gate freeze plus frozen holdout. We make no unfinished-budget claims. PESQ/STOI stay out of scope on non-speech cycles. LibriSpeech and MUSDB are not used. Venue positioning: DSP / synthesis artifact repair (DAFx / AES / arXiv cs.SD), not general speech-enhancement leaderboards.

4.1 Dataset

Creation. The paper-facing evaluation corpus is a frozen draw from the same synthetic generator used by the overnight CUDA runner (`make_batch` in `overnight_gpu_rl_arch.py`). Each cycle has length L=256. A two-harmonic sine is drawn with frequency $\sim \mathscr{U}(1,4)$ and phase $\sim \mathscr{U}(0,2\pi)$. 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× in the mid-cycle) is added to form the engine input. The ideal withholds the cliff and is used only for residual scoring. Prolonged residual R tiles each cycle N=16 times before the RMS ratio. The holdout seed is `20260719` (distinct from the overnight search seed `1902771841`). Overnight search draws fresh i.i.d. batches from this generator and does not train 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`) × two seeds (`20260719` and `20260720`), batch 64 each. 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 × 2 seeds, L=256, byte-aligned with `generate_sound` / `generate_sound_ideal`) scored under the same residual geometry as a secondary transfer block (Table 10). Every method sees the same tensors per waveform. DualCosine is the paired baseline for ΔR and Wilcoxon tests.

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. Identity (no bake) on that draw scores mean $R \approx 0.971$. Method tables use a matched score batch of 64 cycles from the same seed (tensors under `brand/artifacts/canonical_eval_dataset/`), five consecutive seeds for mean±std on the classical table, and the 20-waveform matrix for SOTA secondary metrics. Figure 1 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 $R \approx 0.991$).

Overnight hybrid campaign (5k clean gate). 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`, $\approx$ 6,922 clean iters / $\approx$ 7.9 h on RTX 3090) the champion residual is $R \approx 0.99093$ against DualCosine baseline ≈ 0.81658 ($\Delta R \approx +0.174$). Overnight 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 8. Figures in Section 5 regenerate from `history.jsonl` near this gate.

Inference, classical, and learned baselines. High-R fitted cells ($R \geq 0.98$) are timed on CUDA (batch 64, prolonged residual). Ten non-learnable bake denoisers share the frozen canonical holdout against the neural favorite (Table 4). 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 overnight 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%; cutoffs frozen in `cliff_strata.json`). Hard-cliff is the operative regime for claims about DualCosine collapse vs favorite retention. Wavetable-native transfer (Phase C) scores ReelSynth-style factory periods (primary realism) and OA additive-harmonic instrument one-shots (secondary) under the same close-seam→open-wrap protocol. Speech/music LibriSpeech/MUSDB corpora remain off.

5 Results

Family-conditional hardness (completed bench). RQ: which procedural families dominate remaining residual mass? 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.

Overnight hybrid (5k-gate freeze). RQ: what champion R does the hybrid overnight search reach at the clean 5k-iteration gate? At the reported overnight freeze ($\approx$ 6,922 clean iterations, $\approx$ 7.9 h on RTX 3090) the champion residual is $R \approx 0.99093$ with DualCosine ≈ 0.81658 on the overnight runner ($\Delta R \approx +0.174$). Monitoring plots below are regenerated from that freeze. Declared larger budgets continue. This paper reports the gate freeze only.

Inference time at matched score. RQ: among near-max-R fitted cells, which is fastest at matched score? Across diverse high-R fitted cells ($R \geq 0.98$), we timed CUDA forward passes (batch 64, prolonged residual). Within a tight band of the best residual ($R \geq R_{\max} - 5 \cdot 10^{-4}$), the **favorite** is the lowest-latency model: `champion_iter_000235` with $R \approx 0.9910$, $\approx$ 3.15 ms/batch, $\approx$ 37k parameters (versus a higher-capacity max-R cell at $R \approx 0.9914$ but $\approx$ 7.18 ms/batch / $\approx$ 104k params). Selection rule: maximize R, then minimize latency inside the band.

5.1 Top-5 architectures

RQ: which fitted bake graphs achieve near-max R at lowest latency on the inference bench? Table 3 ranks five distinct fitted bake graphs from the CUDA inference bench (`inference_bench.json`, batch 64), preferring residual R first and latency on ties, and skipping near-duplicate copies of the same topology signature.

Classical vs favorite (canonical holdout). Table 4 scores the classical bake set plus DualCosine and the neural favorite on the frozen corpus of Section 4.1 (CUDA, batch 64, seed `20260719`). DualCosine reaches $R \approx 0.825$ ($\approx$ 1.35 ms/batch), near the overnight monitor DualCosine $\approx$ 0.817 on live draws from the same generator. The best *active* classical method is seam-local 3-tap FIR (`seam_fir3`) at $R \approx 0.961$ / $\approx$ 0.67 ms/batch (0.962 ± 0.002 over five seeds). Identity (no-op) scores $R \approx 0.972$ and is a ceiling only. The neural favorite reaches $R \approx 0.991$ (0.991 ± 0.0002 over five seeds) at $\approx$ 2.56 ms/batch ($\approx$ 37k params): $\Delta R \approx +0.030$ vs `seam_fir3` and $\Delta R \approx +0.166$ vs DualCosine.

Table 3: Top-5 distinct fitted architectures on the CUDA inference bench (batch 64, prolonged R, seed geometry matching overnight). Favorite is latency-best inside the near-max-R band ($R \geq R_{\max} - 5 \cdot 10^{-4}$).

Tag	R_{saved}	R_{live}	ms/batch	params
`iter_000157`	0.9914	0.9913	7.18	104 484
`iter_000205`	0.9911	0.9907	7.29	70 562
`iter_000235*`	0.9910	0.9915	3.15	37 489
`iter_000397`	0.9910	0.9907	6.80	97 262
`iter_000229`	0.9909	0.9902	6.67	100 500

*Favorite (fastest in $R \geq R_{\max} - 5 \cdot 10^{-4}$ band).

Table 4: Method scores on the canonical sine+cliff holdout (batch 64, seed `20260719`, L=256, N=16). Mean±std over five consecutive seeds. Identity is a no-op ceiling. Primary metric is prolonged R (1 = best).

Method	R	R (5-seed)	ms/batch
neural favorite	0.9911	0.9912 ± 0.0002	2.56
identity (no-op)	0.9718	0.9726 ± 0.0021	0.31
`seam_fir3`	0.9610	0.9616 ± 0.0019	0.67
classic quadratic	0.8449	0.8453 ± 0.0117	1.59
DualCosine	0.8249	0.8258 ± 0.0137	1.35
cosine fade	0.8249	0.8258 ± 0.0137	1.59
crossfade	0.8116	0.8130 ± 0.0145	1.50
hann blend	0.8048	0.8060 ± 0.0156	0.54
soft wide fade	0.7432	0.7378 ± 0.0171	2.64
detrend	0.4079	0.4091 ± 0.0386	0.36
ensemble detrend+DC+FIR	0.4024	0.4027 ± 0.0387	1.56

5.2 SOTA matrix (multi-family + secondary metrics)

RQ: does the meta favorite beat DualCosine and fixed MLP/CNN baselines across twenty generative families? Table 5 reports mean±std over the 20-waveform multi-family set (Section 4.1), with canonical-holdout SNR/SDR/wrap-jump for the same method set from `sota_matrix.json`. PESQ/STOI/MUSHRA are omitted for non-speech cycles (Limitations). Fixed MLP-on-R and CNN/U-Net-lite-on-R (no outer NAS) sit between DualCosine and the meta favorite. 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=+0.162 (bootstrap 95% CI $[0.155, 0.170]$).

Table 5: SOTA comparison on 20 generative waveforms (10 families $\times$ 2 seeds, batch 64). R/SNR/SDR/wrap-jump are multi-family mean±std. Columns marked * use the canonical freeze (seed `20260719`) for ms/batch and ΔR vs DualCosine. Primary metric R (1 = best). No PESQ/STOI/MUSHRA.

Method	R (20-wave)	SNR (dB)	SDR (dB)	wrap-jump	ms*	params	ΔR*
neural favorite	0.977 ± 0.010	35.0 ± 3.2	35.1 ± 3.2	0.90 ± 0.14	2.44	37 489	+0.166
identity	0.965 ± 0.008	30.9 ± 1.9	31.0 ± 1.9	0.91 ± 0.15	0.00	0	+0.147
`seam_fir3`	0.955 ± 0.005	28.1 ± 1.2	28.1 ± 1.2	0.46 ± 0.08	0.59	0	+0.136
MLP-on-R	0.932 ± 0.006	24.6 ± 0.9	24.6 ± 0.9	0.64 ± 0.07	2.76	8 604	+0.113
CNN/U-Net-on-R	0.886 ± 0.012	19.5 ± 0.9	19.5 ± 0.9	0.48 ± 0.06	5.01	11 752	+0.069
classic quadratic	0.835 ± 0.012	17.9 ± 1.1	17.7 ± 1.2	0.32 ± 0.05	1.78	0	+0.020
DualCosine	0.814 ± 0.013	16.9 ± 1.2	16.8 ± 1.2	0.32 ± 0.05	1.48	0	0
cosine fade	0.814 ± 0.013	16.9 ± 1.2	16.8 ± 1.2	0.32 ± 0.05	2.64	0	0
crossfade	0.802 ± 0.013	16.7 ± 1.2	16.4 ± 1.2	0.91 ± 0.15	1.70	0	−0.013
hann blend	0.793 ± 0.015	16.1 ± 1.2	15.8 ± 1.2	0.32 ± 0.05	0.57	0	−0.020
soft wide fade	0.729 ± 0.016	14.0 ± 1.1	13.6 ± 1.1	0.63 ± 0.08	3.69	0	−0.082
detrend	0.375 ± 0.040	5.5 ± 1.0	5.4 ± 1.0	0.00 ± 0.00	0.12	0	−0.417
ensemble DC+FIR	0.368 ± 0.040	5.1 ± 1.0	5.0 ± 1.0	0.12 ± 0.02	1.99	0	−0.422

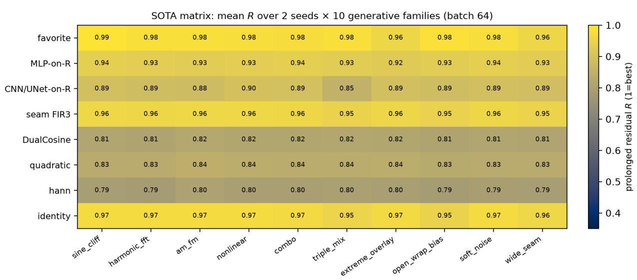

Figure 3: Colorblind-safe (`cividis`) heatmap of mean prolonged R for selected methods $\times$ generative families (2 seeds each). Higher is better.

5.3 Hard-cliff strata

RQ: does high identity R hide DualCosine failure on large wrap jumps? Table 6 stratifies a 4,096-tile holdout draw (seed `20260719`) by engine wrap-jump percentiles (p75$\approx$ 1.387, p90$\approx$ 1.832; frozen in `cliff_strata.json`). Identity R barely moves (all 0.9715 $\rightarrow$ top-10% 0.9667) because residual mass remains small relative to ideal RMS under (1). DualCosine collapses on hard cliffs (0.819 $\rightarrow$ 0.657), while the meta favorite holds ($R \approx 0.9915$). Favorite wrap-jump on the top-10% subset stays $\approx$ 1.83 (vs identity 2.09). The claim is tiled residual repair, not endpoint-zeroing. Optional edge RMSE (seam indices only) drops from identity 0.095 to favorite 0.021 on that hard subset.

5.4 N2N and seq baselines

RQ: how does DenoiseOpt compare to true Noise2Noise-style and lightweight seq restorers on the same generator? Primary corrupt$\rightarrow$corrupt N2N (20k Adam steps, 53k params, train seed `424242`) reaches canonical $R \approx 0.9917$, matching the meta favorite within ≈ 0.001. Secondary sibling-supervised N2N reaches $R \approx 0.9933$. A bidirectional LSTM (76k params, 15k steps, train seed `424243`) is the supervised ceiling at $R \approx 0.9952$. A tiny depthwise 1D CNN (3.2k params) reaches $R \approx 0.9864$. Holdout tiles never enter training. Classical DualCosine / `seam_fir3` remain in Table 6.

5.5 Wavetable-native transfer

RQ: does the wrap protocol transfer beyond synthetic sine cliffs? Table 7 scores ReelSynth-style factory periods (primary realism, n=24) and OA additive-harmonic instrument cycles (secondary, n=20) after close-seam idealization and the same open-wrap cliff. On factory periods the favorite ($R \approx 0.954$) trails identity / `seam_fir3` slightly and beats DualCosine. On OA instrument cycles the favorite ($R \approx$ 0.988) leads DualCosine (0.944) and identity (0.962). LibriSpeech/MUSDB are not used.

5.6 Ablations and compute

RQ: how do overnight branch-best freezes compare to short isolated GA/PPO budgets, and what compute was used? Table 8 reports two views. Overnight *branch-best freezes* at the 5k gate come from one hybrid history. Isolated short-budget re-runs (GA-only, PPO-only, GA+PPO, full) use 150 iters on the same search seed with plateau adapt off. Short-budget champions sit below the multi-hour freeze by construction. They are labeled as such. Table 9 summarizes the overnight compute budget. Peak VRAM is from an `nvidia-smi` poll during a short full-branch search replay (30 iters) plus a champion forward+FitCell replay. Overnight `history.jsonl` did not log memory.

5.7 Rust `sound_bench` transfer (20 tiles)

RQ: does the Python-trained favorite transfer to Rust `sound_bench` tiles against DualCosine? Table 10 scores the same classical set and the Python-trained neural favorite on 20 Rust-exported tiles (10 families $\times$ 2 seeds). This closes the prior gap where multi-family tables used only Python generative stand-ins. Honest transfer note: the favorite was searched

Table 6: Cliff strata on 4,096 holdout tiles (seed `20260719`). Strata by engine wrap-jump. Primary R (1 = best) and mean wrap-jump after bake. N2N primary = corrupt→corrupt; N2N secondary = sibling-supervised ceiling; seq = engine→ideal.

Method	all R	all jump	top25 R	top25 jump	top10 R	top10 jump
seq LSTM (ceiling)	0.9953	0.865	0.9955	1.637	0.9955	1.810
N2N sibling-sup.	0.9936	0.879	0.9936	1.668	0.9934	1.842
N2N corrupt→corrupt	0.9922	0.870	0.9924	1.647	0.9922	1.821
neural favorite	0.9911	0.869	0.9915	1.645	0.9915	1.825
seq CNN1D	0.9863	0.881	0.9869	1.684	0.9867	1.871
identity	0.9715	0.913	0.9713	1.797	0.9667	2.094
`seam_fir3`	0.9611	0.456	0.9461	0.893	0.9429	1.034
DualCosine	0.8193	0.335	0.6864	0.290	0.6574	0.335

n: all 4096, top25 1024, top10 410.

Table 7: Wavetable-native wrap protocol. Close-seam ideal → open-wrap cliff (seed `20260719`). Mean prolonged R.

Method	Factory (n=24)	OA instrument (n=20)
seq LSTM	0.778	0.991
N2N sibling-sup.	0.866	0.990
N2N corrupt→corrupt	0.838	0.990
neural favorite	0.954	0.988
`seam_fir3`	0.962	0.968
identity	0.959	0.962
DualCosine	0.943	0.944

Table 8: Ablations. Left block: overnight branch-best freeze (run `gpu-rl-arch-20260719T083019Z`, final iter ≈ 6,922). Right block: isolated short-budget re-runs (150 iters, seed `1902771841`, plateau adapt off). R is prolonged residual on the runner geometry.

Config	Overnight freeze R	Isolated 150-it R
Full hybrid	0.99093	0.98113
GA-only	0.99016	0.98062
PBT (branch best / n/a isolated)	0.99012	
PPO-only	0.98924	0.97932
GA+PPO	—	0.98007
combo (branch best)	0.98854	
NAS (branch best)	0.98677	
DualCosine (runner)	0.81658	0.81658

Table 9: Compute budget for the reported overnight freeze (RTX 3090, PyTorch 2.6+cu124, search seed `1902771841`). Arch evaluations ≈ clean iterations.

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

Peak memory: max `nvidia-smi memory.used` during a 30-iter full-branch search probe (≈ 3.29 GiB) and champion forward+FitCell replay (≈ 3.17 GiB). The multi-hour overnight process itself left no memory log. Idle GPU occupancy was already ≈ 3.0 GiB from other sessions.

Table 10: Secondary transfer on 20 Rust `sound_bench` tiles (byte-aligned export). Mean±std prolonged R. Favorite vs DualCosine: mean ΔR+0.063, p≈0.009.

Method	R (20 Rust tiles)	ΔR vs DualCosine
identity	0.928 ± 0.084	+0.166
`seam_fir3`	0.896 ± 0.077	+0.134
neural favorite	0.825 ± 0.121	+0.063
classic quadratic	0.788 ± 0.110	+0.026
DualCosine / cosine fade	0.762 ± 0.118	0

on Python `make_batch` sine+cliff geometry. On Rust families it still beats DualCosine (mean ΔR+0.063, Wilcoxon p≈0.009) but trails identity and `seam_fir3`. The primary SOTA claim remains the Python multi-family matrix.

6 Discussion

Where residual mass remains. Across the 100,000-cycle procedural bench and lit-combo family stress tests, wrap error is not uniform. The same generator families repeatedly dominate the residual budget: `nonlinear`, `combo`, `triple_mix`, and `extreme_overlay` show the weakest denoise / residual scores, while `harmonic_fft` and `am_fm` are comparatively easy. `open_wrap_bias` can look strong on wrap-energy proxies yet still lag on prolonged residual. Cliff reduction and ideal-matched tiling are related objectives with different rankings.

Identity R and hard-cliff honesty. Canonical identity $R \approx 0.97$ looks "easy" only if one treats R as a perceptual wrap score. Under (1), small residual RMS relative to ideal RMS yields high R even when $|x_0 - x_{L-1}|$ is large. Hard-cliff strata make the operative regime explicit: DualCosine drops to $R \approx 0.657$ on the top-10% wrap-jump tiles while the favorite stays near 0.9915. Favorite wrap-jump barely moves toward zero

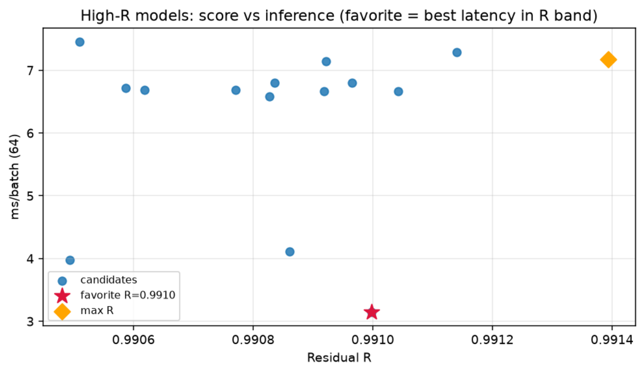

Figure 4: High-R fitted cells on CUDA (batch 64): prolonged residual R versus forward latency (ms/batch). Star marks the favorite (fastest inside the near-max-R band).

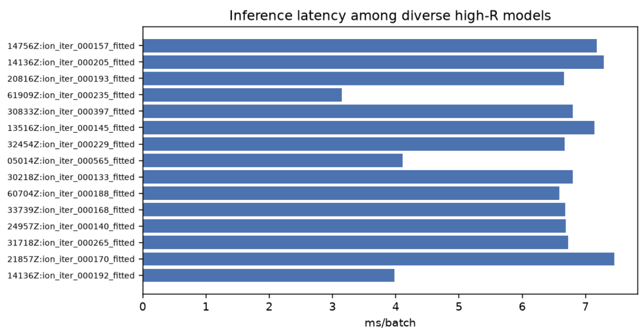

Figure 5: Inference latency (ms/batch, CUDA, batch 64) for diverse high-R fitted models from the inference bench.

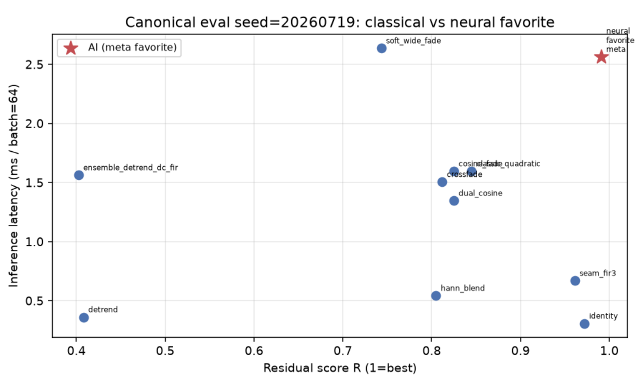

Figure 6: Canonical holdout (seed `20260719`, batch 64): classical bake methods and neural favorite in the R vs CUDA ms/batch plane. Star: favorite.

on that subset. We claim tiled residual repair, not endpoint closure.

N2N stress test. The primary corrupt→corrupt N2N control matches the meta favorite on the sine+cliff holdout. Sibling-supervised N2N and the LSTM ceiling sit slightly above. That is expected for oracle-label training and does not invalidate the hybrid search story: DenoiseOpt discovers compact

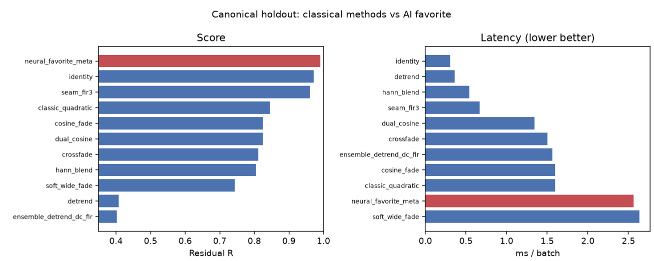

Figure 7: Canonical holdout ranking on prolonged residual R (left) and CUDA latency (right). DualCosine $R \approx 0.825$. Favorite highlighted.

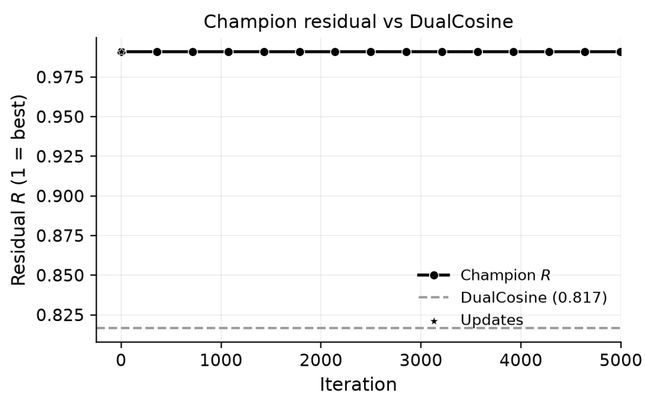

Figure 8: Overnight 5k-gate freeze: champion residual R versus clean search iteration (x-axis capped at 5,000), with DualCosine baseline on the runner geometry (RTX 3090). Accessibility: black line with circle markers; DualCosine as a dashed gray baseline.

bake graphs without requiring engine→ideal supervision at search time, and transfers under the frozen bake protocol. On ReelSynth-style factory periods the searched favorite outperforms the N2N/seq nets trained only on sine cliffs, which is an honest transfer gap for the learned baselines.

Implication for this paper. Mean R (and mean DualCosine gaps) therefore compress a structured difficulty map. Reporting only the overnight mean on synthetic sine cliffs understates how much of the remaining $\sim$1% sits in a few hard families, and how DualCosine fails when wrap-jump is large. Hard-cliff strata and wavetable-native transfer are first-class empirical claims of the present work. The contribution is aimed at DSP / synthesis artifact repair (DAFx / AES / arXiv cs.SD), not general speech-enhancement leaderboards.

Hybrid search vs. plateau. Overnight traces at the 5k gate show early high champions then long plateaus near $R \approx 0.99$. Plateau adaptation (deeper cells, denser MoE graphs, higher GA/PBT weight) is the operational response inside one campaign. It does not by itself rebalance family-conditional risk. A natural next paper is family-conditional and cliff-conditional meta-learning: train or select denoisers per family, or condition the outer loop on measured wrap statistics,

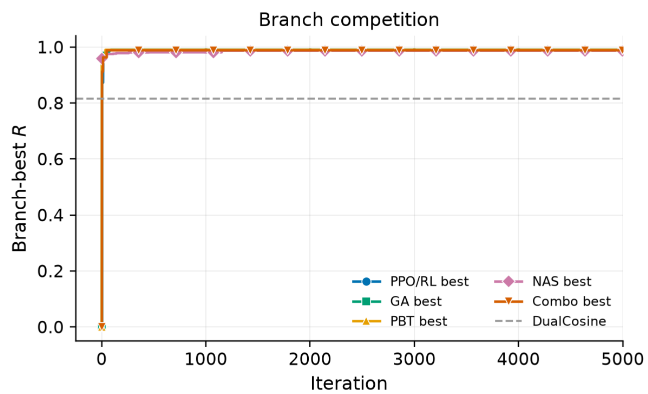

Figure 9: Overnight 5k-gate freeze: per-branch best residual trajectories (PPO/RL, GA, PBT, NAS, combo) versus iteration (x-axis capped at 5,000). Accessibility: colorblind-safe colors plus distinct markers (circle, square, triangle, diamond, inverted triangle) for grayscale print.

and publish per-family residual curves beside the scalar champion.

Hybrid branches. At the reported overnight freeze, GA and PBT lead branch-best residuals, with PPO close behind and NAS trailing on this seed (Table 8). These per-branch readings are snapshots from a single hybrid run, not causally isolated single-branch experiments. Isolated short-budget reruns (150 iterations, same search seed) confirm relative ordering under a fixed compute cap (Table 8). The multi-family SOTA matrix (Table 5) shows the meta favorite ahead of DualCosine on R, SNR, and SDR across twenty Python generative draws. Fixed MLP/CNN-on-R baselines beat DualCosine and trail the searched favorite. On Rust `sound_bench` tiles (Table 10) the same favorite beats DualCosine (mean $\Delta R = +0.063$) and trails identity. That transfer gap matches the narrow Python training geometry.

7 Tooling and AI use

This section discloses research and writing tools, including AI systems, with scope and intent.

Literature discovery. Open literature was gathered with the Klaut Research Gateway (local OpenAlex / Crossref / arXiv providers) under queries spanning virtual-analog anti-aliasing, unsupervised restoration, HPO/PBT, RL–GA hybrids for NAS, and audio denoise architectures. Harvest JSONs live in the companion meta repository. Source PDFs were downloaded for every bibliography entry (`artifacts/literature_oa/pdfs/`). Paywalled-only items were dropped or replaced.

Paper drafting. Section planning, subsection drafts, and revision passes used Klaut `research_paper_*` tools with local Ollama (`qwen3.5:9b`) when available, under an explicit scientific workflow (plan $\rightarrow$ write $\rightarrow$ revise $\rightarrow$ export) into an arXiv-style twocolumn template. AI drafts were treated as scaffolding. Claims, numbers, and citations were checked against harvest records, PDF analyses, and measured artifacts. Final prose was edited for DSP terminology and to remove generic filler.

Why AI was used. (1) Scale literature triage across DSP and AutoML without pretending every high-cite hit is in-domain. (2) Accelerate IMRaD scaffolding so human effort focuses on metric honesty (residual vs DualCosine, family hardness, gate vs final overnight claims). (3) Keep an auditable tool chain (plans, PDF snippet analyses, citation flags) over opaque chat-only writing.

Where AI was *not* used. Model fitting, residual scoring, DualCosine baselines, lit-combo / overnight GPU search, and inference-time benchmarks are conventional code (Rust / PyTorch) with deterministic logs and JSON artifacts. AI did not invent experimental numbers.

Other tools. Rust/`cargo` for the procedural bench and meta binaries. PyTorch CUDA for overnight hybrid search. matplotlib for figures. MiKTeX `pdflatex` for PDF builds. GitHub for versioned papers and artifacts.

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 synthetic periods. Training and search use procedural sine+cliff generators. No private studio stems are required for the reported tables. The main resource cost is GPU time on a consumer RTX 3090 (order of hours for the 5k-gate freeze, Table 9). Authors should not over-claim speech enhancement or production mastering quality from these metrics.

Reproducibility. Holdout seed `20260719`. Overnight search seed `1902771841`. Primary code: https://github.com/reeldemo/reelsynth (`scripts/overnight_gpu_rl_arch.py`, `scripts/bench_sota_matrix.py`, `scripts/metrics_snr_sdr.py`). Paper companion: https://github.com/reeldemo/denoise-opt-meta. Dependencies for the CUDA benches: Python 3.12, PyTorch 2.6+CUDA 12.4, NumPy/matplotlib. Rust/`cargo` builds the separate `sound_bench` hardness probe. JSON artifacts under `brand/artifacts/` (`sota_matrix.json`, `canonical_eval_dataset/`, overnight `history.jsonl`) freeze the numbers in Section 5. PESQ/STOI are omitted for non-speech cycles (domain mismatch). MUSHRA was not run.

8.1 Funding

This research received no external funding.

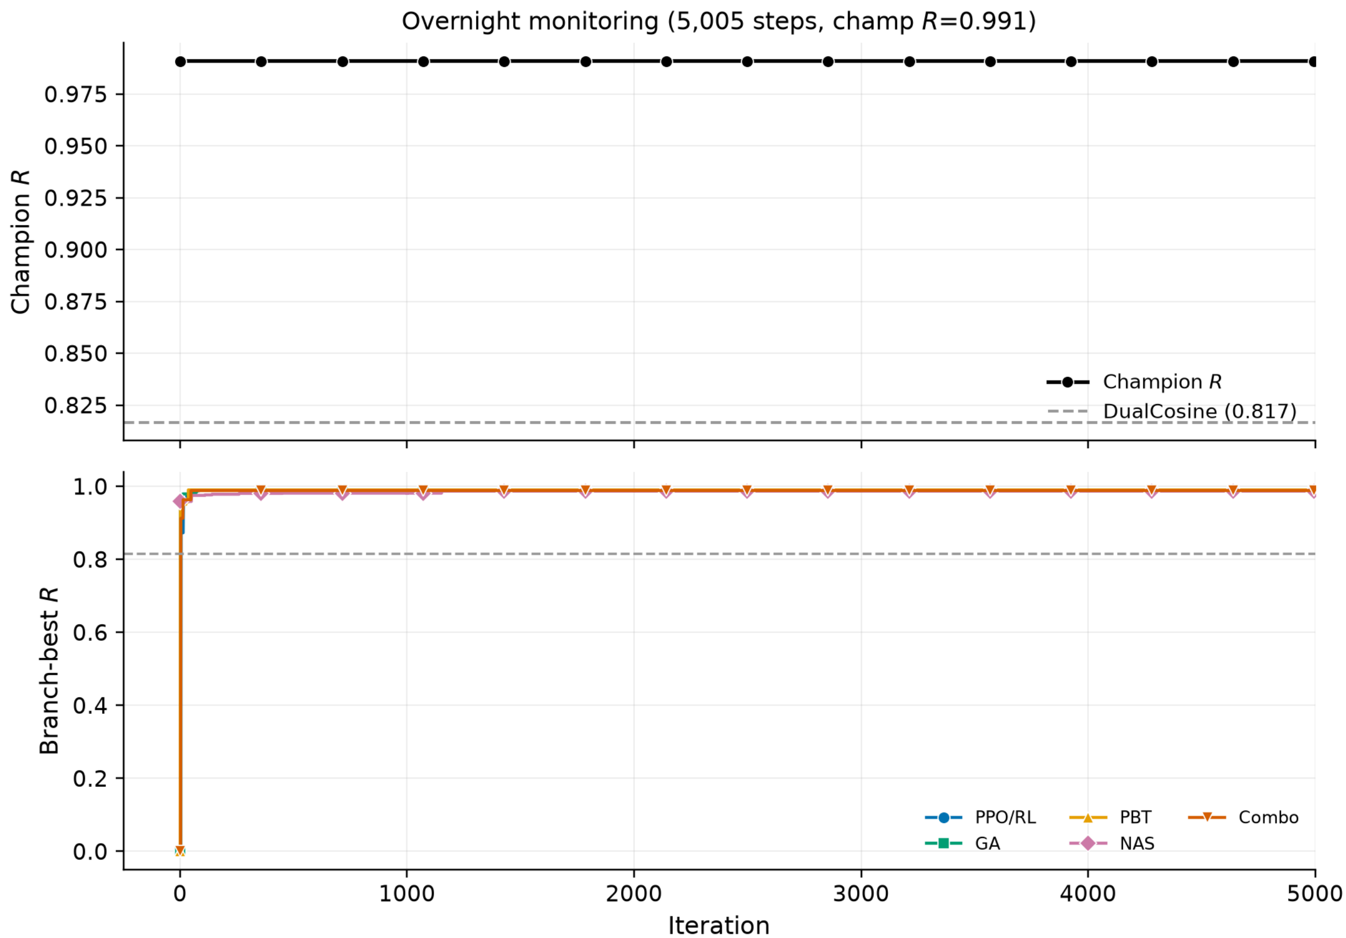

Figure 10: Full-width overnight monitoring panel at the 5k-gate freeze (first 5,000 iterations): champion R (top) and branch-best trajectories (bottom). Markers and Okabe-Ito colors match Figure 9. Source: `history.jsonl`.

Figure 11: Per-trial residual scatter by search branch for the first 5,000 iterations at the 5k-gate freeze, with champion overlay. Branch markers match Figure 9.

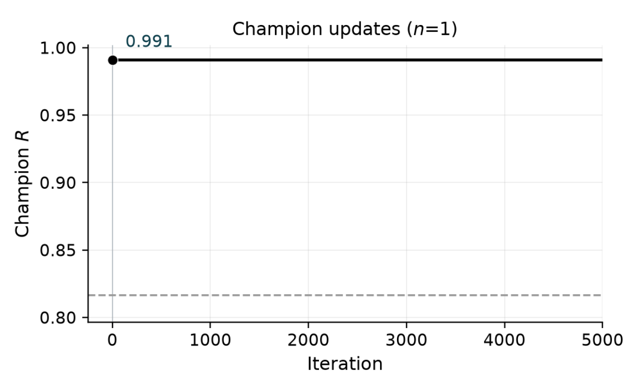

Figure 12: Champion update timeline for the first 5,000 iterations at the 5k-gate freeze. Sparse R labels mark first, mid-peak, and latest improvements for twocolumn legibility.

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, overnight history, fitted champions) are published with the companion repositories: https://github.com/reeldemo/reel synth under `brand/artifacts/`, and https://github.com/reeldemo/denoise-opt-meta under `artifacts/` and `paper/v6/figures/`. All reported tables regenerate from those files plus the scripts named above. No private studio recordings were used.

8.4 Competing interests

None declared.

9 Limitations

Scope and evaluation limits. The prolonged residual R is defined against a procedural no-cliff sibling and does not capture perceptual quality directly. Formal listening tests (MOS, MUSHRA) are outside the scope of this work and are deferred to future studies. PESQ and STOI are excluded because the primary cycles are non-speech synthetic wavetable tiles; speech-quality metrics on sine+cliff families would be a domain mismatch. MUSHRA was not run (requires human listeners). Any perceptual or speech-proxy secondary study would need separately labeled open-access speech snippets and remains out of scope for this narrow-seam paper.

- Residual uses a procedural ideal (no-cliff sibling). Perceptual MOS and listening panels are out of scope.
- Bake metrics leave formal listening tests and speech quality metrics to separate studies.
- **PESQ / STOI / MUSHRA deferred.** Primary cycles are non-speech wavetable tiles. Speech metrics would domain-mismatch sine+cliff families, so PESQ/STOI numbers are absent here (none invented on synthetic tiles). MUSHRA needs human listeners and was *not run*.
- **No LibriSpeech / MUSDB.** Wavetable-native realism uses ReelSynth-style factory periods and OA instrument-like cycles under the wrap protocol. Speech/music enhancement corpora remain off by design.
- The overnight CUDA runner scores Python `make_batch` sine+cliff draws. A separate Rust `sound_bench` export of 20 family tiles (byte-aligned generator) is a secondary transfer block. The Python generative SOTA matrix is the primary multi-family table.
- Per-trial architecture width is capped for overnight throughput. Full Demucs-/diffusion-scale stacks were screened out (domain mismatch vs $L{=}256$ periodic seams).
- We do not claim wrap-closure or hybrid-search convergence theorems. Empirical gains are protocol-bound.
- The 5k-gate freeze reports measured champion and overnight branch-best residuals. Isolated GA/PPO/GA+PPO/full re-runs use a short labeled budget (same search seed). Those short runs are a different protocol from the multi-hour freeze.
- Classical VA/BLEP anchors use author or proceedings PDFs (`stilson1996`, `nam2009polyblep`, `esqueda2016blamp`).

9.1 Outlook: follow-up publication

A self-contained follow-up can treat *which sounds fail* as the primary object of study: stratified residual heatmaps over the ten procedural families, cliff-size conditioning, and family-specialist or mixture-of-experts denoisers selected by the same RL–GA outer loop. That paper would cite the present mean-R hybrid protocol as the baseline and ask whether beating $R > 0.999$ is mainly an architecture problem or a coverage problem over hard families.

10 Conclusion

DenoiseOpt casts wavetable wrap discontinuity as cycle-local seam repair under a prolonged residual R, with a hybrid GA–PPO(+PBT) outer loop over bake cells. At the overnight freeze the champion reaches $R \approx 0.9909$ versus DualCosine ≈ 0.8166. On the frozen holdout a compact favorite reaches $R \approx 0.991$ versus DualCosine $R \approx 0.825$; across twenty multi-family draws, mean $\Delta R \approx +0.162$ (Wilcoxon $p \approx 8.9 \times 10^{-5}$). Hard-cliff strata keep favorite $R \approx 0.9915$ while DualCosine falls to ≈ 0.657. Same-generator N2N / seq baselines bound the supervised ceiling. Wavetable-native transfer covers factory and OA instrument cycles under the wrap protocol. Scope remains wavetable seam restoration for DSP / synthesis venues. Code: https://github.com/reeldemo/reelsynth, https://github.com/reeldemo/denoise-opt-meta.

Acknowledgments

Compute for the overnight CUDA campaign used a local NVIDIA GeForce RTX 3090 workstation. Open literature was retrieved via the Klaut Research Gateway under open-access PDF constraints. No external collaborators contributed text, code, or experimental numbers for this preprint.

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.