



HAL
open science

Learning to Solve Inverse Problems for Perceptual Sound Matching

Han Han, Vincent Lostanlen, Mathieu Lagrange

► **To cite this version:**

Han Han, Vincent Lostanlen, Mathieu Lagrange. Learning to Solve Inverse Problems for Perceptual Sound Matching. IEEE/ACM Transactions on Audio, Speech and Language Processing, 2024, 32, pp.2605-2615. 10.1109/TASLP.2024.3393738 . hal-04580765

HAL Id: hal-04580765

<https://hal.science/hal-04580765>

Submitted on 20 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Learning to Solve Inverse Problems for Perceptual Sound Matching

Han Han, *Student Member, IEEE*, Vincent Lostanlen, *Member, IEEE* and Mathieu Lagrange, *Member, IEEE*

Abstract—Perceptual sound matching (PSM) aims to find the input parameters to a synthesizer so as to best imitate an audio target. Deep learning for PSM optimizes a neural network to analyze and reconstruct prerecorded samples. In this context, our article addresses the problem of designing a suitable loss function when the training set is generated by a differentiable synthesizer. Our main contribution is perceptual–neural–physical loss (PNP), which aims at addressing a tradeoff between perceptual relevance and computational efficiency. The key idea behind PNP is to linearize the effect of synthesis parameters upon auditory features in the vicinity of each training sample. The linearization procedure is massively parallelizable, can be precomputed, and offers a 100-fold speedup during gradient descent compared to differentiable digital signal processing (DDSP). We demonstrate PNP on two datasets of nonstationary sounds: an AM/FM arpeggiator and a physical model of rectangular membranes. We show that PNP is able to accelerate DDSP with joint time–frequency scattering transform (JTFS) as auditory feature while preserving its perceptual fidelity. Additionally, we evaluate the impact of other design choices in PSM: parameter rescaling, pretraining, auditory representation, and gradient clipping. We report state-of-the-art results on both datasets and find that PNP-accelerated JTFS has greater influence on PSM performance than any other design choice.

Index Terms—Inverse problems, parameter estimation, synthesizer, neural network.

I. INTRODUCTION

IMITATION plays a central role in intelligent sonic interactions. In particular, the sensorimotor skills underlying vocal learning are acquired by repeated imitation of tutors. Another example is given by music: by aiming at an exact replication of a teacher’s performance, students gradually refine their level of technical control on the instrument. In both scenarios, two cognitive abilities form a closed loop: listening and search. On one hand, listening evaluates the perceptual fidelity of imitation by comparing each sonic replica with its reference. On the other hand, search gradually refines the execution of the replica at the gestural level.

A. Motivation

Endowing machines with the general-purpose ability to imitate sounds by fine motor control, as opposed to simple loudspeaker playback, has the potential to transform research protocols in robotics and human–computer interaction [1]. In this context, we call *perceptual sound matching* (PSM) the task of retrieving the input parameters of a synthesizer so as

to best imitate a digital audio target. PSM is an emerging subdomain of artificial intelligence (AI) whose application scope is rapidly expanding, thanks to the recent development of highly expressive synthesizers for speech [2], [3], music [4], [5], and bioacoustics [6].

Procedural audio synthesis is found at multiple levels of complexity: from digital audio synthesizer parameters, tangible parameters reflecting physical constructions and gestures, to polyphonic music notations for multi-player ensemble. Consider a target sound producible by a commercial synthesizer in music production and sound design: deriving the optimal set of input parameters that reproduce this sound accelerates and eases the synthesizer control [7], [8]. When the synthesizer is a physical model [9], [10], inferring the physical parameters is an interesting inverse problem that reveals the excitation mechanism as well as physical constructs of the sounding bodies, paving ways for digital musical instrument design (DMI). Even more challenging is the case of automatic orchestration [11], where parameters carry physical meaning of instrument type, playing techniques, dynamics etc.

B. Problem statement

Let us denote a target sound by the signal x , a parametric synthesizer by the function g , and its search space by the set Θ . We define PSM in terms of a function f which takes x as input and returns an element of Θ , denoted by $\tilde{\theta} = f(x)$. For an ideal PSM, the replica $\tilde{x} = g(\tilde{\theta}) = (g \circ f)(x)$ is indiscernible from the target sound x . On the contrary, failure cases of PSM arise when x and \tilde{x} differ beyond the psychoacoustic threshold of just noticeable differences: the minimal change needed for a difference to be perceived. Thus, the computational evaluation of PSM requires defining a psychoacoustic descriptor Φ such that $\Phi(x) = \Phi(\tilde{x})$ if and only if x and \tilde{x} are indiscernible. Throughout this paper, we assume the function $(\Phi \circ g)$ to be differentiable and deterministic. Given Φ , g , and Θ , the optimal PSM f appears as the solution of the following minimization problem:

$$f(x) \in \operatorname{argmin}_{\theta \in \Theta} \|(\Phi \circ g)(\tilde{\theta}) - \Phi(x)\|_2^2. \quad (1)$$

Given x , if there exists some θ such that $x = g(\theta)$, then setting $\tilde{\theta} = \theta$ in the equation above leads to $g(\tilde{\theta}) = x$ and thus $\|(\Phi \circ g)(\tilde{\theta}) - \Phi(x)\|_2^2 = 0$. Thus, θ is an optimal PSM value for $f(x)$. Yet, in full generality, no such θ exists and the optimum of Equation 1 is reached at a nonzero perceptual distance: the associated inverse problem is said to be *ill-posed*.

H. Han, V. Lostanlen, and M. Lagrange are with Nantes Université, École Centrale Nantes, CNRS, LS2N, UMR 6004, F-44000 Nantes, France. e-mail: (han.han; vincent.lostanlen; mathieu.lagrange@ls2n.fr).

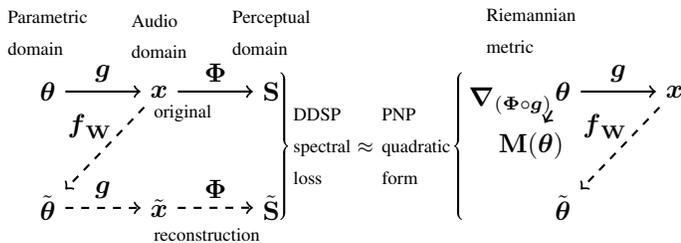


FIG. 1: Graphical outline of the proposed method. Given a known synthesizer g and feature map Φ , we train a neural network f_W to estimate θ and minimize the “perceptual–neural–physical” (PNP) quadratic form $\langle \tilde{\theta} - \theta | \mathbf{M}(\theta) | \tilde{\theta} - \theta \rangle$ where \mathbf{M} is the Riemannian metric associated to $(\Phi \circ g)$. Transformations in solid (resp. dashed) lines can (resp. cannot) be cached during training.

The search space Θ of a synthesizer is theoretically infinite; in practice, it may contain up to 10^9 elements [12]. Furthermore, the computation of a multidimensional descriptor Φ that is suitable to represent timbre dissimilarities is costly, even on GPU hardware. For these two reasons, retrieving the optimal PSM value $f(x)$ by exhaustive evaluation of $(\Phi \circ g)$ over Θ is intractable for any problem of realistic size. Instead, prior publications have resorted to iterative methods such as genetic algorithms [13] or gradient descent [14] so as to update $\tilde{\theta}$ from a random guess until it reaches convergence. Though reliable solutions may still be obtained this way, systems relying on iterative methods require step-wise reevaluation of the perceptual descriptor and new optimization steps each time given a different sound query. Moreover, given a nonconvex objective, there is no guarantee of converging to the same optimal solution across queries.

On the contrary, training a deep neural network (DNN) for PSM circumvents the aforementioned problems of uncertain and costly predictions. Denoting its weights by \mathbf{W} , the prediction of $\tilde{\theta}$ from x incurs a single forward pass: i.e., $\tilde{\theta} = f_W(x)$; the computational cost of this prediction depend solely on the DNN architecture and not on Φ nor g . Yet, training f_W to minimize $\|(\Phi \circ g \circ f_W)(x) - \Phi(x)\|_2^2$ remains constrained by the cost of $(\Phi \circ g)$. Our article offers a generic solution to mitigate that constraint.

C. Key idea

We propose a novel loss objective termed “perceptual–neural–physical” (PNP) loss. PNP allows working with computationally heavy Φ that would otherwise induce intractable training time, under the conditions where Φ and g are differentiable. It is also required that θ be known to supervise the training with PNP loss.

PNP is the optimal quadratic approximation of a given L^2 perceptual loss. It adopts a bilinear form whose kernel matrix is the Riemannian metric formed by the differentiable map $(\Phi \circ g)$. These kernel matrices may be precomputed and cached in memory before training. Furthermore, we propose a corrective term added onto the kernel matrix, as a mechanism to reduce the condition number of the kernel matrix in case where the inverse problem is ill-conditioned.

D. Contributions

We train models against various state-of-the-art loss functions to perform sound matching tasks on two differentiable synthesizers - an AM/FM arpeggiator and a rectangular drum physical model. The former is a pure tone chirp signal with amplitude-modulated temporal envelope, an extension to our previous work [15]. The latter is an inharmonic transient signal with uneven exponential decay envelope on each partial, which better reflects sounds in the real world.

We demonstrate experimentally that optimizing the DNN against our proposed PNP loss based on JTFS consistently achieves the best performance on both sound matching tasks. Meanwhile we show that the tasks are nontrivial given the poor performance with previous state of the art method using MSS spectral loss. Finally, we extend [15] by presenting the learning dynamics and an ablation study of our proposed model to discuss the key factors contributing to empirical success (or failure). We supplement this work with open-source code.¹

II. RELATED WORK

A. Differentiable Synthesizers

Synthesizers have longstanding tradition in digital music production. They range in design principles but mostly fall into a few families - subtractive, additive, frequency modulation, and wavetable synthesis. Using filters to amplify or diminish certain frequency components, using oscillators to modulate another oscillator’s pitch, phase, and amplitude envelope, digital synthesizers elicit a rich variety of timbres. Powerful as they are, the concept of differentiable synthesizers only came in the advent of differentiable digital signal processing (DDSP) model [16]. DDSP models are encoder-decoder architectures capable of generating high-fidelity audio by passing learned encodings through expressive synthesizers. Optimizing against a perceptual loss in the form of Equation. 1, DDSP model’s parameter updates via back propagation depend on computing $\nabla_{(\Phi \circ g)}(\theta)$. This computation is made easy by automatic differentiation (AD), the automatic evaluation of compounds of partial derivatives via the chain rule, during forward pass through the function $(\Phi \circ g)$. Consequently, in order to be compatible with AD, the forward function composes differentiable functions.

To enable learning with perception-based objectives, the recent years have seen a surge of newly implemented differentiable synthesizers. Such an implementation is straightforward for certain synthesizers, ones that can be expressed as a cascade of linear or nonlinear continuous functions. Additive synthesizers [17], subtractive synthesizers [18], harmonic plus noise synthesis [16], FIR and IIR filters with closed form impulse responses, all fall under this category.

Other synthesizers with discrete elements, thresholding or nonlinear analog effects may be adapted into differentiable implementations as well. Facing the infinite discrete choices of routing and patching in FM synthesis, [7] implemented a differentiable six-oscillator FM synthesizer with fixed patching for each target musical instrument, in order to prescribe

¹Companion website: https://github.com/lylyhan/perceptual_neural_physical

harmonic spectra and ease the retrieval of continuous control parameters to regenerate instrument sounds. Time-varying nonlinear effects particular to analog audio devices that lack an analytic expression have also been modeled by learned neural net, for example phasor effects in [19]. Other interests lie in enabling the differentiability of synthesis based on physical models, due to its physical interpretability and proximity to realistic sounds. Commonly involving solving a partial differential equation, such synthesis algorithms grow increasingly sophisticated and time-consuming as the underlying physical model increases in dimensionality of solution space and complexity of initial conditions. [20] and [21] trained respectively recurrent and convolution-based neural network to model time-stepping solution from finite difference method and resonator responses from finite element method. We refer to [22] for a comprehensive review of differentiable synthesizers.

Whether through direct conversion or data-driven simulation via learned filters, enabling the differentiability of synthesis algorithms opens up opportunities for optimizations informed by audio samples. Nevertheless it is important to note that, in face of sophisticated synthesis algorithms g and cumbersome perceptual representations Φ , efficient methods are needed to address the computational and memory overload during training.

B. Perceptually motivated time–frequency representations

Spectrotemporal modulations (STM) are multiresolution, oriented oscillation rates in the time–frequency domain. As one of the evidences that identify STMs as the underpinning for auditory perception, it is shown in neurophysiological research that systematic degradation of STMs in complex signals such as speech and music, are correlated with a loss of intelligibility [23]. Relatedly, spectrotemporal receptive fields (STRF) reflect each mammal’s selectivities to STMs of differing scales (amplitude modulation) and rates (frequency modulation). They are measurable through exposing the mammals’ primary auditory cortex (A1) under systematically varied, parametrized auditory stimuli, such as ripples [24]. Based on these experimental results, [25] among others have proposed computational idealizations of the A1 processing aiming to simulate STRF. These models translate acoustic signals to basic aspects of physiological responses, paving way towards constructing a computational proxy for indicating sonic differences. More recently, [26] applied metric learning onto STM model responses and demonstrated its correlation to similarity judgements of musical sounds.

Despite STRF’s physical relevance, its numerical representation is cumbersome and RAM-intensive. It is also challenging to interface with common deep learning framework as a learning objective, due to a lack of differentiable implementation. One alternative approach is the joint time–frequency scattering transform (JTFS) proposed in [27]. JTFS is an operator that analyzes spectrotemporal modulation via joint wavelet convolutions, nonlinearities, and low-pass filters. Despite being close to STRF, its design results in more lightweight and efficient perceptual encodings without compromising the comprehensiveness of the captured information.

JTFS succeeds as audio feature in various state-of-the-art classification tasks thanks to its time-shift invariance property, and is able to recover complex audio signals via gradient descent [27], an indication of its high efficiency and preservation of information resolution. Meanwhile, in the spirit of [26], [28] reported strong correlation between the Euclidean distance in JTFS coefficients, and human subjects’ similarity judgements on musical instruments’ playing techniques. Recently, [29] published a differentiable implementation of JTFS, making it possible to be adopted as a learning objective.

C. P-loss and Differentiable Spectral Loss

Given a DNN $f_{\mathbf{W}}$ to retrieve $\tilde{\theta}$ from x , choosing the training objective largely influences learning dynamics and in turn the converged solution.

When neither Φ nor g is differentiable, a common approach termed parameter loss, or “P-loss” for short, optimizes $\|\theta - \tilde{\theta}\|^2$ the mean squared error between the ground truth and estimated parameters. P-loss is a fast alternative as it circumvents the backward pass through Φ and g . Yet, it only finds perceptually acceptable solution under the premise that $\|\tilde{\theta} - \theta\| \propto \|(\Phi \circ g)(\tilde{\theta}) - (\Phi \circ g)(\theta)\|$; that is, a proximity in parameter space is equivalent to a proximity in the perceptual space. Indeed, parameter and perceptual loss surface share a common global minimum \mathbf{W}^* such that $f_{\mathbf{W}^*}(x) = \theta$. However, the smoothness of the error surfaces should also align to converge to similar local minima [30]. In practice, a certain operator h is often applied onto θ to better approximate the relative distance in a perceptual dimension. Common operations include min–max scaling, logarithmic transformation, and other scaling methods based on known ranges or distributions of the parameter dimensions. The goal is to have the error in reparameterized space vary consistently as the perceptual discrepancy. These domain specific knowledge should also condition the sampling of datasets to ensure the success of a data-driven approach [13].

The benefit of P-loss is not limited to its role as an inexpensive surrogate to perceptual loss. Being a convex objective in the parameter space, convergence of \mathbf{W} is easily achieved when parameters are normalized into similar ranges. On the contrary, perceptual loss does not guarantee this convergence. A typical spectral representation used in DDSP is the multi-scale spectrogram distance (MSS). Given a target sound x and estimated sound $\tilde{x} = (g \circ f_{\mathbf{W}})(x)$, MSS² is defined by:

$$\mathcal{L}_x^{\text{MSS}}(\mathbf{W}) = \sum_{k=5}^{10} \|\Phi_{\text{STFT}}^k(x) - \Phi_{\text{STFT}}^k(\tilde{x})\|_2^2, \quad (2)$$

where Φ_{STFT}^i refers to short-time Fourier transforms computed with window sizes of 2^k . Using $\mathcal{L}_x^{\text{MSS}}$ as spectral loss is computationally inexpensive, does not require the presence of ground truth θ , and more importantly reflects distance between harmonic and stationary sounds of the same pitch. However, its adequacy is compromised when comparing sounds with

²The original definition of MSS exist combined an L^1 norm between magnitudes and an L^1 norm between log-magnitudes. However, we choose to retain the L^2 norm instead, for consistency with Equation 1 and 3.

different pitches [31], inharmonicity [15] or nonstationarities [14]. When representing these sounds, MSS no longer accurately distinguishes large perceptual differences, and therefore fails to provide informative gradient vectors.

Due to the complementary nature between the above two approaches, choosing the right learning objective requires careful consideration on the characteristics of sounds and preprocessing of parameter space in each specific application. Encouragingly, attempts at merging the above two approaches have shown to be successful [17]. Nevertheless, these approaches usually require separate pre-training and fine-tuning stages, introducing auxiliary hyperparameters in the training scheme.

III. METHODS

A. Taylor expansion

Given a spectral representation Φ and a synthesizer g , a DDSF-style L^2 spectral loss can be expressed as

$$\begin{aligned} \mathcal{L}_\theta^{\text{DDSP}}(\mathbf{W}) &= \|\Phi(x) - \Phi \circ g \circ f_{\mathbf{W}}(x)\|_2^2 \\ &= \|(\Phi \circ g)(\theta) - (\Phi \circ g)(\tilde{\theta})\|_2^2 \end{aligned} \quad (3)$$

By relying on the assumption that the function $(\Phi \circ g)$ is continuously differentiable, we perform its second-order Taylor expansion around $\tilde{\theta}$ and obtain

$$\begin{aligned} (\Phi \circ g)(\tilde{\theta}) &= (\Phi \circ g)(\theta) + \nabla_{(\Phi \circ g)}(\theta) \cdot (\tilde{\theta} - \theta) \\ &\quad + O(\|\tilde{\theta} - \theta\|_2^2), \end{aligned} \quad (4)$$

where the Jacobian matrix $\nabla_{(\Phi \circ g)}(\theta)$ has $P = \dim(\Phi \circ g)(\theta)$ columns and $J = \dim \theta$ rows. From Equation (4), we deduce³:

$$\mathcal{L}_\theta^{\text{DDSP}}(\mathbf{W}) = \langle \tilde{\theta} - \theta | \mathbf{M}(\theta) | \tilde{\theta} - \theta \rangle + O(\|\tilde{\theta} - \theta\|_2^3) \quad (5)$$

where, for every $\theta \in \Theta$,

$$\mathbf{M}(\theta) = \nabla_{(\Phi \circ g)}^\top(\theta) \nabla_{(\Phi \circ g)}(\theta) \quad (6)$$

is a positive semidefinite matrix with J rows and J columns. The optimal quadratic approximation of $\mathcal{L}_\theta^{\text{DDSP}}$ near \mathbf{W} is

$$\mathcal{L}_\theta^{\text{PNP}}(\mathbf{W}) = \langle \tilde{\theta} - \theta | \mathbf{M}(\theta) | \tilde{\theta} - \theta \rangle \quad (7)$$

The main interest in PNP resides in its computational efficiency. Since Φ is computationally expensive, gradient backpropagation through $\mathcal{L}_\theta^{\text{DDSP}}$ is impractically slow. According to the chain rule, optimizing against conventional spectral loss requires evaluating at every iteration the gradient of each path in the spectral representation \mathbf{S}_p with respect to each neural net weight scalar \mathbf{W}_i . Using PNP, however, we decompose the above gradient computation into the matrix multiplication of two Jacobian matrices of $P \times J$ and $J \times |W|$ entries.

$$\nabla_{\Phi \circ g \circ f_{\theta}}(\mathbf{W}) = \nabla_{(\Phi \circ g)}(\tilde{\theta}) \nabla_{f_{\theta}}(\mathbf{W}) \quad (8)$$

We precompute only once $\nabla_{(\Phi \circ g)}(\theta)$. Note that in practice, $J^2 \ll JP$: spectral representation lies in much higher dimension than synthesis parameters. Therefore, despite the high dimensionality of $\nabla_{(\Phi \circ g)}(\theta)$, its inner product $\mathbf{M}(\theta)$ is lightweight to be stored in memory and loaded on the fly while training.

³The bracket notation $\langle y | \mathbf{A} | y \rangle$ is equivalent to $y^\top \mathbf{A} y$.

B. Riemannian geometry

$\mathbf{M}(\theta)$ defines an inner product on the tangent space of $(\Phi \circ g)$ at each point θ . Alternatively viewed, it is the Riemannian metric induced by the smooth manifold composing differentiable map $(\Phi \circ g)$ and the open set $\Theta \subset \mathbb{R}^J$. Such a metric allows for measuring properties such as curvature, volume, and length in the manifold space. Using Riemannian metric to investigate and exploit information geometry in a manifold space is not new. In this context, a relevant prior publication in [32], used the inverse determinant of Riemannian metric to infer data importance and inform geometry-aware data sampling.

Let $\sigma_j^2(\theta)$, $v_j(\theta)$ be the paired eigenvalues and orthonormal eigenvectors of $\mathbf{M}(\theta)$. For each v_j , one has: $\mathbf{M}(\theta)v_j(\theta) = \sigma_j(\theta)^2 v_j(\theta)$. Hence, Equation (7) may be rewritten as

$$\mathcal{L}_\theta^{\text{PNP}}(\mathbf{W}) = \sum_{j=1}^J \sigma_j^2(\theta) |\langle \tilde{\theta} - \theta | v_j(\theta) \rangle|^2 \quad (9)$$

Indeed, if $\mathbf{M}(\theta)$ is an identity matrix, $\sigma_j(\theta)^2 = 1$, $v_j(\theta)$ is the standard basis, and Equation (9) falls back to P-loss. Yet, a Riemannian metric $\mathbf{M}(\theta)$ projects parameter error vectors onto $v_j(\theta)$, scale by $\sigma_j(\theta)^2$, and recombine them to form a loss objective which takes into consideration the non-Euclideanness of the multidimensional perceptual space.

C. Levenberg-Marquardt algorithm

In practice, $\mathbf{M}(\theta)$ is not necessarily positive definite. Under cases where θ has interdependent dimensions, or when certain dimensions of the manifold space $(\Phi \circ g)$ are invariant to changes in θ , $\mathbf{M}(\theta)$ is rank deficient. Such a condition in the Riemannian metric poses a serious problem, as implied in the aforementioned metrics of inverse determinant: the determinant would collapse to zero, the metrics would blow up to infinity, implying an explosion of information containment. Consequently even though $\mathcal{L}_\theta^{\text{PNP}}$ remains strictly non-negative, designating it as a loss function faces a challenge. Since $\mathbf{M}(\theta)$ is positive semidefinite, a Cholesky decomposition leads to $\mathbf{M}(\theta) = \mathbf{A}(\theta)^\top \mathbf{A}(\theta)$. It follows that minimizing Eq. 7 is equivalent to minimizing $\|\mathbf{A}(\theta)(\tilde{\theta} - \theta)\|_2^2$. The corresponding ordinary least squares (OLS) estimator via Gauss-Newton Method is

$$\begin{aligned} \tilde{\theta} &= (\mathbf{A}(\theta)^\top \mathbf{A}(\theta))^{-1} \mathbf{A}(\theta)^\top \mathbf{A}(\theta) \theta \\ &= \mathbf{M}^{-1}(\theta) \mathbf{A}(\theta)^\top \mathbf{A}(\theta) \theta. \end{aligned} \quad (10)$$

Here, the inversion of $\mathbf{M}(\theta)$ is ill-conditioned when $\mathbf{M}(\theta)$ has a high condition number

$$\kappa(\theta) = \frac{\max_{1 \leq j \leq J} \sigma_j(\theta)}{\min_{1 \leq j \leq J} \sigma_j(\theta)}. \quad (11)$$

More precisely, a rank deficient matrix has at least one eigenvalue of zero, rendering its condition number to infinity and its inverse undefined. In the Levenberg-Marquardt algorithm, a regularization term is added to the matrix before inversion to avoid this potential instability.

$$\tilde{\theta} = (\mathbf{A}(\theta)^\top \mathbf{A}(\theta) + \lambda \mathbf{I})^{-1} \mathbf{A}(\theta)^\top \mathbf{A}(\theta) \theta \quad (12)$$

Each eigenvalue of \mathbf{M} increases by the corresponding eigenvalue of the added diagonal matrix, in this case a constant factor λ . Consequently, the condition number κ decreases and a numerically stable inverse can be found. We take inspiration from LMA and adopt a similar corrective term approach to improve PNP. We define:

$$\begin{aligned} \mathcal{L}_{\tilde{\theta}, \lambda}^{\text{PNP}}(\mathbf{W}) &= \langle \tilde{\theta} - \theta | \mathbf{M}(\theta) + \lambda \mathbf{I} | \tilde{\theta} - \theta \rangle \\ &= \langle \tilde{\theta} - \theta | \mathbf{M}(\theta) | \tilde{\theta} - \theta \rangle + \lambda \|\tilde{\theta} - \theta\|^2 \end{aligned} \quad (13)$$

The corrective term simultaneously shifts the eigenvalues σ_j^2 by λ thereby improving the condition number of \mathbf{M} , and acts as a L^2 regularization term for the inverse problem. In the limit of $\lambda \rightarrow \infty$, λ balances the magnitudes of σ_j^2 such that (13) approximates a P-loss regime.

Similar to LMA, we initialize λ at a large value $\lambda_{\max} = \max_{1 \leq j \leq J} \sigma_j^2(\theta)$, the largest eigenvalue of all $\mathbf{M}(\theta)$. We iteratively reduce it whenever epoch validation loss decreases, until reaching a low threshold (typically zero). This way, the optimization starts in the P-loss regime and gradually transitions to the perceptual loss regime. Another reason why this damping is beneficial is that the higher order error term in (4) ceases to be negligible when $(\tilde{\theta} - \theta)$ is large. As the DNN weights are randomly initialized, $\tilde{\theta}$ is not guaranteed to approximate θ at the beginning of training. The network initialization is therefore crucial to the reliability of such a spectral loss approximation.

IV. APPLICATION

A. Joint time–frequency scattering transform (Φ)

Let $\psi(t)$ be the mother wavelet, a complex-valued zero-average bandpass filter. Through dilating and scaling ψ in time by a factor of 2^λ , we may construct a filterbank of multiscale wavelets with a constant quality factor, i.e. constant ratio between their center frequencies and bandwidths

$$\psi_\lambda(t) = 2^\lambda \psi(2^\lambda t), \quad (14)$$

where $\lambda \in \mathbb{R}$ are the log-frequency variables of the wavelet filterbank. Its center frequencies are logarithmically spaced in the hearing range with Q number of filters per octave, and J number of octaves. Here, we set $Q = 12$. Given a signal x , we first obtain the modulus scalogram after complex convolution with $\psi_\lambda(t)$, then we apply a low-pass filter $\phi_T(t)$ of time support T which introduces invariance to time shifts within $\tau \leq T$. This results in the first-order scattering coefficients $\mathbf{S}_1 \mathbf{x}$.

$$\mathbf{S}_1 \mathbf{x} = |x * \psi_\lambda(t)| * \phi_T(t) \quad (15)$$

To recover the high frequency information lost in the time-averaging operation and further capture spectrotemporal modulations in the signal, we apply a second wavelet transform. We define two separate filterbanks $\psi_\mu(t)$ and $\psi_l(\lambda)$ in the time and log-frequency domain respectively. Their center frequencies refer roughly to the modulation scale and rate in STRF [25]. We adopt the 2D wavelets resulting from the outer product between them:

$$\psi_{\mu, l, s} = \psi^{(t)}(2^\mu t) \psi^{(f)}(s 2^l \lambda) \quad (16)$$

, where $s = \pm 1$ is the “spin” parameter indicating the upward or downward directionality of the wavelet. We take $Q = 1$ for the second order wavelet filterbanks, as the modulation frequency decomposition does not require the same frequency resolution. The second order coefficients are obtained in the same way after modulus nonlinearities and timeaveraging with a spectrotemporal low pass filter.

$$\mathbf{S}_2 \mathbf{x} = |\mathbf{S}_1 \mathbf{x} * \psi_{\mu, l, s}(t, \lambda)| * \phi_{T, F}(t, \lambda) \quad (17)$$

The convolution along frequency axis makes it robust to preserve frequency-dependent temporal structures and time–frequency warping even after the compulsory low pass filter for time-shift invariance. Finally, JTFS concatenates $\mathbf{S}_1 \mathbf{x}$ and $\mathbf{S}_2 \mathbf{x}$ to form the final time-invariant representation $\text{JTFS}(\mathbf{x}) = [\mathbf{S}_1 \mathbf{x}, \mathbf{S}_2 \mathbf{x}]$.

Previously demonstrated to improve discriminability among signals with small scattering coefficients in classification tasks [29], we further apply a stabilized logarithmic transformation:

$$\Phi_{\text{JTFS}}(\mathbf{x}) = \log \left(1 + \frac{\text{JTFS}(\mathbf{x})}{\varepsilon} \right), \quad (18)$$

where we have set the hyperparameter ε to 10^{-3} .

B. AM/FM arpeggiator (g)

Nonstationary signals governed by regularities in the mesostructural level are challenging to be effectively distinguished by first-order time–frequency representation focusing on local variations [14]. The AM/FM arpeggiator is an amplitude-modulated ascending chirp signal:

$$g(\theta) = \sin \left(\frac{2\pi f_c}{\gamma \log 2} (2^{\gamma t} - 1) \right) \sin(2\pi f_m t) \phi(\gamma t), \quad (19)$$

where $\theta = \{\log f_c, \log f_m, \log \gamma\}$.

Modified by chirp rate γ in octaves per second, a carrier signal sweeps exponentially over the frequency interval $[f_c, f_c 2^{\gamma T}]$ in a period of time T . Then, a slow oscillation sinusoid of frequency f_m in Hz is multiplied to modulate the amplitude envelope. Finally, we apply a Gaussian window of width $0.2/\gamma$ in seconds onto the signal.

To construct the dataset, we sample each parameter from their respective intervals: $f_c \in [512, 1024]$ Hz, $f_m \in [4, 16]$ Hz and $\gamma \in [0.5, 4]$ Hz, and generate sounds at a sampling rate of 2^{13} Hz. Specifically, we divide each dimension of the 3-D cube formed by these intervals into 30 equal steps in the logarithmic space, making $30^3 = 27 \cdot 10^3$ three-dimensional grids. Each audio sample is thus randomly sampled from a point in each grid. To ensure varied distribution in the training and validation set, we take the middle $14^3 \approx 2.7 \cdot 10^3$ grids as the validation set. The training and testing sets remain in the same distribution at the peripheral of the sampling cube. The train/test/val split is 8 : 1 : 1.

C. Drum sound synthesizer (g)

Another type of nonstationary signals can be exemplified by inharmonic transient sounds, such as those coming from a percussive instrument. We synthesize percussive sounds via physical model of a 2D membrane residing in the Cartesian

coordinate system \mathbf{u} , where a random position's vertical displacement $\mathbf{X}(t, \mathbf{u})$ is governed by a fourth-order damped partial differential equation. For $t \geq 0$,

$$\left(\frac{\partial^2 \mathbf{X}}{\partial t^2}(t, \mathbf{u}) - c^2 \nabla^2 \mathbf{X}(t, \mathbf{u}) \right) + S^4 (\nabla^4 \mathbf{X}(t, \mathbf{u})) + \frac{\partial}{\partial t} \left(d_1 \mathbf{X}(t, \mathbf{u}) + d_3 \nabla^2 \mathbf{X}(t, \mathbf{u}) \right) = 0 \quad (20)$$

The higher order spatial derivatives and first-order time derivative account for materiality, air-induced frictions and energy dissipation respectively, introducing in the resulting solution inharmonicities and uneven decay envelopes among partials. Furthermore, the characteristics of such a solution are parameterized by material stiffness S , traveling wave speed c , frequency-independent damping coefficient d_1 and frequency-dependent damping coefficient d_3 . For compatibility with the adopted coordinate system, we focus on a rectangular drum of side lengths l by $l\alpha$ with its boundaries at zero at all time. We simulate the excitation by setting the initial condition at $t_0 = 0$ to be $\mathbf{X}(t_0, \mathbf{u}) = (0.4l, 0.4l\alpha) = 0.03$ meters, and 0 otherwise.

We adopt the functional transformation method (FTM) [33], a solver to the above PDE, as our synthesizer $\mathbf{g}(\boldsymbol{\theta})$, where $\boldsymbol{\theta} = \{S, c, d_1, d_3, \alpha\}$. FTM disentangles time-space solution by performing Laplace and Sturm-Liouville transforms and thereby transforming the variables into their respective functional spaces. There, algebraic solutions in the functional space domain are obtained and then inverse transformed back into the time-space domain. The solution follows a modal synthesis form:

$$\mathbf{x}(t) = \mathbf{X}(t, \mathbf{u}) = \sum_{\mathbf{m} \in \mathbb{N}^2} K_{\mathbf{m}}(\mathbf{u}, t) \exp(\sigma_{\mathbf{m}} t) \sin(\omega_{\mathbf{m}} t). \quad (21)$$

The coefficients $K_{\mathbf{m}}(\mathbf{u}, t)$, $\sigma_{\mathbf{m}}$, $\omega_{\mathbf{m}}$ are respectively modal amplitudes, modal exponential decay rates, and modal frequencies, where $\mathbf{m} = (m_1, m_2)$. They are derived from the original PDE parameters in the following ways.

$$\omega_{\mathbf{m}}^2 = \left(S^4 - \frac{d_3^2}{4} \right) \Gamma_{\mathbf{m}}^2 + \left(c^2 + \frac{d_1 d_3}{2} \right) \Gamma_{\mathbf{m}} - \frac{d_1^2}{4} \quad (22)$$

$$\sigma_{\mathbf{m}} = \frac{d_3}{2} \Gamma_{\mathbf{m}} - \frac{d_1}{2} \quad (23)$$

$$K_{\mathbf{m}}(\mathbf{u}, t) = y_u^{\mathbf{m}} \delta(t) \sin\left(\frac{\pi m_1 u_1}{l}\right) \sin\left(\frac{\pi m_2 u_2}{l\alpha}\right) \quad (24)$$

where $\Gamma_{\mathbf{m}} = \pi^2 m_1^2 / l^2 + \pi^2 m_2^2 / (l\alpha)^2$, and $y_u^{\mathbf{m}}$ is the \mathbf{m}^{th} coefficient associated to the eigenfunction $\sin(\pi m \mathbf{u} / l)$ that decomposes $y_u(\mathbf{u})$. Lastly, to relate better input parameters to sonic characteristics, we reparameterize $\boldsymbol{\theta} = \{S, c, d_1, d_3, \alpha\}$ to $\boldsymbol{\theta} = \{\log \omega_1, \tau_1, \log p, \log D, \alpha\}$ in the following way:

$$\omega_1 = \sqrt{\frac{\beta^4}{\alpha^2} S^4 + \frac{\beta}{\alpha} c^2 - \frac{1}{4} \left(\frac{\beta}{\alpha} d_3 - d_1 \right)^2} \quad (25)$$

$$\tau_1 = \frac{2}{d_1 - \frac{\beta}{\alpha} d_3} \quad (26)$$

$$p = \frac{d_3}{\beta d_3 - \alpha d_1} \quad (27)$$

$$D = \frac{1}{\alpha \omega_1} \sqrt{S^4 - \frac{d_3^2}{4}} \quad (28)$$

Min-Max	Log Transform	Activation
✓	✓	Tanh
✓	×	Tanh
×	✓	None
×	×	LeakyReLU

TABLE I: Activation functions used by the last linear layer of $\mathbf{f}_{\mathbf{w}}$ under different reparametrizations

where $\beta = 1 + 1/\alpha$. For simplicity, we denote the first mode in both dimension $m_1 = m_2 = 1$ as $\mathbf{m} = 1$. Fundamental frequency ω_1 ranges between 40 Hz and 1kHz; duration τ_1 , between 400 ms and 3 s; inhomogeneous damping rate p , between 10^{-5} and 0.2; frequential dispersion D , between 10^{-5} and 0.3; and aspect ratio α , between 10^{-5} and 1. These predefined ranges motivated by perception can be converted back to the PDE parameters in the following way:

$$S = \left((D\omega_1\alpha)^2 + \left(\frac{p\alpha}{\tau_1} \right)^2 \right)^{\frac{1}{4}} \quad (29)$$

$$c = \sqrt{\frac{\alpha}{\tau_1^2} \left(\frac{1}{\beta} - p^2 \beta \right) + \alpha \omega_1^2 \left(\frac{1}{\beta} - D^2 \beta \right)} \quad (30)$$

$$d_1 = \frac{2}{\tau_1} (1 - p\beta) \quad (31)$$

$$d_3 = -\frac{2p\alpha}{\tau_1} \quad (32)$$

The dataset is constructed the same way as described in section IV-B. We take $m_1 = m_2 = 10$ and synthesize ~ 3 s sounds at a sampling rate of 22050 Hz. The hyperdimensional cube is divided into 10 equal steps per parameter dimension, forming 100k samples in total. The train/test/validation split is 8 : 1 : 1.

D. Reparametrization ($\boldsymbol{\theta}$)

To retrieve $\boldsymbol{\theta}$, the decoder $\mathbf{f}_{\mathbf{w}}$ commonly ends with a linear layer of the same size activated by certain function. Scaling of ground truth parameters is a key factor in designating such an activation function and in turn the ease of optimization. In this work, we examine the effect of two modes of reparametrization on the neural net's learning dynamics - min-max scaling and logarithmic transformation.

Min-max scaling is a common technique to normalize parameters into a confined range: typically, $[0, 1]$ or, as in our case, $[-1, 1]$.

$$h_{\text{minmax}}(\boldsymbol{\theta}) = -1 + 2 \frac{\boldsymbol{\theta} - \boldsymbol{\theta}_{\min}}{\boldsymbol{\theta}_{\max} - \boldsymbol{\theta}_{\min}} \quad (33)$$

Logarithmic transformation $h_{\log}(\boldsymbol{\theta}) = \log(\boldsymbol{\theta})$ applies log operations onto parameters that influence perception in an approximately logarithmic scale and or have too small ranges. In our case, these parameters are f_c , f_m , and γ for the AM/FM arpeggiator; and ω_1 , p , and D for the drum synthesizer.

Both scaling operations above are considered as domain knowledge that modify the optimization landscape such that a more stable convergence is achieved. To finetune an output vector of entries lying or affecting the loss objective in drastically different ranges, the Jacobian matrix $\nabla_{(\Phi \circ \mathbf{g})}(\tilde{\boldsymbol{\theta}})$ needs to produce gradient updates with magnitudes of the according

sensitivities. This is hard to achieve as, in deep learning, the commonly adopted first-order optimizers do not consider curvature of the optimization landscape. This poses a risk of perpetually improving one parameter’s accuracy at the expense of detrimenting the other. Furthermore, optimizing at once parameters of different ranges is problematic for unweighted loss functions. Consider optimizing \mathbf{f}_W against P-loss, an error in the parameter with large range contributes more significantly to the overall loss, while those with small ranges have minimal effect. Consequently the network naturally tends to prioritize reducing errors in the largest parameter.

We apply none, one of or both scaling methods to the ground truth parameters, while adapting the final layer accordingly as specified in Table I.

E. Learned Decoder (\mathbf{f}_W)

A suite of DNN architecture known as EfficientNet [34], uses compound scaling method to find the optimal balance for widths, depths and input resolutions of subsequent convolutional blocks. The network achieved state-of-the-art classification result on both image [34] and audio tasks [29] using significantly less parameters. We adopt the most light weight version EfficientNet-B0 as our encoder \mathbf{f}_W architecture, summing to 4M trainable parameters. The network takes in constant Q transform (CQT) coefficients of each audio sample, and outputs a vector of dimension $\dim \theta$. The final linear layer vary in its activation functions according to the chosen reparametrization, specified in Table I. We leave to future work the question of scaling our experiment benchmark beyond B0 and up to B7.

V. EXPERIMENTS

A. Competing Approaches

For each sound synthesizer, we train \mathbf{f}_W against 3 loss objectives: P-loss as [35], $\mathcal{L}_{\theta,\lambda}^{\text{PNP}}$ based on Φ_{JTFS} with adaptive damping as [15], and MSS as [16]. Each network is exposed to 25600 samples per epoch for a total of 70 epochs. For P-loss and PNP loss we train with a batch size of 256 samples. For MSS loss, we need to reduce the batch size to 64 samples due to the memory constraints of $\Phi \circ \mathbf{g}$. To eliminate the effect of differing batch sizes on experimental results, we also trained our models using P-loss and PNP loss with a batch size of 64 samples. As these supplementary experiments did not indicate significant degradation nor narrowed the advantageous leap from the results of MSS loss, we choose not to include them in Table II. In making our best efforts to replicate [16], [17], all above experiments are trained with Adam optimizer and a scheduled learning rate initialized at 0.001 that reduces by a factor of 0.1 whenever validation loss plateaus for three epochs. Finally in the spirit of [17], we propose to pretrain with P-loss for 35 epochs, and finetune with $\mathcal{L}_{\theta,\lambda}^{\text{PNP}}$ for 35 epochs, using weight decay and gradient clipping in addition to the above optimizer setup. The full algorithm is written explicitly in Algorithm 1. The ground truth parameters are pre-processed with both min-max scaling and logarithmic transformation, unless specified otherwise.

B. Evaluation

We adopt two metrics to evaluate our experiments: MSS and JTFS distance. Both being perceptual metrics, the former assesses local frequency component alignments at various time scales. The latter examines additional differences in frequency and temporal modulation rates. For each trained model, we report both evaluation metrics averaged over the entire test set in Table II.

C. Discussion

Our proposed model consistently outperforms the other methods on both metrics. P-loss and $\mathcal{L}_{\theta,\lambda}^{\text{PNP}}$ also converge stably to suboptimal solutions. On the contrary, MSS fails to induce consistent gradient updates that retrieve accurate parameters, even under the help of improved model initialization.

1) *Perceptual distance between nonstationary sounds:* The large discrepancy between using MSS versus JTFS-based loss objectives in retrieving parameters from nonstationary sounds is expected. It is also the pivotal factor contributing to the success of accurate parameter retrieval, as shown in the ablation study of Figure 2.

First, prior work such as [31], [36] have shown that spectrogram distances are inaccurate in representing distance between sinusoids of different frequencies. This knowledge manifests in the difficulties encountered during retrieving f_c from AM chirp signals and $\log \omega_1$ from FTM percussive sounds. In particular, [15] has shown how supplying ground truth frequency aids the retrieval of other parameters.

Second, AM chirp signals are governed by long-range periodicity in time (f_m) otherwise referred to as the “meso-structures”, which are imperceptible by first order time-frequency analysis with short time support. [14] has visualized the loss surfaces and gradient fields to retrieve f_m, γ from the same AM chirp sounds with known f_c using MSS and JTFS loss objectives. It is observed that while MSS is able to provide decently reliable gradient updates on γ , its updates on f_m are uninformative. On the contrary, the second order coefficients of JTFS analyze precisely the temporal modulation rate and thus provide informative distance measures that indicate errors in f_m .

Third, in terms of frequency evolution, AM chirp signals have ascending single-tone frequency with their slopes dictated by γ . Drum sounds have $m_1 m_2 = 100$ partials spaced out in exact or progressive deviation from the harmonic series, determined by dispersion D . While neither of these characteristics correspond exactly to a frequency modulation, single-tone frequency evolution is sufficient to be captured by first order time-frequency analysis and the partials spacings indeed have underlying frequency periodicity. That being said, inharmonic partial spacings pose challenge in modulation rate extraction for JTFS.

Finally, although all synthetic sounds in this dataset are aligned in time, real-world sounds are rarely captured with such precision. While JTFS is already a much more adequate objective in representing nonstationary sounds governed by spectrotemporal modulations, its potential can be further exploited with its time-shift invariance property, for example when applied onto more realistic sounds that misalign in time.

	Loss	Φ	AM/FM arpeggiator		FTM drum synthesizer	
			JTFS distance	MSS	JTFS distance	MSS
Han <i>et al.</i> 2020 [35]	P-loss	—	5.2 ± 0.3	0.36 ± 0.01	47 ± 7	0.97 ± 0.07
Engel <i>et al.</i> 2020 [16]	$\mathcal{L}_\theta^{\text{DDSP}}$	Φ_{MSS}	47 ± 3	1.1 ± 0.09	87 ± 6	1.3 ± 0.09
Masuda <i>et al.</i> 2023 [17]	P-loss and $\mathcal{L}_\theta^{\text{DDSP}}$	Φ_{MSS}	14 ± 8	0.51 ± 0.07	90 ± 5	1.3 ± 0.06
Han <i>et al.</i> 2023 [15]	$\mathcal{L}_\theta^{\text{PNP}}$	Φ_{JTFS}	5.5 ± 0.7	0.37 ± 0.04	59 ± 1	1.09 ± 0.01
Ours	P-loss and $\mathcal{L}_\theta^{\text{PNP}}$	Φ_{JTFS}	3.2 ± 2	0.26 ± 0.06	32 ± 1	0.78 ± 0.03

TABLE II: Average JTFS and MSS distance on the test set. Five loss objectives are trained to retrieve parameters from two sound synthesizers: AM/FM arpeggiator (left) and FTM drum synthesizer (right). The top four models are attempted reproductions of published papers referenced at the leftmost column, with slight variations in the training procedure: [16] and [17] have used a weighted combination of MSS and log-MSS L1 distance whereas we use an MSS L^2 distance; [35] and [15] did not apply a “reduce learning rate on plateau” policy.

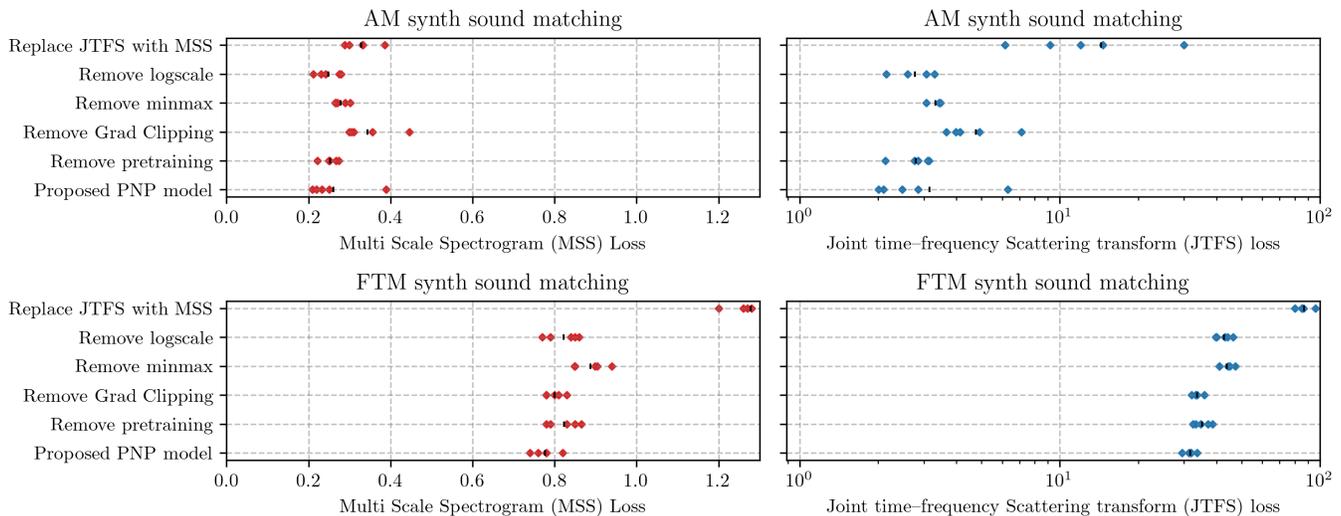


FIG. 2: Ablation study of the proposed model. We demonstrate the deterioration of both JTFS metrics (left column) and MSS metrics (right column) under individual removal of one of various elements in the algorithm. For both metrics, lower values indicate a better perceptual accuracy. Bottom row is the proposed model detailed in Algorithm 1. Each row above shows the extent of degradation when disabling each mechanism. The use of JTFS as opposed to MSS is essential to model performance. Each model is randomly initialized with varying seeds and trained five times (red diamonds), with black vertical lines indicating the average metric across the trials.

2) *Learning dynamics*: To trace the convergence of each loss objective over time, we train a model with P-loss for 35 epochs, then use either $\mathcal{L}_{\theta,\lambda}^{\text{PNP}}$ or $\mathcal{L}_\theta^{\text{DDSP}}$ to finetune the model, or continue training with P-loss for another 35 epochs. The experiments’ optimization setups correspond to respectively [35], [17] and our proposed approach with weight decay and gradient clipping. We observe in Figure 3 that model trained with P-loss soon converges and reaches plateau in both MSS and P-loss metrics. At the start of the fine-tuning stage, we see an abrupt leap in MSS metrics, most likely due to a mismatched learning rate. We expect that the “reduce learning rate on plateau” mechanism would automatically address this discrepancy. Indeed despite the initial leap, $\mathcal{L}_{\theta,\lambda}^{\text{PNP}}$ is able to catch up with P-loss model in parameter error and continue to improve upon MSS metrics. Nevertheless, models trained with $\mathcal{L}_\theta^{\text{DDSP}}$ deteriorate instantly and fail to achieve a good solution.

3) *Effect of reparametrization*: Both min–max and logarithmic transformation are domain-specific knowledge that supposedly aids linearizing parameter errors in the perceptual error domain. We hope, through the inclusion of reparametrization, to provide a continuous optimization landscape where the same error magnitudes at different parameter regions reflect similar perceptual errors. As expected we observe in

Figure 2 that both scaling operations help ease the retrieval. For both AM/FM arpeggiator and FTM drum synthesizer, min–max scaling contributes more to the improvement of model convergence than does logarithmic transformation. Even though removing logarithmic transformation results in inadequate scaling that further deviates from linear correlation and poses serious challenge in training the penultimate layer of \mathbf{f}_W , even if the logarithmic transformation reduces that range discrepancy. This potentially explains the relatively larger degradation in removing min–max scaling when retrieving from FTM drum sounds.

4) *Effect of optimizer*: The addition of weight decay and gradient clipping has greatly improved the convergence across all models, reducing both metrics by up to 2 folds. As shown in Algorithm 1 at each iteration, \mathbf{W} is weight decayed, and then updated by an exponential moving average of the gradient clipped by a scalar ρ . Both mechanisms aim to circumvent the presence of unusually large gradient updates in the current minibatch that causes unstable convergence or oscillations around local minimas.

Algorithm 1 Proposed algorithm for deep neural network training for perceptual sound matching with regularized JTFS-based PNP loss. Includes all additional refinements: P-loss pretraining, adaptive damping, gradient clipping, and reparametrization (both logarithmic and minmax).

Input: Neural network architecture f with weights W , pairs of parameters and sounds $\{(\theta_n, \mathbf{x}_n)\}_{n=1}^N$, minibatch size B , steps per epoch for training set S_{tr} , and validation set S_{val} , hyperparameters $\eta, \beta_1, \lambda, \varepsilon, \rho$.

Set $\eta \leftarrow 0.001, \beta_1 \leftarrow 0.965, m_{T/2} \leftarrow 0, \epsilon \leftarrow 10^{-15}, \rho \leftarrow 0.04, W \sim \text{Normal i.i.d.}$

for sample $n = 1$ to N **do**

$\bar{\theta}_n \leftarrow h(\theta_n)$ ▷ Section IV-D
 $\mathbf{M}(\bar{\theta}_n) \leftarrow \nabla_{(\Phi \circ g \circ h^{-1})}(\bar{\theta}_n) \nabla_{(\Phi \circ g \circ h^{-1})}(\bar{\theta}_n)$

for epoch $t = 1$ to $T/2$ **do**

for step $s = 1$ to S_{tr} **do**

$\mathcal{L}_p \leftarrow \frac{1}{BS_{tr}} \sum_{n \sim \mathbf{U}(1, N)} \|\bar{\theta}_n - \mathbf{f}_W(\mathbf{x}_n)\|_2^2$ ▷ P-loss

$\mathbf{W} \leftarrow \mathbf{W} - \frac{\eta}{BS_{tr}} \sum_{n \sim \mathbf{U}(1, N)} \nabla_{\mathbf{W}} \mathcal{L}_p$

$\mathcal{L}_p^t \leftarrow \frac{1}{BS_{val}} \sum_{n \sim \mathbf{U}(1, N)} \|\bar{\theta}_n - \mathbf{f}_W(\mathbf{x}_n)\|_2^2$

if $\mathcal{L}_p^t \geq \mathcal{L}_p^{t-3}$ **then**

$\eta \leftarrow 0.1\eta$ ▷ Learning rate scheduler

Set $\lambda \leftarrow \lambda_{\max}(\mathbf{M}(\bar{\theta}_n))$ ▷ Section III-C

Set $\eta \leftarrow 0.001$

for epoch $t = T/2 + 1$ to T **do**

for step $s = 1$ to S_{tr} **do**

$\delta_n \leftarrow \mathbf{f}_W(\mathbf{x}_n) - \bar{\theta}_n$
 $\mathcal{L}_{PNP} \leftarrow \frac{1}{BS_{tr}} \sum_{n \sim \mathbf{U}(1, N)} \langle \delta_n | \mathbf{M}(\bar{\theta}_n) + \lambda \mathbf{I} | \delta_n \rangle$

▷ PNP

$m_t \leftarrow \beta_1 m_{t-1} + (1 - \beta_1) \nabla_{\mathbf{W}} \mathcal{L}_{PNP}$ ▷ Momentum

$\mathbf{W} \leftarrow \mathbf{W} - 0.1\eta \mathbf{W}$ ▷ Weight decay

$\mathbf{W} \leftarrow \mathbf{W} - \eta \cdot \text{clip}(m_t / \varepsilon, \rho)$ ▷ Gradient Clipping

$\mathcal{L}_{PNP}^t \leftarrow \frac{1}{BS_{val}} \sum_{n \sim \mathbf{U}(1, N)} \langle \delta_n | \mathbf{M}(\bar{\theta}_n) + \lambda \mathbf{I} | \delta_n \rangle$

if $\mathcal{L}_{PNP}^t \geq \mathcal{L}_{PNP}^{t-3}$ **then**

$\eta \leftarrow 0.1\eta$

if $\mathcal{L}_{PNP}^t \leq \mathcal{L}_{PNP}^{t-1}$ **then**

$\lambda \leftarrow 0.2\lambda$ ▷ Adaptive damping

return \mathbf{f}_W

5) *Blending parameter and perceptual loss:* Despite the slight improvement, pretraining with P-loss is not instrumental in our proposed method as shown in Figure 2. Indeed, as λ is initialized to be a big number, the corrected \mathbf{M} has uniform eigenvalues and thus train the decoder in the P-loss regime. The evolution of subsequent reduction of λ also resembles a gradual transition to optimizing with a perceptual objective. Common practice of blending multiple learning objectives require hyperparameter tuning. In our case this is done semi-automatically by singular value decomposition of \mathbf{M} . As λ_0 is draw from data, namely initialized with the largest eigenvalue of all $\mathbf{M}(\theta)$ in the dataset, the only hyperparameter needed is the damping rate that designate the rate of transitioning.

Experiments demonstrated that, for g generating nonstation-

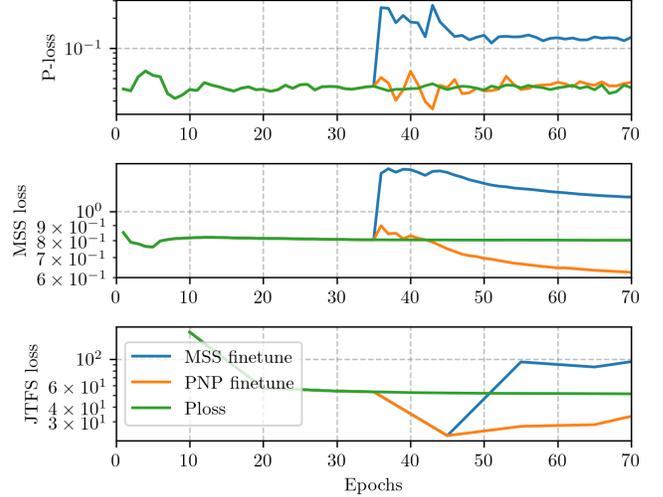


FIG. 3: Learning dynamics of P-loss. Pretraining with P-loss followed by finetuning with JTFS-based PNP loss and MSS loss. The curves show progression of three measures (average P-loss, MSS, and JTFS distance on validation set) over the course of 70 epochs training time. P-loss converges rapidly yet reaches plateau on all metrics. JTFS-based PNP finetuning keeps improving upon the two perceptual metrics without deteriorating on P-loss metrics. On the contrary, MSS finetuning fails to converge to an optimal solution after a spike in metrics due to a re-initialized learning rate.

ary sounds, complex Φ such as JTFS is indispensable to the success of PSM. Reparametrization played a more important role in FTM drum synth parameter retrieval, which has difficult parametrization with more diverse parameter ranges and varied perceptual scaling. Removing pretraining with P-loss has little to no effect on model performance. The degradations among JTFS and MSS metrics are in general consistent. However, due to the presence and absence of second-order time–frequency analysis, the optimal solutions in JTFS metrics are not always the best in MSS metrics. This tendency is especially pronounced in AM synth, possibly due to its construction’s exact alignment with JTFS analysis.

VI. CONCLUSION

We have proposed perceptual neural physical (PNP) loss, a novel learning objective for PSM balancing computational efficiency and perceptual relevance. PNP has achieved state-of-the-art sound matching results on two differentiable synthesizers of nonstationary sounds. Important improvements on the results involve the addition of adaptive damping mechanism, the use of JTFS as perceptual feature, reparametrization, gradient clipping and weight decay.

PNP is the optimal quadratic approximation of L^2 perceptual loss in the synthesis parameter space. It precomputes and stores perceptual effects in the form of kernels in a bilinear formulation, therefore allows working with otherwise cumbersome differentiable implementations of auditory features. Its adaptively damped variation enables semi-automatic transition between parameter loss and perceptual loss regimes.

Applying PNP objective calls for two prerequisites: a synthetic training set for supervised learning, and a differentiable

synthesizer and perceptual representation for automatic differentiation. With the upsurge of direct and adapted implementations of differentiable synthesizers, the latter requirement is increasingly easy to fulfill in practice. In contrast, the former requirement remains a limitation that refrains PNP from solving PSM with real world data. To address this current lack of supervision, future work, such as domain adaptation, is necessary to transfer the knowledge acquired by neural networks from exposure to synthetic data to applications on real world data.

ACKNOWLEDGMENT

This work was granted access to the HPC resources of IDRIS under the allocation 20XX-AD011014253 made by GENCI. V. Lostanlen is supported by ANR project nIrvAna (ANR-23-CE37-0025-04).

REFERENCES

- [1] P.-Y. Oudeyer, *Self-Organization in the Evolution of Speech*, ser. Studies in the Evolution of Language. Oxford University Press, Jan. 2006, vol. 6.
- [2] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter waveform models for statistical parametric speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, 2020.
- [3] Z. Liu, K. Chen, and K. Yu, "Neural Homomorphic Vocoder," in *Proceedings of Interspeech 2020*, 2020.
- [4] L. Renault, R. Mignot, and A. Roebel, "DDSP-Piano: a Neural Sound Synthesizer Informed by Instrument Knowledge," *AES - Journal of the Audio Engineering Society Audio-Acoustics-Application*, vol. 71, no. 9, Sept 2023.
- [5] A. Wiggins and Y. E. Kim, "A differentiable acoustic guitar model for string-specific polyphonic synthesis," in *Proceedings of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, WASPAA*.
- [6] A. Amador and G. B. Mindlin, "Synthetic birdsongs as a tool to induce, and listen to, replay activity in sleeping birds," *Frontiers in Neuroscience*, vol. 15, 2021.
- [7] F. Caspe, A. McPherson, and M. Sandler, "DDX7: Differentiable fm synthesis of musical instrument sounds," in *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, 2022.
- [8] O. Barkan, D. Tsiris, O. Katz, and N. Koenigstein, "Inversynth: Deep estimation of synthesizer parameter configurations from audio signals," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 12, 2019.
- [9] C. Traube and J. Smith, "Extracting the fingering and the plucking points on a guitar string from a recording," in *Proceedings of the IEEE Workshop on the Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2001.
- [10] J. Riionheimo and V. Välimäki, "Parameter estimation of a plucked string synthesis model using a genetic algorithm with perceptual fitness calculation," *EURASIP Journal on Advances in Signal Processing*, 2003.
- [11] C. E. Cella, L. Dzwonczyk, A. Saldarriaga-Fuertes, H. Liu, and H.-C. Crayencour, "A Study on Neural Models for Target-Based Computer-Assisted Musical Orchestration," in *2020 Joint Conference on AI Music Creativity (CSMC + MuMe)*, ser. Proceedings of the Joint Conference on AI Music Creativity (AIMC), Oct. 2020.
- [12] J. Turian, J. Shier, G. Tzanetakis, K. McNally, and M. Henry, "One billion audio sounds from gpu-enabled modular synthesis," in *Proceedings of the International Conference on Digital Audio Effects (DAFx)*, 2021.
- [13] M. Yee-King and M. Roth, "A comparison of parametric optimisation techniques for musical instrument tone matching," *Audio Engineering Society Convention 130*, vol. 2, Jan 2011.
- [14] C. Vahidi, H. Han, C. Wang, M. Lagrange, G. Fazekas, and V. Lostanlen, "Mesostructures: Beyond Spectrogram Loss in Differentiable Time-Frequency Analysis," *Journal of the Audio Engineering Society*, 2023.
- [15] H. Han, V. Lostanlen, and M. Lagrange, "Perceptual-Neural-Physical Sound Matching," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Jun. 2023.
- [16] J. Engel, L. H. Hantrakul, C. Gu, and A. Roberts, "DDSP: Differentiable Digital Signal Processing," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.
- [17] N. Masuda and D. Saito, "Improving semi-supervised differentiable synthesizer sound matching for practical applications," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, 2023.
- [18] Da-Yi Wu, Wen-Yi Hsiao, Fu-Rong Yang, Oscar Friedman, Warren Jackson, Scott Bruzenak, Yi-Wen Liu and Yi-Hsuan Yang, "DDSP-based Singing Vocoders: A New Subtractive-based Synthesizer and A Comprehensive Evaluation," 2022.
- [19] A. Carson, S. King, C. Valentini Botincho, and S. Bilbao, "Differentiable grey-box modelling of phaser effects using frame-based spectral processing," in *Proceedings of the International Conference on Digital Audio Effects (DAFx)*, ser. Proceedings of the International Conference on Digital Audio Effects (DAFx), Sept 2023.
- [20] S. Schlecht, J. Parker, M. Schäfer, and R. Rabenstein, "Physical modeling using recurrent neural networks with fast convolutional layers," in *Proceedings of the International Conference on Digital Audio Effects (DAFx)*, Sept 2022.
- [21] R. Diaz, B. Hayes, C. Saitis, G. Fazekas, and M. Sandler, "Rigid-body sound synthesis with differentiable modal resonators," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun 2023.
- [22] B. Hayes, J. Shier, G. Fazekas, A. McPherson, and C. Saitis, "A review of differentiable digital signal processing for music & speech synthesis," 2023.
- [23] P. Albouy, L. Benjamin, B. Morillon, and R. Zatorre, "Distinct sensitivity to spectrotemporal modulation supports brain asymmetry for speech and melody," *Science*, vol. 367, Feb 2020.
- [24] D. A. Depireux, J. Z. Simon, D. J. Klein, and S. A. Shamma, "Spectrotemporal response field characterization with dynamic ripples in ferret primary auditory cortex," *Journal of Neurophysiology*, vol. 85, no. 3, 2001.
- [25] T. Chi, P. Ru, and S. Shamma, "Multiresolution spectrotemporal analysis of complex sounds," *The Journal of the Acoustical Society of America*, vol. 118, Sept 2005.
- [26] E. Thoret, B. Caramiaux, P. Depalle, and S. Mcadams, "Learning metrics on spectrotemporal modulations reveals the perception of musical instrument timbre," *Nature Human Behaviour*, vol. 5, Mar 2021.
- [27] J. Andén, V. Lostanlen, and S. Mallat, "Joint time-frequency scattering," *IEEE Transactions on Signal Processing*, vol. PP, May 2019.
- [28] V. Lostanlen, C. El-Hajj, M. Rossignol, G. Lafay, J. Andén, and M. Lagrange, "Time-frequency scattering accurately models auditory similarities between instrumental playing techniques," *EURASIP Journal on Audio, Speech, and Music Processing*, no. 1, 2021.
- [29] J. Muradeli, C. Vahidi, C. Wang, H. Han, V. Lostanlen, M. Lagrange, and G. Fazekas, "Differentiable time-frequency scattering in kymatio," in *Proceedings of the International Conference on Digital Audio Effects (DAFx)*, 2022.
- [30] M. J. Yee-King, L. Fedden, and M. d'Inverno, "Automatic programming of vst sound synthesizers using deep networks and other techniques," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 2, 2018.
- [31] J. Turian and M. Henry, "I'm sorry for your loss: Spectrally-based audio distances are bad at pitch," *Proceedings of the "I Can't Believe It's Not Better" Workshop (ICBINB)*, 2020.
- [32] C. Chadebec, E. Thibeau-Sutre, N. Burgos, and S. Allasonnière, "Data augmentation in high dimensional low sample size setting using a geometry-based variational autoencoder," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, 2023.
- [33] L. Trautmann and R. Rabenstein, *Digital Sound Synthesis by Physical Modeling Using the Functional Transformation Method*. Springer, 2003.
- [34] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proceedings of the International conference on Machine Learning (ICML)*, 2019.
- [35] H. Han and V. Lostanlen, "wav2shape: Hearing the Shape of a Drum Machine," in *Proceedings of Forum Acusticum*, 2020.
- [36] B. Hayes, C. Saitis, and G. Fazekas, "Sinusoidal frequency estimation by gradient descent," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.