



HAL
open science

Unsupervised Harmonic Parameter Estimation Using Differentiable DSP and Spectral Optimal Transport

Bernardo Torres, Geoffroy Peeters, Gaël Richard

► **To cite this version:**

Bernardo Torres, Geoffroy Peeters, Gaël Richard. Unsupervised Harmonic Parameter Estimation Using Differentiable DSP and Spectral Optimal Transport. IEEE International Conference on Acoustics, Speech and Signal Processing, Apr 2024, Seoul, South Korea. hal-04358467v1

HAL Id: hal-04358467

<https://hal.science/hal-04358467v1>

Submitted on 21 Dec 2023 (v1), last revised 11 Jan 2024 (v3)

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.

Unsupervised Harmonic Parameter Estimation Using Differentiable DSP and Spectral Optimal Transport

Bernardo Torres Geoffroy Peeters Gaël Richard

LTCI, Telecom Paris, Institut Polytechnique de Paris

Abstract

In neural audio signal processing, pitch conditioning has been used to enhance the performance of synthesizers. However, jointly training pitch estimators and synthesizers is a challenge when using standard audio-to-audio reconstruction loss, leading to reliance on external pitch trackers. To address this issue, we propose using a spectral loss function inspired by optimal transportation theory that minimizes the displacement of spectral energy. We validate this approach through an unsupervised autoencoding task that fits a harmonic template to harmonic signals. We jointly estimate the fundamental frequency and amplitudes of harmonics using a lightweight encoder and reconstruct the signals using a differentiable harmonic synthesizer. The proposed approach offers a promising direction for improving unsupervised parameter estimation in neural audio applications.

Index Terms: differentiable signal processing, machine learning, optimal transport, frequency estimation

1 Introduction

Spectral Modeling Synthesis (SMS) [1], a foundation in signal processing, aims to represent audio signals through a combination of sinusoidal and noise components that define clear spectral structures visible in short-term spectral analysis. A significant portion of work on SMS has been dedicated to modeling harmonically related sinusoidal signals, where a fundamental frequency parameterizes a harmonic template [2]. The task is threefold: identifying the fundamental frequency, estimating harmonic amplitudes, and capturing transients and aperiodicities. While existing methods offer suitable solutions, they often fall short in dynamic scenarios where frequency amplitudes vary over time or in data-driven systems trained end-to-end via gradient descent.

Recent advances in Differentiable Digital Signal Processing (DDSP) [3] have sought to bridge traditional signal processing with the capabilities of deep learning. These methods often depend on an external pitch estimator for determining the fundamental frequency, subsequently used for conditioning synthesizers [4]. This approach favors harmonic alignment but constrains their flexibility and adaptability. On the other hand, Deep Neural Networks (DNNs) have outperformed traditional DSP-based methods in pitch estimation but still, for the most part, rely on supervised learning. Additionally, they might suffer from poor generalization due to data distribution mismatches [5], and poor integration with other tasks.

This work was funded by the European Union (ERC, HI-Audio, 101052978). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Some noteworthy approaches, such as SPICE [6] have emerged that tackle pitch estimation through analysis-by-synthesis and self-supervised learning. However, most existing methods focus either on pitch or amplitude estimation, rarely both. The seminal DDSP paper [3] and subsequent works [7] have shown that unsupervised approaches to joint estimation is particularly challenging, and notable recent works address this issue directly through pre-training on synthetic data [8], or using a damped sinusoidal model [9].

Recent literature has begun to question the efficacy of the commonly used audio-to-audio reconstruction losses [10]. Our work grows out of this branch and turns towards the theory of optimal mass transport for a novel perspective for comparing spectral representations.

Optimal transport has been employed in various audio tasks, such as Non-negative Matrix Factorization (NMF) based transcription [11], Blind-Source separation [12], estimating the pitch of in-harmonic signals [13], performing spectral interpolation [14] and defining distances between power spectral densities (PSDs) of time series [15]. To the best of our knowledge, this is the first work to propose an optimal transport-based cost for differentiable spectral comparison and empirical risk minimization in a neural audio system.

This paper is organized as follows. Firstly, Section 2 delves into the limitations of DDSP models and *vertical* spectral losses. Section 3 introduces our alternative loss function, followed by a description of our proposed autoencoding task and experiments (Section 4) and results (Section 5). Section 6 provides some concluding remarks.

2 DDSP and the Multi-Scale Spectral loss

DDSP synthesizers allow fine-grained spectral modeling [3]. When provided with accurate fundamental frequency (f_0), these models have a high degree of harmonic alignment by design. DDSP-inspired works have capitalized on this property by using f_0 conditioning from external pitch trackers [4] and optimizing spectral losses.

In time-frequency analysis, larger window sizes offer finer frequency discrimination but at the cost of temporal resolution. Conversely, smaller window sizes are prone to spectral leakage, although they have better time resolution. By using multiple resolutions of the Short-Time Fourier Transform (STFT), the Multi-Scale Spectral loss (MSS) hopes to achieve the best of both worlds [3, 16]:

$$\mathcal{L}_{\text{MSS}}(s, \hat{s}) = \sum_{\gamma \in \Gamma} \mathcal{L}_{\text{lin}}^{\gamma} + \mathcal{L}_{\text{log}}^{\gamma}, \quad (1)$$

where (s, \hat{s}) are time domain signals. The linear and log spectral losses are $\mathcal{L}_{\text{lin}}^{\gamma} = \|\mathbf{S}_{\gamma} - \hat{\mathbf{S}}_{\gamma}\|_1$ and $\mathcal{L}_{\text{log}}^{\gamma} = \|\log(\mathbf{S}_{\gamma}) - \log(\hat{\mathbf{S}}_{\gamma})\|_1$, respectively, where $\mathbf{S}_{\gamma} = |\text{STFT}_{\gamma}(s)|$ is the magnitude STFT at scale (window size) γ . The set Γ usually consists of powers of 2. The linear loss is more effective during the initial training stages [17]

by emphasizing large amplitude differences, while the logarithmic loss may help in discerning more subtle frequency patterns.

Following the terminology in [15], we refer to losses that operate point-wise between spectra with identical support as *vertical* losses. Conventional spectral losses fall under this category, comparing using an L_p norm the amplitude of a given frequency bin of a reference frame spectrum to the amplitude of the same bin on a target. If the spectra do not overlap, a *vertical* comparison may not properly measure the amount of energy displacement on the frequency axis.

When attempting to learn a pitch tracker with the neural audio model in an end-to-end fashion, the MSS (a *vertical* loss) has been reported as ineffective [3, 7], which may be due to its poor sense of gradient orientation with respect to an oscillator’s frequency [10, 18] and many notable local minima [9].

A *horizontal* comparison of spectra should faithfully represent spectral energy displacement. In this context, we aim to validate an alternative loss function that enables a more nuanced, *horizontal* comparison of spectral content, thereby guiding neural audio models toward more accurate frequency estimation. The framework of Optimal Transport offers precisely that.

3 Optimal transport

Let $\alpha = \sum_{i=1}^n \mathbf{a}_i \delta_{x_i}$ and $\beta = \sum_{j=1}^m \mathbf{b}_j \delta_{y_j}$ be discrete measures with weights \mathbf{a}_i at positions x_i and weights \mathbf{b}_j at positions y_j , where δ_x is the Dirac measure at location x . Let additionally (\mathbf{a}, \mathbf{b}) belong to the space of probability vectors, *i.e.* $\mathbf{a} \in \Sigma_n$ and $\mathbf{b} \in \Sigma_m$, for $\Sigma_n = \{\mathbf{a} \in \mathbb{R}_+^n; \sum_{i=1}^n \mathbf{a}_i = 1\}$. Given a cost function $c(x_i, y_j)$ representing the cost of transporting a unit of mass from x_i to y_j , the discrete Optimal Transport (OT) problem is [19]:

$$\mathcal{L}_c(\alpha, \beta) = \min_{\mathbf{P}} \sum_{i=1}^n \sum_{j=1}^m \mathbf{P}_{ij} c(x_i, y_j), \quad (2)$$

where $\mathbf{P} \in \mathbb{R}_+^{n \times m}$ is the transport plan describing the amount of mass to be transported from bin i to bin j , subject to mass conservation constraints $\{\mathbf{a}_i = \sum_{j=1}^m \mathbf{P}_{ij}\}_{i=1, \dots, n}$ and $\{\mathbf{b}_j = \sum_{i=1}^n \mathbf{P}_{ij}\}_{j=1, \dots, m}$. The optimal transport map, \mathbf{P}^* , gives the optimal way to transport α to β . Note that $\mathcal{L}_c(\alpha, \beta)$ depends on both the weights (\mathbf{a}, \mathbf{b}) , and the supporting positions (x_i, y_j) [19].

For cost functions $c(x_i, y_j) = |x_i - y_j|^p$, $\mathcal{L}_c(\alpha, \beta)^{1/p}$ is the p -Wasserstein distance between α and β [20]. This distance allows for the comparison of distributions when they do not overlap, which is a limitation of L_p norm-based distances and commonly employed divergences (such as the Kullback–Leibler).

3.1 Optimal transport in one dimension

While being challenging for dimensions $d > 1$, the OT problem in 1D for *general* measures $(\alpha, \beta) \in \mathcal{M}(\mathcal{X})$ (the set of Radon measures on the space \mathcal{X}) has a closed-form solution for the p -Wasserstein distance [19, 21]:

$$\mathcal{W}_p(\alpha, \beta)^p = \int_0^1 |F_\alpha^{-1}(r) - F_\beta^{-1}(r)|^p dr, \quad (3)$$

where $F_\alpha : \mathbb{R} \rightarrow [0, 1]$ is the cumulative distribution function (CDF) $F_\alpha(x) = \int_{-\infty}^x d\alpha$ and $F_\alpha^{-1} : [0, 1] \rightarrow \mathbb{R}$ is its pseudoinverse:

$$F_\alpha^{-1}(r) = \inf \{x \in \mathbb{R} : F_\alpha(x) \geq r\}, \quad (4)$$

the generalized quantile function. When dealing with discrete measures, Equation 3 can be approximated by a discrete sum [22]:




Fig. 1: Cumulative sum and quantile functions of two measures α and β representing harmonic spectra with different f_0 s across the Nyquist range [0-0.5]. The grey lines depict the pointwise differences between the inverse CDFs as per Eq. 5, with the shaded region representing the Wasserstein metric \mathcal{W}_1 . Green connectors denote the optimal plan (\mathbf{P}_{ij}^*) , weighted by their respective magnitudes.

$$\mathcal{W}_p(\alpha, \beta)^p = \sum_{i=1}^n |F_\alpha^{-1}(r_i) - F_\beta^{-1}(r_i)|^p (r_i - r_{i-1}), \quad (5)$$

where F is computed using the step function $u(\cdot)$: $F_\alpha(x) = \sum_{i=1}^n u(x - x_i) \mathbf{a}_i$. The set r can be taken to be the ordered set of quantiles of both α and β [22]. Figure 1 illustrates this process.

3.2 Spectral optimal transport

The Wasserstein-Fourier distance was defined by Cazelles et al. [15] as the \mathcal{W}_2 distance between Power Spectral Densities of stationary time-series signals. We refer hereafter to Spectral Optimal Transport (SOT) as a more general sense of *horizontal* (measuring energy displacement) comparison of discrete time-frequency spectra by approximating \mathcal{W}_p using Equation 5.

Let $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}_+^{L \times M}$ be a normalized time-frequency transform so that each frame $\Phi(s)_j$ sums to 1. We define the *horizontal* SOT loss \mathcal{L}_{SOT} as the temporal mean (over L time frames) of the p -Wasserstein distance computed on the frequency axis of Φ :

$$\mathcal{L}_{\text{SOT}}(s, \hat{s}) = \frac{1}{L} \sum_{j=0}^{L-1} \mathcal{W}_p(\Phi(s)_j, \Phi(\hat{s})_j)^p, \quad (6)$$

where s and \hat{s} are time-domain signals. Recall that \mathcal{W}_p depends on the weights (\mathbf{a}, \mathbf{b}) and their positions $(\{x_i\}_{i=1, \dots, M}, \{y_j\}_{j=1, \dots, M})$, which here represent respectively the amplitude vectors of the frame spectra and the frequencies of bins (i, j) of Φ .

SOT for simple sinusoids: As demonstrated in [15], the Wasserstein-Fourier distance of two frequency-shifted signals is proportional to the frequency shift. We illustrate this in Figure 2, where we show the value of different spectral losses between two sinusoidal signals as a function of their frequency difference. We compute the losses in the time-frequency domain and average over time, using the SOT formulation in Eq. 6 with \mathcal{W}_2 . SOT points towards a minimal frequency difference, while the *vertical* L_1 and L_2 losses converge smoothly only for small differences. This serves as a motivating example for studying the SOT loss for more challenging spectra.




Fig. 2: Normalized *spectral* Single-Scale (SS), Multi-Scale (MSS) and proposed SOT \mathcal{W}_2 losses as a function of the frequency shift between two 16 KHz sinusoidal signals of 4096 samples. The reference sinusoid has a frequency of 4000 Hz. SS is computed for window size $\gamma = 1024$ and MSS with scales $\Gamma = \{2^k\}_{k=6, \dots, 11}$.

4 Methodology and experiments

Our approach employs an encoder for joint frame-wise estimation of fundamental frequency and harmonic amplitudes of a signal, taking as input the Constant-Q Transform (CQT). The signal is reconstructed in the time domain using a DDSF harmonic synthesizer [3], and trained end-to-end via an audio reconstruction loss. The underlying assumption of this approach is that the instantaneous frequency varies slowly enough for accurate estimation from spectral frames (local stationarity). Through this toy autoencoding task, we study the interplay between *vertical* losses (such as the MSS) and *horizontal* SOT variants under various conditions.

4.1 Analysis/synthesis autoencoder

Input representation: We compute the CQT from the input waveform on the fly [23]. The CQT’s logarithmic scaling along frequencies naturally relates pitch shifting in semitones as frequency-axis translations. The lower frequency bound of the CQT is set to 32.7 Hz (the C1 note) and 285 bins are used to cover the whole 8 kHz spectrum (amounting to 3 bins per semitone). A hop size of 256 is used, which defines the estimation rate of the system.

Encoder: Our encoder adopts the lightweight architecture from [24] (46K parameters), featuring 1D convolutions along the frequency axis. A Toeplitz linear layer¹ maps the output feature map from the encoder to 285 pitch logits, which are mapped to a categorical distribution over pitch classes $\mathbf{p} \in \mathbb{R}^{285}$ through a softmax operation [8]. The employed architecture ensures translation invariance, meaning that a pitch shift in the input (CQT) corresponds to a shift in the output pitch probabilities by an equal amount of bins.

We set the number of estimated harmonics H to 20. Harmonic amplitudes are obtained through a feed-forward linear layer, which maps the the output feature map from the encoder to H amplitude values \mathbf{c} , followed by an exponentiated sigmoid activation from DDSF [3]. Unlike [3], we treat each spectral frame as independent.

Pitch regression: The pitch is computed as the expectation of the pitch distribution, or the *soft-argmax*: $\hat{f}^u = \sum_{i=1}^{285} f_i^u \mathbf{p}_i$; where f_i^u is the frequency of the i -th pitch class, logarithmically mapped to the $[0, 1]$ range. Each bin f_i^u corresponds to a frequency bin in the input CQT. A temperature parameter $\tau = 0.1$ is used to pre-scale the logits to facilitate a unimodal distribution [25].

Harmonic synthesis: We use a DDSF harmonic synthesizer [3]:

$$s(n) = \sum_{h=0}^{H-1} \mathbf{c}_h(n) \sin(\phi_h(n)), \quad (7)$$

¹We point the reader to [24] for a detailed explanation of the architecture

parameterized at the sample-level by f_0 and amplitudes \mathbf{c} . The cumulative sum $\phi_h(n) = \sum_{t=0}^n f_h(t)$ computes the phase at time n , for $f_h = hf_0$. We set initial phases to 0 and use the same frame rate to sample rate interpolation of f_0 and \mathbf{c} as in DDSF.

4.2 Data

We generate a synthetic dataset of 4000 examples, where each example is a 16 kHz signal with $N = 4096$ samples. Harmonic synthesis (Eq. 7) is performed with a fixed instantaneous frequency during the signal length. The number of harmonics of each signal is selected randomly from [1-8]. The f_0 values are randomly drawn from the range [40-1950] Hz and the harmonic amplitudes from [0.4-1]. The dataset is split into 70% training, 20% validation, and 10% testing.

4.3 Baselines

The baselines are trained with the MSS loss only, for scales $\Gamma = \{2048, 1024, 512, 256, 128, 64\}$ and hop size respecting 75% overlap. MSS-Lin uses only \mathcal{L}_{lin} and MSS-LogLin uses $\mathcal{L}_{\text{log}} + \mathcal{L}_{\text{lin}}$.

4.4 Spectral Optimal Transport experiments

We use \mathcal{W}_2 in \mathcal{L}_{SOT} for computing the Wasserstein distances between frame spectra, employing the squared magnitude STFT with a *flattop*² window of size γ and hop size of 256 as the time-frequency transform Φ . In all SOT experiments, MSS was also included as a *vertical* loss (only \mathcal{L}_{lin}) to penalize \mathcal{W}_2 minima with low spectral overlap. This can happen, for instance, when the number of active harmonics is overestimated. This also addresses instability issues observed with the 1D Wasserstein loss alone [26], possibly due to the normalization of the spectra. The final objective is $\mathcal{L}_{\text{SOT}} + \lambda \mathcal{L}_{\text{lin}}$, with $\lambda = 0.05$. We additionally employ a frame-wise frequency cutoff in the spectrum of the estimated signal, limiting energy to lower frequencies, which helped to significantly reduce octave errors.

Frequency cutoff: The process starts with proportional normalization of spectra: the target frame spectrum $\Phi(s)_j$ sums to 1, while the reconstructed frame $\Phi(\hat{s})_j$ is normalized to a sum potentially exceeding 1 if its energy is greater than the original. For frame j , f_{cut} is determined to be the frequency where the cumulative sum of the reconstructed spectrum matches the total normalized energy of the reference³ ($F_{\Phi(\hat{s})_j}(f_{\text{cut}}) = 1$).

SOT Variants: We experiment with three main \mathcal{L}_{SOT} variants, varying window size and frequency scaling: (I) SOT-512 uses $\gamma = 512$; (II) SOT-512-LS is similar to SOT-512 but additionally scales frequencies logarithmically before computing \mathcal{W}_2 (i.e. $x_i \rightarrow \log(x_i), y_j \rightarrow \log(y_j)$), thereby reducing the frequency ground cost for higher frequencies; (III) SOT-2048 uses $\gamma = 2048$.

Ablations: We provide two ablations to SOT-2048. First, we examined the special case SOT-SingleScale using a single STFT resolution $\Gamma = \{512\}$ for \mathcal{L}_{lin} , and $\lambda = 0.1$. Lastly, the frequency cutoff is removed in configuration SOT-NoCut.

4.5 Training details

We employed the Adam optimizer with a learning rate of 1×10^{-4} and a batch size of 64 for $25k$ steps. To ensure robustness against

²Using a window with less prominent side lobes helped with stability.

³In practical terms, rather than applying a filtering operation, we ignore indices (i) for quantile (r) values over one in the computation of Equation 5.

	Mean (STD)				Median			
	LSD [db] ↓	RPA [%] ↑	RCA [%] ↑	OD	LSD	RPA	RCA	OD
MSS-Lin	46.4 (21.4)	20.2 (44.6)	26.9 (42.7)	-2.3 (1.3)	58	0.2	3.9	-2.8
MSS-LogLin	80.5 (15.1)	1.4 (2.7)	4.0 (4.5)	-0.9 (1.9)	82.6	0.1	3.2	-0.6
SOT-NoCut	70.6 (31.8)	23.7 (30.3)	46.0 (36.4)	-0.4 (0.4)	77.6	20	45	-0.2
SOT-SingleScale	97.9 (32.5)	14.1 (25.5)	28.6 (32.6)	0.0 (1.0)	101.1	4.7	11.6	0.3
SOT-512	40.5 (23.5)	42.9 (39.4)	62.3 (42.6)	-0.8 (0.7)	26.6	63.6	75.2	-0.4
SOT-512-LS	25.9 (2.5)	55.4 (36.1)	86.8 (16.2)	-0.7 (0.9)	<u>25</u>	<u>63.7</u>	<u>95.6</u>	-0.4
SOT-2048	23.5 (3.5)	75.0 (43.2)	99.2 (1.6)	-0.5 (0.9)	24.5	99.7	99.8	0

Table 1: Experimental results for the joint f_0 and harmonic amplitude estimation task on synthetic data. We report the mean, standard deviation and median evaluation metrics for 5 separate training runs. Best results are indicated in bold and second-best are underlined.

weight initialization, each model was trained five times using different seeds for the pseudo-random number generator. Code and training recipes are available online⁴.

4.6 Evaluation

Evaluation is done using a combination of pitch estimation and signal reconstruction metrics. We employ Raw Pitch Accuracy (RPA), Raw Chroma Accuracy (RCA) [27, 28], and Log Spectral Distance (LSD). RPA measures the percentage of frames with estimated pitch close to the ground truth within 0.5 semitones. RCA is similar but allows octave errors. The LSD is computed as:

$$\text{LSD}(s, \hat{s}) = \frac{1}{LM} \left\| \log(\mathbf{S}_\gamma) - \log(\hat{\mathbf{S}}_\gamma) \right\|_F^2, \quad (8)$$

where $\|\cdot\|_F$ denotes the Frobenius norm of a matrix, $\mathbf{S}_\gamma = |\text{STFT}_\gamma(s)| \in \mathbb{R}^{L \times M}$ denotes the magnitude STFT of s for window size γ , with L frames and M frequency bins. We use $\gamma = 1024$ and a hop size of 256. The Mean Octave Difference (OD) is also included to indicate, on average, the mean deviation (in octaves) of the estimated pitch from the ground truth, highlighting the nature of the octave errors. For each run, the checkpoint with the lowest validation reconstruction score was selected for testing. We report the mean metrics on the test set.

5 Results

Table 1 summarizes the experimental results for our trained models on the synthetic dataset. The high standard deviation (STD) reveals a significant dependence on initialization. The initial pitch probability distribution, determined by the neural networks’ weight initialization, was a key factor impacting the model’s success.

This sensitivity to initialization is particularly pronounced in models trained only with the Multi-Scale Spectral loss (MSS-Lin and MSS-LogLin). Out of five runs, only one achieved optimal LSD, RPA, and RCA. This variability is evidenced by the large standard deviation and low median values. Interestingly, logarithmic compression (MSS-LogLin) resulted in worse performance.

In contrast, models using the Spectral Optimal Transport (SOT) loss demonstrated greater robustness, converging to satisfactory solutions on average. Notably, the model employing a larger window size (SOT-2048) outperformed all others, with a mean of 23.5 dB LSD, 75% RPA and 99.2% RCA.

Frequency scaling and frequency cutoffs helped performance. Logarithmic frequency scaling (SOT-512-LS) led to an improvement over SOT-512. Logarithmic scaling makes the SOT loss less punitive for high-frequency deviations, aligning with our focus on

estimating the fundamental frequency. The use of the frequency cutoff was found to be beneficial when the model overestimates the signal’s energy or is initialized with many active harmonics. This prevents the gradient from skewing towards an incorrect f_0 after the normalization step needed for SOT loss.

The SOT-SingleScale model exhibited suboptimal results in LSD, RPC, and RCA, but maintained an average octave difference of zero, indicating that the pitch solution it converged to was within the octave range of the ground truth. This is a nuanced failure mode in which the model retains more or less harmonics than actually present in the reference signal. Obtaining perfect spectral alignment by minimizing \mathcal{W}_2 in \mathcal{L}_{SOT} requires matching the number of harmonics with the reference, implying that a mismatch may lead to inaccurate f_0 predictions. The introduction of a *vertical* loss was intended to penalize solutions with low spectral alignment, but using only a single scale has proven to be insufficient.

While promising, our results also raise questions on the interplay between MSS and SOT losses. Preliminary findings suggest that these loss functions may be optimizing different aspects of the estimator and could be working in a conflicting manner. For instance, the *horizontal* loss might favor frequency shifts in the spectra, while the *vertical* loss may lower the amplitudes, possibly causing it to get stuck in a local minimum. Moreover, the stability of these losses in more complex scenarios, such as with many harmonics, transients, or inharmonicity remains an open question. The SOT loss was found to be highly sensitive to small frequency variations, such as in spectral side lobes or noise, and in our preliminary experiments, it struggled with more realistic datasets such as Nsynth [29].

6 Conclusion

In this work, we leverage optimal transport theory to develop a training objective for audio reconstruction tasks. The proposed Spectral Optimal Transport (SOT) loss, which compares discrete spectra *horizontally* (measuring energy displacement), overcomes some limitations of traditional spectral losses. We have shown that it has the potential for frequency localization and joint estimation of fundamental frequency (f_0) and harmonic amplitudes. Results on a synthetic dataset have demonstrated the efficacy of our method compared to the commonly employed Multi-Scale Spectral (MSS) loss.

However, our findings also reveal that the proposed loss is very sensitive to small spectral variations and might not suffice, in its current state, for accurate f_0 matching in real signals. This underscores the need for a deeper understanding of the trade-offs involved.

⁴<https://github.com/torresbf/1d-spectral-optimal-transport>

7 References

- [1] Xavier Serra and Julius Smith, “Spectral modeling synthesis: A sound analysis/synthesis system based on a deterministic plus stochastic decomposition,” *Computer Music Journal*, vol. 14, no. 4, pp. 12–24, 1990.
- [2] R McAuley and T Quatieri, “Speech analysis/synthesis based on a sinusoidal representation,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 34, pp. 744–754, 1986.
- [3] Jesse H. Engel, Lamtharn Hantrakul, Chenjie Gu, and Adam Roberts, “DDSP: Differentiable digital signal processing,” in *ICLR*. 2020, OpenReview.net.
- [4] Ben Hayes, Jordie Shier, György Fazekas, Andrew McPherson, and Charalampos Saitis, “A review of differentiable digital signal processing for music & speech synthesis,” *arXiv preprint arXiv:2308.15422*, 2023.
- [5] Max Morrison, Caedon Hsieh, Nathan Pruyne, and Bryan Pardo, “Cross-domain Neural Pitch and Periodicity Estimation,” 2023.
- [6] Beat Gfeller, Christian Frank, Dominik Roblek, Matt Sharifi, Marco Tagliasacchi, and Mihajlo Velimirović, “SPICE: Self-supervised pitch estimation,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 1118–1128, 2020.
- [7] Da-Yi Wu, Wen-Yi Hsiao, Oscar Friedman, Warren Jackson, Scott Bruzenak, Yi-Wen Liu, Yi-Hsuan Yang, and Fu-Rong Yang, “DDSP-based singing vocoders: A new subtractive-based synthesizer and A comprehensive evaluation,” *ISMIR*, 2022.
- [8] Jesse Engel, Rigel Swavely, Lamtharn Hanoi Hantrakul, Adam Roberts, and Curtis Hawthorne, “Self-supervised pitch detection by inverse audio synthesis,” in *ICML 2020 Workshop on Self-Supervision in Audio and Speech*, 2020.
- [9] Ben Hayes, Charalampos Saitis, and György Fazekas, “Sinusoidal frequency estimation by gradient descent,” in *ICASSP*, 2023, pp. 1–5.
- [10] Joseph Turian and Max Henry, “I’m Sorry for Your Loss: Spectrally-Based Audio Distances Are Bad at Pitch,” 2020.
- [11] Rémi Flamary, Cédric Févotte, Nicolas Courty, and Valentin Emiya, “Optimal spectral transportation with application to music transcription,” *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [12] Antoine Rolet, Vivien Seguy, Mathieu Blondel, and Hiroshi Sawada, “Blind source separation with optimal transport non-negative matrix factorization,” *EURASIP Journal on Advances in Signal Processing*, vol. 2018, no. 1, pp. 1–16, 2018.
- [13] Filip Elvander, Stefan Ingi Adalbjornsson, Johan Karlsson, and Andreas Jakobsson, “Using optimal transport for estimating inharmonic pitch signals,” in *ICASSP*. 2017, pp. 331–335, IEEE.
- [14] Trevor Henderson and Justin Solomon, “Audio transport: A generalized portamento via optimal transport,” *arXiv preprint arXiv:1906.06763*, 2019.
- [15] Elsa Cazelles, Arnaud Robert, and Felipe Tobar, “The Wasserstein-Fourier distance for stationary time series,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 709–721, 2020.
- [16] Xin Wang, Shinji Takaki, and Junichi Yamagishi, “Neural source-filter waveform models for statistical parametric speech synthesis,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 402–415, 2019, Publisher: IEEE.
- [17] Sercan Ö Arık, Heewoo Jun, and Gregory Diamos, “Fast spectrogram inversion using multi-head convolutional neural networks,” *IEEE Signal Processing Letters*, vol. 26, no. 1, pp. 94–98, 2018.
- [18] Naotake Masuda and Daisuke Saito, “Improving Semi-Supervised Differentiable Synthesizer Sound Matching for Practical Applications,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 31, pp. 863–875, 2023.
- [19] Gabriel Peyré and Marco Cuturi, “Computational optimal transport,” *Foundations and Trends in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.
- [20] Soheil Kolouri, Se Rim Park, Matthew Thorpe, Dejan Slepcevic, and Gustavo K Rohde, “Optimal mass transport: Signal processing and machine-learning applications,” *IEEE signal processing magazine*, vol. 34, no. 4, pp. 43–59, 2017.
- [21] Filippo Santambrogio, “Optimal transport for applied mathematicians,” *Birkhäuser, NY*, vol. 55, no. 58-63, pp. 94, 2015.
- [22] R. Flamary, N. Courty, A. Gramfort, Mokhtar Z. Alaya, A. Boisbunon, S. Chambon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, Léo G., N. T.H. Gayraud, H. Janati, A. Rakotomamonjy, I. Redko, A. Rolet, A. Schutz, V. Seguy, D. J. Sutherland, R. Tavenard, A. Tong, and T. Vayer, “POT: Python optimal transport,” *Journal of Machine Learning Research*, vol. 22, no. 78, pp. 1–8, 2021.
- [23] Kin Wai Cheuk, Hans Anderson, Kat Agres, and Dorien Herremans, “nnAudio: An on-the-Fly GPU Audio to Spectrogram Conversion Toolbox Using 1D Convolutional Neural Networks,” *IEEE Access*, vol. 8, pp. 161981–162003, 2020.
- [24] Alain Riou, Stefan Lattner, Gaëtan Hadjeres, and Geoffrey Peeters, “Pesto: Pitch estimation with self-supervised transposition-equivariant objective,” in *ISMIR*, 2023.
- [25] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry, “End-to-end learning of geometry and context for deep stereo regression,” in *ICCV*, 2017, pp. 66–75.
- [26] Manuel Martinez, Makarand Tapaswi, and Rainer Stiefelhagen, “A closed-form gradient for the 1D earth mover’s distance for spectral deep learning on biological data,” in *ICML 2016 Workshop on Computational Biology (CompBio@ ICML16)*, 2016.
- [27] Graham E. Poliner, Daniel P. W. Ellis, A. F. Ehmman, Emilia Gómez, Sebastian Streich, and Beesuan Ong, “Melody transcription from music audio: Approaches and evaluation,” *IEEE Trans. Speech Audio Process.*, vol. 15, no. 4, pp. 1247–1256, 2007.
- [28] Colin Raffel, Brian McFee, Eric J. Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, and Daniel P. W. Ellis, “MIR.EVAL: A transparent implementation of common MIR metrics,” in *ISMIR*, 2014, pp. 367–372.
- [29] Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Mohammad Norouzi, Douglas Eck, and Karen Simonyan, “Neural audio synthesis of musical notes with wavenet autoencoders,” in *ICML*. 2017, pp. 1068–1077, PMLR.