

One or Two Frequencies? The Scattering Transform Answers

Vincent Lostanlen

Music and Audio Research Lab
New York University
New York, NY, USA
vincent.lostanlen@nyu.edu

Alice Cohen-Hadria

IRIT
University of Toulouse — CNRS
Toulouse, France
alice.cohen-hadria@irit.fr

Juan Pablo Bello

Music and Audio Research Lab
New York University
New York, NY, USA
jpbello@nyu.edu

Abstract—With the aim of constructing a biologically plausible model of machine listening, we study the representation of a multicomponent stationary signal by a wavelet scattering network. First, we show that renormalizing second-order nodes by their first-order parents gives a simple numerical criterion to assess whether two neighboring components will interfere psychoacoustically. Secondly, we run a manifold learning algorithm (Isomap) on scattering coefficients to visualize the similarity space underlying parametric additive synthesis. Thirdly, we generalize the “one or two components” framework to three sine waves or more, and prove that the effective scattering depth of a Fourier series grows in logarithmic proportion to its bandwidth.

Index Terms—Audio systems, Amplitude modulation, Continuous wavelet transform, Fourier series, Multi-layer neural network.

I. INTRODUCTION

In the mammalian auditory system, cochlear hair cells operate like band-pass filters whose equivalent rectangular bandwidth (ERB) grows in proportion to their center frequency. Given two sine waves $t \mapsto \mathbf{y}_1(t) = a_1 \cos(f_1 t + \phi_1)$ and $t \mapsto \mathbf{y}_2(t) = a_2 \cos(f_2 t + \phi_2)$ of respective frequencies $f_1 > 0$ and $f_2 > 0$, we perceive their mixture as a musical chord insofar as \mathbf{y}_1 and \mathbf{y}_2 belong to disjoint critical bands. However, if $a_2 \ll a_1$ or $f_2 \approx f_1$, then the tone \mathbf{y}_2 is said to be *masked* by \mathbf{y}_1 . In lieu of two pure tones, we hear a “beating tone”: i.e., a locally sinusoidal wave whose carrier frequency is $\frac{1}{2}(f_1 + f_2)$ and whose modulation frequency is $\frac{1}{2}|f_1 - f_2|$. In humans, the resolution of beating tones involves physiological processes beyond the cochlea, i.e., in the primary auditory cortex.

The scattering transform (\mathbf{S}) is a deep convolutional operator which alternates constant- Q wavelet decompositions and the application of pointwise complex modulus, up to some time scale T . Broadly speaking, its first two layers (\mathbf{S}_1 and \mathbf{S}_2) resemble the functioning of the cochlea and the primary auditory cortex, respectively. In the context of audio classification, scattering transforms have been successfully employed to represent speech [2], environmental sounds [13], urban sounds [20], musical instruments [10], rhythms [8], and playing techniques [24]. Therefore, the scattering transform simultaneously enjoys a diverse range of practical motivations, a firm rooting in wavelet theory, and a plausible correspondence with neurophysiology.

This article discusses the response of the scattering transform operator to a complex tone input $\mathbf{y} : t \mapsto \mathbf{y}_1(t) + \mathbf{y}_2(t)$, depending on the sinusoidal parameters of \mathbf{y}_1 and \mathbf{y}_2 . In this respect, we follow a well-established methodology in nonstationary signal processing, colloquially known as: “One or two frequencies? The X Answers”, where X is the nonlinear operator of interest. The key idea is to identify transitional regimes in the response of X with respect to variations in relative amplitude ($\frac{a_2}{a_1}$), relative frequency ($\frac{f_2}{f_1}$), and relative phase ($\phi_2 - \phi_1$). Prior publications have done so for X being the empirical mode decomposition [19], the synchrosqueezing transform [25], and the singular spectrum analysis operator [9]. We extend this line of research to the case where X is the scattering transform in dimension one.

II. WAVELET-BASED RECURSIVE INTERFEROMETRY

Let $\boldsymbol{\psi} \in \mathbf{L}^2(\mathbb{R}, \mathbb{C})$ a Hilbert-analytic filter with null average, unit center frequency, and an ERB equal to $1/Q$. We define a constant- Q wavelet filterbank as the family $\boldsymbol{\psi}_\lambda : t \mapsto \lambda \boldsymbol{\psi}(\lambda t)$. Each wavelet $\boldsymbol{\psi}_\lambda$ has a center frequency of λ , an ERB of λ/Q , and an effective receptive field of $(2\pi Q/\lambda)$ in the time domain. In practice, the frequency variable λ gets discretized according

to a geometric progression of common ratio $2^{\frac{1}{Q}}$. Consequently, every continuous signal \mathbf{y} that is bandlimited to $[f_{\min}, f_{\max}]$ activates a number of $Q \log_2(\frac{f_{\max}}{f_{\min}})$ wavelets $\boldsymbol{\psi}_\lambda$ at most.

We define the scalogram of \mathbf{y} as the squared complex modulus of its constant- Q transform (CQT):

$$\mathbf{U}_1 \mathbf{y} : (t, \lambda_1) \longrightarrow \int_{-\infty}^{+\infty} \mathbf{y}(t') \boldsymbol{\psi}_{\lambda_1}(t - t') dt' . \quad (1)$$

Likewise, we define a second layer of nonlinear transformation for \mathbf{y} as the “scalogram of its scalogram”:

$$\mathbf{U}_2 \mathbf{y} : (t, \lambda_1, \lambda_2) \longrightarrow \mathbf{y} * \boldsymbol{\psi}_{\lambda_1} * \boldsymbol{\psi}_{\lambda_2}(t), \quad (2)$$

where the asterisk denotes a convolution product. This construct may be iterated for every integer m by “scattering” the multivariate signal $\mathbf{U}_m \mathbf{y}$ into all wavelet subbands $\lambda_m < \lambda_{m-1}$:

$$\mathbf{U}_{m+1} \mathbf{y} : (t, \lambda_1 \dots \lambda_{m+1}) \longrightarrow \mathbf{U}_m \mathbf{y}(t, \lambda_1 \dots \lambda_m) * \boldsymbol{\psi}_{\lambda_{m+1}}(t, \lambda_1 \dots \lambda_m) . \quad (3)$$

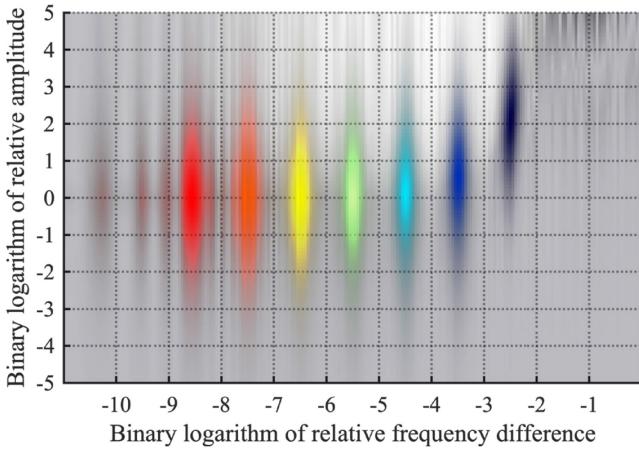


Fig. 1. Superimposed heatmaps of second-order masking coefficients $\Psi_2 y$ after a scattering transform of two sine waves y_1 and y_2 , measured around the frequency f_1 as a function of relative amplitude $\frac{a_2}{a_1}$ and relative frequency difference $\frac{|f_2 - f_1|}{f_1}$. The color of each blot denotes the resolution λ_2 at the second layer. Wavelets have an asymmetric profile (Gammatone wavelets) and a quality factor $Q = 4$. The second layer covers an interval of nine octaves below f_1 . For the sake of clarity, we only display one interference pattern per octave.

Note that the original definition of the scattering transform adopts the complex modulus ($|z| = \sqrt{z\bar{z}}$) rather than its square ($|z|^2 = z\bar{z}$) as its activation function. This is to ensure that \mathbf{U}_m is a non-expansive map in terms of Lipschitz regularity. However, to simplify our calculation and spare an intermediate stage of linearization of the square root, we choose to employ a measure of power rather than amplitude. This idea was initially proposed by [6] in the context of marine bioacoustics.

Every layer m in this deep convolutional network composes an invariant linear system (namely, the CQT) and a pointwise operation (the squared complex modulus). Thus, by recurrence over the depth variable m , every tensor $\mathbf{U}_m y$ is equivariant to the action of delay operators. In order to replace this equivariance property by an invariance property, we integrate each \mathbf{U}_m over some predefined time scale T , yielding the invariant scattering transform:

$$\mathbf{S}_m y : (t, p) \rightarrow \int_{-\infty}^{+\infty} \mathbf{U}_m(t^t, p) \varphi_T(t - t^t) dt^t, \quad (4)$$

where the m -tuple $p = (\lambda_1 \dots \lambda_m)$ is a scattering path and the signal φ_T is a real-valued low-pass filter of time scale T .

III. AUDITORY MASKING IN A SCATTERING NETWORK

Given $n \in \{1, 2\}$, the convolution between every sine wave y_n and every wavelet Ψ_{λ_1} writes as a multiplication in the Fourier domain. Because Ψ_{λ_1} is Hilbert-analytic, the analytic part $\Psi_{\lambda_1}^a = y_n + i\Psi_{\lambda_1} y_n$ of the real signal y_n is preserved in the CQT:

$$(y_n * \Psi_{\lambda_1})(t) = \frac{1}{2} \Psi_{\lambda_1}(f_n) y_n^a(t). \quad (5)$$

By linearity of the CQT, we expand the interference between y_1 and y_2 by heterodyning:

$$\begin{aligned} (\Psi_1 + \Psi_2) * \Psi_{\lambda_1}(t) &= \frac{1}{2} \Psi_{\lambda_1}^2(f_1)^2 \frac{1}{a_1 + \frac{1}{2}} \Psi_{\lambda_1}^2(f_2) \frac{a_2^2}{a_2^2} \\ &+ R \Psi_{\lambda_1} \frac{f_1}{\lambda_1} \Psi_{\lambda_1}^* \frac{f_2}{\lambda_1} \frac{a_1 a_2}{a_1 a_2} \cos((f_2 - f_1)t + (\phi_2 - \phi_1)). \end{aligned} \quad (6)$$

Because the wavelet Ψ has a null average, the two constant terms in the equation above are absorbed by the first layer of the scattering network, and disappear at deeper layers. However, the cross term, proportional to $a_1 a_2$, is a “difference tone” of fundamental frequency $\Delta f = |f_2 - f_1|$.

The authors of a previous publication [3] have remarked that this difference tone elicits a peak in second-order scattering coefficients for the path $p = (\lambda_1, \lambda_2) = (f_1, |f_2 - f_1|)$. In the following, we generalize their study to include the effect of the relative amplitude $\frac{a_2}{a_1}$, the wavelet shape Ψ , the quality factor Q , and the time scale of local stationarity T .

Equation 6 illustrates how the scalogram operator \mathbf{U}_1 converts a complex tone (two frequencies f_1 and f_2) into a simple tone (one frequency $|f_2 - f_1|$). For this simple tone to carry a nonnegligible amplitude in \mathbf{U}_2 , three conditions must be satisfied. First, the rectangular term $a_1 a_2$ must be nonnegligible in comparison to the square terms a_1^2 and a_2^2 . Secondly, there must exist a wavelet Ψ_{λ_1} whose spectrum encompasses both frequencies f_1 and f_2 . Said otherwise, λ_1 must satisfy the inequalities $\frac{|f_n| - 1}{\lambda_1} \ll \frac{1}{Q}$, both for $f_n = f_1$ and for $f_n = f_2$. Thirdly, the frequency difference $|f_2 - f_1|$ must belong to the passband of some second-order wavelet Ψ_{λ_2} . Yet, in practice, to guarantee the temporal localization of scattering coefficients and restrict the filterbank to a finite number of octaves, the scaling factor of every Ψ_{λ_m} is upper-bounded by the temporal constant T . Therefore, the period $\frac{2\pi}{|f_2 - f_1|}$ of the difference tone should be under the pseudo-period of the wavelet with support T ; i.e., a pseudo-period of $Q T$. Hence the third condition: $|f_2 - f_1| \ll \frac{2\pi Q}{T}$.

One simple way of quantifying the amount of mutual interference between signals y_1 and y_2 is to renormalize second-order coefficients by their first-order “parent” coefficients:

$$\frac{\Psi_2 y(t, \lambda_1, \lambda_2)}{\Psi_1 y(t, \lambda_1)} = \frac{S_2 y(t, \lambda_1, \lambda_2)}{S_1 y(t, \lambda_1)} \quad (7)$$

This operation, initially proposed by [2], is conceptually analogous to classical methods in adaptive gain control, notably per-channel energy normalization (PCEN) [14].

In accordance with the “one or two frequencies” methodology, Figure 1 illustrates the value of this ratio of energies in the subband $\lambda_1 = f_1$, for different values of relative amplitude $\frac{a_2}{a_1}$ and relative frequency difference $\frac{|f_2 - f_1|}{f_1}$. We fixed $f_2 < f_1$ without loss of generality. As expected, we observe that, for $a_2 \approx a_1$ and a relative frequency difference between $\frac{Q}{f_1 T}$ and $\frac{1}{Q}$, second-layer wavelets Ψ_{λ_2} resonate with the difference tone as a result of the interference between signals y_1 and y_2 .

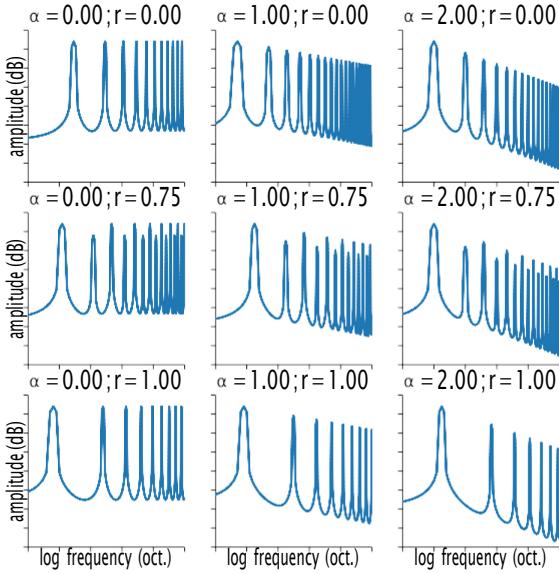


Fig. 2. Log-magnitudes of synthetic musical tones as a function of wavelet log-frequency ($\log \lambda$). Ticks of the vertical (resp. horizontal) axis denote relative amplitude (resp. frequency) intervals of 10 dB (resp. one octave). Parameters α and r denote the Fourier decay exponent and the relative odd-to-even amplitude difference r respectively. See Equation 8 for details.

IV. APPLICATION TO MANIFOLD LEARNING

To demonstrate the ability of the scattering transform to characterize auditory masking, we build a dataset of complex tones according to the following additive synthesis model:

$$\mathbf{y}_{\alpha,r}(t) = \sum_{n=1}^{\infty} \frac{1 + (-1)^n r}{n^\alpha} \cos(n f_1 t) \varphi_T(t), \quad (8)$$

where φ_T is a Hann window of duration T . This additive synthesis model depends upon two parameters: the Fourier decay α and the relative odd-to-even amplitude difference r . Figure 2 displays the CQT log-magnitude spectrum of $\mathbf{y}_{\alpha,r}$ for different values of α and r . In practice, we set T to 1024 samples, N to 32 harmonics, and f_1 between 12 and 24 cycles.

Our synthetic dataset comprises 2500 audio signals in total, corresponding to 50 values of α between 0 and 2 and 50 values of r between 0 and 1, while f_1 is an integer chosen uniformly at random between 12 and 24. We extract the scattering transform of each signal $\mathbf{y}_{\alpha,r}$ up to order $M = 2$, with $Q = 1$ and $J = 8$, by means of the Kymatio Python package [4]. Concatenating QJ first-order coefficients with $\frac{1}{2}Q^2J(J-1)$ second-order coefficients yields a representation in dimension 37.

For visualization purposes, we bring the 37-dimensional space of scattering coefficients to the dimension three by means of the Isomap algorithm for unsupervised manifold learning [22]. The appeal behind Isomap is that pairwise Euclidean distances in the 3-D point cloud approximate the corresponding geodesic distances over the K -nearest neighbor graph associated to the dataset. Throughout this paper, we set the number of neighbors to $K = 100$ and measure neighboring relationships by comparing high-dimensional \mathcal{L}^2 distances.

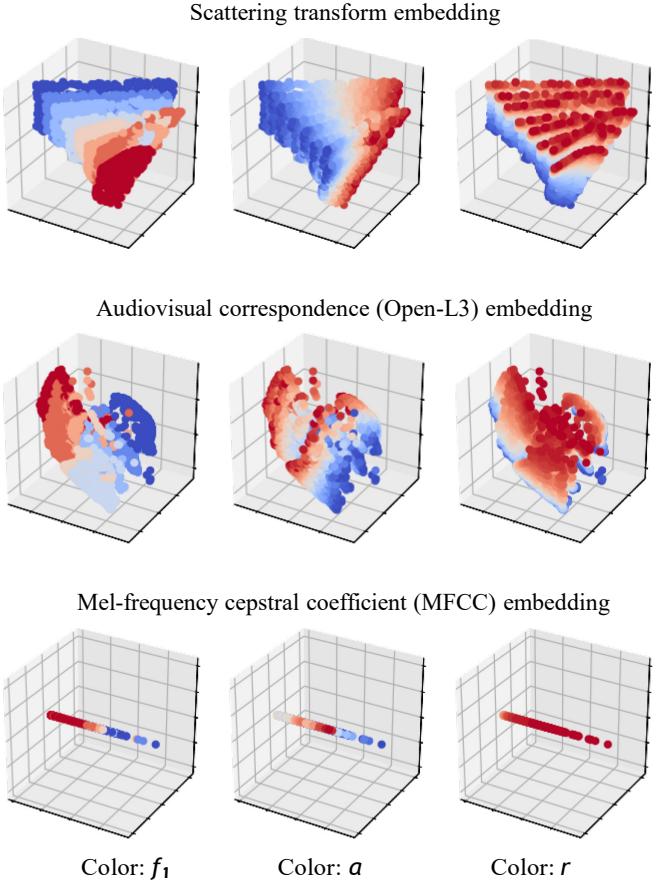


Fig. 3. Isomap embedding of synthetic musical notes, as described by their scattering transform coefficients (top); their Open-L3 coefficients (center); and their mel-frequency cepstral coefficients (MFCC, bottom). The color of a dot, ranging from red to blue via white, denotes the fundamental frequency f_1 (left), the Fourier decay exponent α (center), and the relative odd-to-even amplitude difference r (right) respectively. Note that all methods are unsupervised: triplets (f_1, α, r) are not directly supplied to the models, but only serve for color grading post hoc. See Section IV for details.

Crucially, in the case of the scattering transform, these \mathcal{L}^2 distances are provably stable (i.e., Lipschitz-continuous) to the action of diffeomorphisms [16, Theorem 2.12].

Figure 3 (top) illustrates our findings. We observe that, after scattering transform and Isomap dimensionality reduction, the dataset appears as a 3-D Cartesian mesh whose principal components align with f_1 , α , and r respectively. This result demonstrates that the scattering transform is capable of disentangling and linearizing multiple factors of variability in the spectral envelope of periodic signals, even if those factors are not directly amenable to diffeomorphisms.

As a point of comparison, Figure 3 presents the outcome of Isomap on alternative feature representations: Open-L3 embedding (center) and mel-frequency cepstral coefficients (MFCC, bottom). The former results from training a deep convolutional network (convnet) on a self-supervised task of audiovisual correspondence, and yields 6177 coefficients [7]. The latter results from a log-mel-spectrogram representation, followed by a discrete cosine transform (DCT) over the mel-

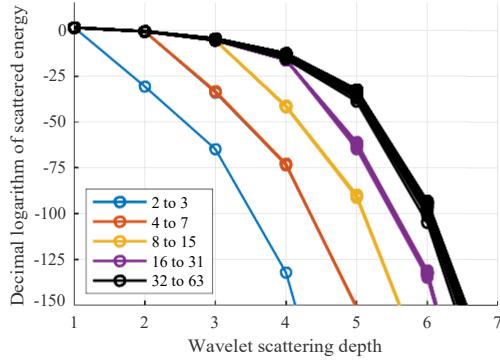


Fig. 4. Energy decay as a function of wavelet scattering depth m , for mixtures of N components with equal amplitudes, equal phases, and evenly spaced frequencies. The color of each line plot denotes the integer part of $\log_2 N$. In this experiment, wavelets have a sine cardinal profile (Shannon wavelets) and a quality factor equal to $Q = 1$. Each filterbank covers seven octaves.

frequency axis, and yields 12 coefficients. We compute MFCC with librosa v0.7 [18] default parameters.

We observe that Open-L3 embeddings correctly disentangles boundary conditions (r) from fundamental frequency (f_1), but fails to disentangle Fourier decay (α) from f_1 . Instead, correlations between r and f_1 are positive for low-pitched sounds (12 to 16 cycles) and negative for high-pitched sounds (16 to 24 cycles). Although this failure deserves a more formal inquiry, we hypothesize that this stems from the small convolutional receptive field of the L^3 -Net: 24 mel subbands, i.e., roughly half an octave around 1 kHz.

Moreover, in the case of MFCC, we find that the variability in fundamental frequency (f_1) dominates the variability in spectral shape parameters (α and r), thus yielding a rectilinear embedding (top). This observation is in line with a previous publication [11], which showed statistically that MFCCs are overly sensitive to frequency transposition in complex tones.

From this qualitative benchmark, it appears that the scattering transform is a more interpretable representation of periodic signals than Open-L3, while incurring a smaller computational cost. However, in the presence of aperiodic signals such as environmental sounds, Open-L3 outperforms the scattering transform in terms of classification accuracy with linear support vector machines [5]. To remain competitive, the scattering transform must not only capture heterodyne interference, but also joint spectrotemporal modulations [1]. In this context, future work will strive to combine insights from multiresolution analysis and deep self-supervised learning.

V. BEYOND PAIRWISE INTERFERENCE: FULL-DEPTH SCATTERING NETWORKS

In speech and music processing, pitched sounds are rarely approximable as a mixture of merely two components. More often than not, they contain ten components or more, and span across multiple octaves in the Fourier domain. Thus, computing the masking coefficient at the second layer only provides a crude description of the timbral content within each critical band. Indeed, \mathbf{S}_2 encodes pairwise interference between

sinusoidal components but fails to characterize more intricate structures in the spectral envelope of \mathbf{y} .

To address this issue, we propose to study the scattering transform beyond order two, thus encompassing heterodyne structures of greater multiplicity. For the sake of mathematical tractability, we consider the following mother wavelet, hereafter called “complex Shannon wavelet” after [15, Section 7.2.2]:

$$\psi : t \mapsto \frac{\exp(2it) - \exp(it)}{2\pi it} \quad (9)$$

The definition of a scattering transform with complex Shannon wavelets requires to resort to the theory of tempered distributions. We refer to [21] for further mathematical details.

The following theorem, proven in the Appendix, describes the response of a deep scattering network in the important particular case of a periodic signal with finite bandwidth.

Theorem V.1. *Let $\mathbf{y} \in C^\infty(\mathbb{R})$ a periodic signal of fundamental frequency f_1 . Let ψ the complex Shannon wavelet as in Equation 9 and \mathbf{U}_1 its associated scalogram operator as in Equation 1. If \mathbf{y} has a finite bandwidth of M octaves, then its scattering coefficients $\mathbf{U}_m \mathbf{y}$ are zero for any $m > M$.*

This result is in agreement with the theorem of exponential decay of scattering coefficients [23]. Note, however, that [23] expresses an upper bound on the energy at fixed depth for integrable signals, while we express an upper bound on the depth at fixed bandwidth for periodic signals.

We apply the theorem above to the case of a signal containing N components of equal amplitudes, equal phases, and evenly spaced frequencies: $\mathbf{y} : t \mapsto \sum_{n=1}^{N-1} a_n \cos(nf_1 t + \phi_1)$. Figure 4 illustrates the decay of scattered energy as a function of depth. The conceptual analogy between depth and scale was originally proposed by [17] in a theoretical effort to clarify the role of hierarchical symmetries in convnets.

Although our findings support this analogy, we note that computing a scattering transform with $M = \log_2 T$ layers is often impractical. However, if the Fourier series in \mathbf{y} satisfies a self-similarity assumption, it is possible to match the representational capacity of a full-depth scattering network while keeping the depth to $M = 2$. Indeed, spiral scattering performs wavelet convolutions over time, over log-frequency, and across octaves, thereby capturing the spectrotemporal periodicity of Shepard tones and Shepard-Risset glissandos [12]. Further research is needed to integrate broadband demodulation into deep convolutional architectures for machine listening.

VI. CONCLUSION

In this article, we have studied the role of every layer in a scattering network by means of a well-established methodology, colloquially known as “one or two components” [19]. We have come up with a numerical criterion of psychoacoustic masking; demonstrated that the scattering transform disentangles multiple factors of variability in the spectral envelope; and proven that the effective scattered depth of Fourier series is bounded by the logarithm of its bandwidth, thus emphasizing the importance of capturing geometric regularity across temporal scales.

REFERENCES

- [1] J. Andén, V. Lostanlen, and S. Mallat. “Joint time-frequency scattering”. In: *IEEE Transactions on Signal Processing* 67.14 (2019), pp. 3704–3718.
- [2] J. Andén and S. Mallat. “Deep scattering spectrum”. In: *IEEE Trans. Signal Process.* 62.16 (2014), pp. 4114–4128.
- [3] J. Andén and S. Mallat. “Scattering representation of modulated sounds”. In: *Proc. DAFX*. 2012.
- [4] M. Andreux et al. “Kymatio: Scattering transforms in Python”. In: *JMLR* 21.60 (2020), pp. 1–6.
- [5] R. Arandjelovic and A. Zisserman. “Look, listen and learn”. In: *Proceedings of the IEEE International Conference on Computer Vision*. 2017, pp. 609–617.
- [6] R. Balestriero and H. Glotin. “Linear time complexity deep Fourier scattering network and extension to nonlinear invariants”. In: *arXiv* 1707.05841 (2017).
- [7] J. Cramer et al. “Look, listen, and learn more: Design choices for deep audio embeddings”. In: *Proc. ICASSP*. IEEE. 2019, pp. 3852–3856.
- [8] D. Haider and P. Balazs. “Extraction of Rhythmic Features with the Gabor Scattering Transform”. In: *Proc. CMMR*. 2019.
- [9] J. Harmouche et al. “Une ou deux composantes: la réponse de l’analyse spectrale singulière”. In: *Actes du colloque GRETSI*. 2015.
- [10] V. Lostanlen, J. Andén, and M. Lagrange. “Extended playing techniques: The next milestone in musical instrument recognition”. In: *Proc. DLfM*. 2018.
- [11] V. Lostanlen and C.-E. Cella. “Deep convolutional networks on the pitch spiral for musical instrument recognition”. In: *Proc. ISMIR*. 2016.
- [12] V. Lostanlen and S. Mallat. “Wavelet scattering on the pitch spiral”. In: *Proc. DAFX*. 2016.
- [13] V. Lostanlen et al. “Relevance-based quantization of scattering features for unsupervised mining of environmental audio”. In: *EURASIP J. Audio Speech Mus. Process.* 2018.1 (2018), p. 15.
- [14] V. Lostanlen et al. “Per-Channel Energy Normalization: Why and How”. In: *IEEE Signal Proc. Let.* 26.1 (2019), pp. 39–43.
- [15] S. Mallat. *A wavelet tour of signal processing: The sparse way*. Associated Press, 2008.
- [16] S. Mallat. “Group invariant scattering”. In: *Comm. Pure Appl. Math.* 65.10 (2012), pp. 1331–1398.
- [17] S. Mallat. “Understanding deep convolutional networks”. In: *Phil. Trans. R. Soc.* 374.2065 (2016), p. 20150203.
- [18] B. McFee et al. *librosa: 0.7.2*. Jan. 2020. doi: 10.5281/zenodo.3606573.
- [19] G. Rilling and P. Flandrin. “One or two frequencies? The empirical mode decomposition answers”. In: *IEEE Trans. Signal Process.* 56.1 (2008), pp. 85–95.
- [20] J. Salamon and J. P. Bello. “Feature learning with deep scattering for urban sound analysis”. In: *Proc. EUSIPCO*. IEEE. 2015, pp. 724–728.
- [21] R. S. Strichartz. *A guide to distribution theory and Fourier transforms*. World Scientific Publishing Company, 2003.
- [22] J. B. Tenenbaum, V. De Silva, and J. C. Langford. “A global geometric framework for nonlinear dimensionality reduction”. In: *science* 290.5500 (2000), pp. 2319–2323.
- [23] I. Waldspurger. “Exponential decay of scattering coefficients”. In: *Proc. SampTA*. IEEE. 2017, pp. 143–146.
- [24] C. Wang et al. “Adaptive time–frequency scattering for periodic modulation recognition in music signals”. In: *Proc. ISMIR*. 2019.
- [25] H.-T. Wu, P. Flandrin, and I. Daubechies. “One or two frequencies? The synchrosqueezing answers”. In: *Adv. Adapt. Data Anal.* 3.01–02 (2011), pp. 29–39.

APPENDIX: PROOF OF THEOREM V.1

Proof. We reason by induction over the depth variable M . The base case ($M = 1$) leads to $\mathbf{U}_1 \mathbf{y}(t, \lambda) = 1$ if $\lambda < f_1 \leq 2\lambda$ and zero otherwise. Because ψ has one vanishing moment, it follows that $\mathbf{U}_2 \mathbf{y}$ is zero, and likewise at deeper layers. To prove the induction step at depth M , to decompose \mathbf{y} into a low-pass approximation ($\mathbf{y} * \mathbf{g}_M$) spanning the subband $[0; 2^M f_1]$ and a high-pass detail ($\mathbf{y} * \mathbf{h}_M$) spanning the subband $[2^M f_1; 2^{(M+1)} f_1]$. Denoting by c_n the complex-valued Fourier coefficients of \mathbf{y} , we have at every time $t \in \mathbb{R}$:

$$\mathbf{y}(t) = (\mathbf{y} * \mathbf{g}_M)(t) + (\mathbf{y} * \mathbf{h}_M)(t) \\ = \sum_{|n| \leq 2^M} c_n \exp(i n f_1 t) + \sum_{|n| > 2^M} c_n \exp(i n f_1 t) \quad (10)$$

On one hand, the coarse term ($\mathbf{y} * \mathbf{g}_M$) has a bandwidth of M octaves. Therefore, by the induction hypothesis, we have $\mathbf{U}_m(\mathbf{y} * \mathbf{g}_M) = 0$ for $m > M$, and *a fortiori* for $m > (M+1)$. On the other hand, we consider the complex Shannon scalogram of ($\mathbf{y} * \mathbf{h}_M$) in some subband $\lambda > 0$:

$$\mathbf{y} * \mathbf{h}_M * \psi_\lambda(t) \leq \sum_{n=1+2^M}^{2^{M+1}} c_n c_n^* \exp(i(n - k) f_1 t) \\ = \sum_{n=1+2^M}^{2^{M+1}} c_n c_{n-2^M}^* \exp(i(n - k) f_1 t) \quad (11)$$

In the double sum above, all integer differences of the form $(n - k)$ range between $-(2^M - 1)$ and $(2^M - 1)$. Thus, $\mathbf{y} * \mathbf{h}_M * \psi_\lambda$ is a periodic signal of fundamental frequency f_1 spanning M octaves. Furthermore, because $\mathbf{h}_M = \psi_{2^M}$,

$\mathbf{y} * \mathbf{h}_M * \psi_\lambda$ has a smaller bandwidth than $\mathbf{y} * \mathbf{h}_M * \psi_{2^M}$; i.e., M octaves or less. By the induction hypothesis, we have:

$$\forall \lambda, \mathbf{U}_{M+1} |\mathbf{y} * \mathbf{h}_M * \psi_\lambda|^2 (\lambda_1, \dots, \lambda_{M+1}) = 0. \quad (12)$$

In the equation above, we recognize the scattering path $p = (\lambda, \lambda_1, \dots, \lambda_{M+1})$ of \mathbf{U}_{M+2} . Finally, because the scattering transform is a nonexpansive operator [16, Prop. 2.5], we have the inequality:

$$\mathbf{U}_{M+2}(\mathbf{y}) \leq \mathbf{U}_{M+2}(\mathbf{y} * \mathbf{g}_M) + \mathbf{U}_{M+2}(\mathbf{y} * \mathbf{h}_M) = 0, \quad (13)$$

which implies $\mathbf{U}_{M+2} \mathbf{y} = 0$, and likewise at deeper layers. We conclude by induction that the theorem holds for any M . \square