

INTERPRETABLE SUPER-RESOLUTION VIA A LEARNED TIME-SERIES REPRESENTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

We develop a novel interpretable and learnable Wigner-Ville Distribution (WVD) that produces a super-resolved quadratic signal representation for time-series analysis. Our approach has two main hallmarks. First, it interpolates between known Time-Frequency Representations (TFRs) in that it can reach super-resolution with increased time and frequency resolution beyond what the Heisenberg uncertainty principle prescribes and thus beyond commonly employed TFRs. Second, it is interpretable thanks to an explicit low-dimensional and physical parametrization of the WVD. We demonstrate that our approach allows to learn highly adapted TFRs and is ready and able to tackle various large-scale classification tasks, where we reach state-of-the-art performance compared to baseline and learned TFRs.

1 INTRODUCTION

With the recent deep learning advances [LeCun et al. \(2015\)](#); [Goodfellow et al. \(2016\)](#) there has been an exponential growth in the use of Deep Networks (DNs) on various time-series. However, the vast majority of DNs do not directly observe the time-series data but instead a handcrafted, a priori designed representation. Indeed, the vast majority of state-of-the-art methods combine DNs with some variant of a *Time-Frequency Representation* TFR [Lattner et al. \(2019\)](#); [Purwins et al. \(2019\)](#); [Liu et al. \(2019\)](#). A TFR is an *image* representation of a time-series obtained by convolving the latter with a filter-bank, such as wavelets or localized complex sinusoids (e.g., Gabor transform). Different filter-banks lead to different TFR families.

The TFR-DN combination is powerful due to three major reasons: (i) the TFR contracts small transformations of the time-series internal events such as translation, time and/or frequency warping [Bruna & Mallat \(2013\)](#) leading to more stable learning and faster DN convergence; (ii) the image representation allows to treat TFRs as a computer vision task where current DNs excels; (iii) given a coherent choice of TFR, the representation of the features of interests, such as phonemes for speech [Waibel et al. \(1989\)](#), form very distinctive shapes in the TFR image, with dimensionality much smaller than the event’s time-series representation. In fact, a single coefficient of the TFR can encode information of possibly thousands of contiguous bins in the time-series representation [Logan et al. \(2000\)](#). Those TFR benefits allow for great performance gains in a vast majority of task and dataset. Nevertheless, the choice of TFR has the potential to dim, or amplify, the above benefits further pushing the performance gains.

Choosing the “best” TFR is a long lasting research problem in signal processing [Coifman & Wickerhauser \(1992\)](#); [Jones & Baraniuk \(1994\)](#); [Donoho \(1994\)](#). While TFR selection and adaptation was originally driven by signal reconstruction and compression, the recent developments of large supervised time-series datasets have led to novel learnable solutions that roughly fall into four camps. First, methods relying on the Wavelet Transform (WT) [Meyer \(1992\)](#). A WT is TFR with a constant-Q filter-bank based on dilations of a mother wavelet. In [Balestriero et al. \(2018\)](#) the learnability of the mother wavelet is introduced by means of cubic spline parameterization of the mother wavelet and learning the shape by learning the values of the spline and its derivative at the knots position. Second, methods relying on band-pass filters without center-frequency to bandwidth (Q) constraint such as the Short-Time Fourier Transform (STFT) [Allen \(1977\)](#). The first method [Khan & Yener \(2018\)](#) performs this by independently learning the center frequencies and bandwidths of a collection of Morlet wavelets (learning of those coefficients independently breaks the constant-Q property). Another method [Ravanelli & Bengio \(2018\)](#) relies on learning the start and cutoff frequency of a bandpass sinc filter apodized with an hamming window. Those methods are thus similarly learning the location of the bandpass but use a different apodization window of a complex sine (Gaussian

or hamming). Third, Zeghidour et al. (2018) proposes to learn Mel filters that are applied onto a spectrogram (modulus of STFT). Those filters linearly combine adjacent filters in the frequency axis which can be interpreted as learning a linear frequency subsampling of the spectrogram; learning the apodization window used to produce the spectrogram has also been developed in Jaillet & Torrèsani (2007); Pei & Huang (2012). Finally, there are also methods relying on unconstrained DN layers applied on the time-series but which are pre-trained such that the induced representation (layer output) resembles an a priori determined target TFR. This has been done for chirplet transforms Baraniuk & Jones (1996) in Glotin et al. (2017) and for Mel-Spectrograms in Çakir & Virtanen (2018).

All the above methods for choosing the “best TFR” for a DN application suffer from at least one of the three following limitations: (i) the inability to interpolate between different TFR families due to family specific parameterization of the learnable filter-banks; (ii) the inability to maintain interpretability of the filter-bank/TFR after learning; (iii) the inability to reach super-resolution in time and frequency to allow more precise representation. There is thus a need to *provide a universal learnable formulation able to interpolate between and within TFRs while preserving interpretability of the learned representation and with the ability to reach super-resolution*. We propose such a representation in this paper by a specific parametrization of the Wigner-Ville Distribution (WVD) Wigner (1932); Flandrin (1998). We validate our method on multiple large scale datasets of speech, bird and marine bioacoustic and general sound event detection, and demonstrate that the proposed representation outperforms current learnable TFR techniques as well as fixed baseline TFRs when combined with various DNs.

Our contributions are:

[C1] We develop a Wigner-Ville Distribution based TFR with explicit interpretable parameterization that can interpolate between any known TFR and reach super-resolution (Sec. 3.1) and derive various properties such as invariance and covariance of the representation and its stability to parameter perturbations (Sec. 3.2).

[C2] We provide an efficient implementation allowing to compute the proposed representation solely by means of Short-Time Fourier Transforms (Sec. 4.1). This allows GPU friendly computation and applicability of the method to large scale time-series dataset. We study the method complexity and provide a detailed pseudo-code and Python implementation (Sec. 4.2). The Python code will be released upon completion of the review process.

[C3] We validate our model and demonstrate how the proposed method outperforms other learnable TFRs as well as fixed expert based transforms on various dataset and across multiple DN architectures. We interpret the learned representations hinting at the key features of the signals needed to solve the task at hand (Sec. 5.1).

2 BACKGROUND ON TIME-FREQUENCY REPRESENTATIONS

Fourier and Spectrogram. Motivated by the understanding of physical phenomena, mathematical analysis tools have been created, notably the Fourier transform Bracewell & Bracewell (1986). This representation corresponds to a signal expansion into the orthogonal family of complex exponentials as $\mathcal{F}_x(\omega) = \int_{-\infty}^{\infty} x(t)e^{-i\omega t}dt$, which provides a powerful representation for stationary signals. In many situations, it does not seem reasonable to assume the entire series to be stationary. For non stationary signal analysis, where the observed signal carries different information throughout its duration, co-existence of the time and frequency variables in the representation is needed. One solution is offered by the Short Time Fourier Transform (STFT) Allen (1977) defined as follows

$$\text{STFT}_{x,w}(t, f) = \int_{-\infty}^{\infty} w(t - \tau)x(\tau)e^{-if\tau}d\tau, \quad (1)$$

with w an apodization window which vanishes when moving away from 0. This representation thus only assumes stationarity within the effective support of w which we denote by σ_T . The squared modulus of the STFT is called the spectrogram as $\text{SP}_x(t, \omega) = |\text{STFT}_x(t, \omega)|^2$. The simplicity and efficiency of its implementation makes the spectrogram one of the most widely used TFR for non-stationary signals. Nonetheless, the spectrogram has a fundamental antagonism between its temporal and frequency resolution which depends on the apodization window spread σ_T . In a spectrogram, large σ_t allows high frequency resolution and poor time resolution and conversely for small σ_t .

Most applications employ a Gabor (or truncated/approximated) apodization window, in which case $w(u) = \frac{1}{\sqrt{2\pi\sigma_t}} e^{-t^2/(2\sigma_t^2)}$. Such spectrograms are denoted as Gabor transforms [Gabor \(1946\)](#). We will denote such a Gaussian window by g_{σ_t} .

Wavelet Transform. A wavelet filter-bank B is obtained by dilating a mother filter ψ_0 with various scales s which correspond to frequencies f via $s = 2^{S(1-\frac{f}{\pi})}$ with S the largest scale to be analyzed. Application of those dilated filters onto a signal leads to the wavelet transform WT [Mallat \(2008\)](#)

$$\text{WT}_x(t, s) = (x \star \psi_s)(t), \text{ where } \psi_s(t) = \frac{1}{\sqrt{s}} \psi_0\left(\frac{t}{s}\right),$$

with $s > 0$. The relative position of the mother wavelet center frequency is not relevant as the scales can be adapted as desired, let consider here that ϕ_0 is placed at the highest frequency to be analyzed. As opposed to the spectrogram, the resolution of the WT varies with frequencies as the filters have a constant bandwidth to center frequency ratio. In a WT, high frequency atoms ψ_λ with λ close to 1 are localized in time offering a good time resolution but low frequency resolution. Conversely, for low frequency atoms with $\lambda \gg 1$, the time resolution is poor but the frequency resolution is high. As natural signals tend to be of small time duration when they are at high frequencies and of longer duration at low frequencies [Daubechies \(1990\)](#), the scalogram is one of the most adapted TFR for natural biological signals [Meyer \(1992\)](#).

Wigner-Ville Transform and Cohen Class. The Wigner-Ville (WV) transform (or quasi probability or distribution) [Wigner \(1932\)](#) describes the state of a quantum particle by a quasi-probability distribution in the phase space formulation of quantum mechanics instead of being described by a vector on the Hilbert space [Moyal \(1949\)](#). It was then leveraged outside of quantum physics as a quadratic time frequency representation for signal processing [Ville \(1948\)](#). The transform combines complex sine filters as the Fourier transform and auto-correlations of the signal as follows

$$\text{WV}_x(t, \omega) = \int_{-\infty}^{\infty} x\left(t - \frac{\tau}{2}\right) x^*\left(t + \frac{\tau}{2}\right) e^{-i\omega\tau} d\tau. \quad (2)$$

Due to its form, computing the WV is demanding as it provides a representation that is highly dimensional in time and in frequency. Nevertheless, and as opposed to the spectrogram or scalogram, the WV has perfect time and frequency resolution. To focus on specific physical properties and obtain different representations, it is possible to convolve the WV with a 2-dimensional low pass filter Π as :

$$C_x(t, f; \Pi) = (\text{WV}_x \star \Pi)(t, f), \quad (3)$$

where \star represents the 2-dimensional convolution. This filtering lessen the time and frequency resolution while removing the present interferences due to the auto-correlation of the different events present in the signal. The collection of representations $\{C_x(\cdot, \cdot; \Pi), \Pi \in \mathbb{L}^2(\mathbb{R}^2), \Pi \geq 0\}$ defines the *Cohen class* [Cohen \(1989\)](#). For example, any spectrogram lives in this class, as well as many other known TFRs.

This paper extends this representation to allow learning and computational efficiency while providing interpretability of the produced representation.

3 LEARNABLE WIGNER-VILLE BASED SIGNAL REPRESENTATION

In this section we develop our proposed signal representation which builds upon the Wigner-Ville Distribution. We first define our representation and study some key properties to finally propose a physics based parametrization that will allow for interpretability and robust learning.

3.1 INTERPRETABLE, UNIVERSAL AND STABLE SIGNAL REPRESENTATION

We now define our transformation, coined the *K-transform*, which corresponds to applying a kernel $\Phi[t, f] \in \mathbb{L}^2(\mathbb{R} \times \mathbb{R}^+)$ onto the Wigner-Ville transform $\text{WV}_x \in \mathbb{L}^2(\mathbb{R} \times \mathbb{R}^+)$ of the signal $x \in L^2(\mathbb{R})$ (recall (2)), for each time $t \in \mathbb{R}$ and frequency $f \in \mathbb{R}$ as follows. This transform corresponds to replacing the time-frequency-invariant convolution of Cohen's class (recall (3)) with a more general time-frequency-varying convolution.

Definition 1. *The K-transform of a signal x with kernel Φ is defined as*

$$K_{x, \Phi}(t, f) = \langle \text{WV}_x, \Phi[t, f] \rangle. \quad (4)$$

The K -transform is thus defined by taking an inner product of the Wigner-Ville representation with a 2-dimensional filter given by $\Phi[t, f]$ for each time t and frequency f . It is real valued as the Wigner-Ville representation is real and so is the kernel Φ . In particular, we are interested at an interpretable kernel Φ and thus propose to parameterize it as a 2-dimensional Gaussian function given by

$$\Phi[t, f](\tau, \omega) = \mathcal{N}([\tau, \omega]^T; \mu(t, f), \Sigma(t, f)), \quad (5)$$

where \mathcal{N} the Gaussian multivariate function [DeGroot & Schervish \(2012\)](#); we also explicit the mean and covariance parameters given by

$$\mu(t, f) = (\mu_t(t, f), \mu_f(t, f))^T, \Sigma(t, f) = \begin{pmatrix} \sigma_t(t, f)^2 & \rho(t, f) \\ \rho(t, f) & \sigma_f(t, f)^2 \end{pmatrix}. \quad (6)$$

Such parametrization of the kernel employed onto the WV arises naturally in various contexts as it corresponds (through some specific mean and covariance configurations) to the analytical kernel associated to known TFRs such as the chirplet transform, the Wavelet with Morlet wavelet or the Gabor transform [Nuttall \(1988\)](#); [Flandrin & Rioul \(1990\)](#); [Jeong & Williams \(1990\)](#); [Baraniuk & Jones \(1996\)](#); [Talakoub et al. \(2010\)](#). Second, the explicit parametrization of the mean and covariance matrix (6) allows to interpret those coefficients as to what are the physical properties of the kernel.

Interpretability. A great advantage of the Gaussian formulation and explicit parametrization from (6) resides in its simplicity to be interpreted and analyzed. In fact, we directly obtain that the vector $\mu(t, f)$ encodes the center frequency and time position information. The covariance matrix encodes the frequency and time bandwidths of the filter as well as the covariance coefficient ρ which encodes the chirpness (how instantaneous frequency of the filter changes with time). Due to the low dimensionality of the parameter vector

$$\theta(t, f) = (\mu_t, \mu_f, \sigma_t, \sigma_f, \rho)(t, f), \quad (7)$$

comparing learned representations can also be done in a straightforward manner by comparing $\theta(t, f)$ and the other parameter realization $\theta'(t, f)$. In fact, we now demonstrate how the representation is contractive w.r.t. those parameters, that is, how close parameters imply close K-transforms. Let denote directly the representation obtained with the Gaussian kernel $K_{x, \theta}$.

Lemma 1. *Given two parameter vectors $\theta(t, f)$ and $\theta'(t, f)$, the distance between the two induced representations is upper bounded by the distance between their parameters as*

$$\|K_{x, \theta} - K_{x, \theta'}\|_{L^2(\mathbb{R}^2)}^2 \leq \kappa \|x\|_{L^2(\mathbb{R})}^2 \left(\int_{t, f} \|\theta(t, f) - \theta'(t, f)\|_2^2 \right)^{-\frac{1}{2}}.$$

with κ the Lipschitz constant of a standard 2D Gaussian (≈ 0.2422) (proof in [App. C.1](#)).

In particular, for unit norm signal x the representation will be contractive w.r.t. the θ parameters. As a result, given a fixed signal x , close parameters θ and θ' imply close representations $K_{x, \theta}$ and $K_{x, \theta'}$ and thus it is sufficient to compare the vectors of parameters between learned representation or between learned and known representation to assert on what information is encoder and how as studied in [Table 3](#).

Universality of the Representation and Continuous Interpolation. We now provide the following result that demonstrates how the formulation from (4) is sufficient to span any common TFR. In particular, due to the Gaussian parametrization, not all TFR are reachable by varying θ , we thus state the following result for our special case but prove the more general case of arbitrary kernel Φ in the proof.

Proposition 1. *Any Gabor transform, Gabor wavelet transform, Gabor chirplet transform (with arbitrary chirpness, from $\rho(t, f)$) and their signal adapted variants can be reached by the K -transform. (proof in [App. C.2](#)).*

The above result is crucial to understand that it is not only possible to reach various TFRs such as constant-Q transforms or spectrograms with Gaussian apodization window but also any representation that has been adapted to the signal [Abramovich et al. \(1998\)](#), that is, with a time-dependent set of filter-bank. We depict in [Fig. 1](#) some examples of Gaussian parametrization. To explicit the above let consider one case of signal based representation given by wavelet packet trees [Ramchandran & Vetterli \(1993\)](#). In this formulation, the representation we denote as P_x will adapt for each time and

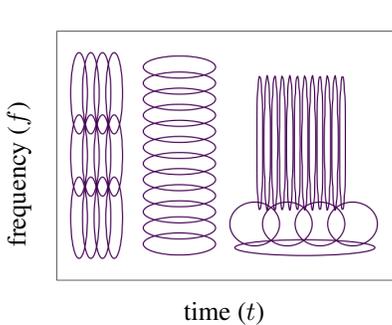


Figure 1: Example of Gaussian parametrization leading to a STFT with high time resolution and poor frequency resolution (left) or STFT with poor time resolution and high frequency resolution (middle) and finally the wavelet (or constant Q-transform) case with adaptive time frequency resolution, for a chirplet transform consider this last case with a non diagonal covariance matrix, with chirpness ρ , the support of each effective Gaussian is depicted. The K-transform uses a 2D-Gaussian for computing each $K(t, f)$ coefficients, with learnable mean and covariance, thus allowing to continuously interpolate between known TFR or learn a new one based on the data and task at hand.

frequency which wavelet family with the given translation and dilation is used to obtain $P_x(t, f)$ such that the final representation P_x is optimal in some sense, for example with minimal Entropy. In this case, the final representation is adapted to the signal with filters varying with time and not just with frequencies. Such representation do not belong in the Cohen class but are reachable with the proposed formulation.

Another crucial property that is beneficial to learning and stability of the representation is to ensure that when moving the parameter θ from one TFR to another, the representation moves continuously and regularly which is formalized in the next result following directly from Cor. 1.

Theorem 1. *The representation $K_{x,\Phi}$ moves continuously with Φ and allow to continuously interpolate between any desired TFR (proof in App. C.3).*

As a result, the proposed representation will be stable during learning where the kernel Φ (or θ) will be incrementally updated based on some rule. Continuity of the representation combine with contractivity from Cor. 1 ensure that small updates of the parameters will also lead to small updates in the representation and thus not lead to unstable learning/gradient based updates.

3.2 EQUIVARIANCE AND INVARIANCE PROPERTIES

Recall that the K-transform does not impose any dependency between different time and frequency kernels $\Phi[t, f]$ as opposed to most WV based transforms such as the affine Cohen class Flandrin & Rioul (1990); Flandrin (1993); Daubechies & Planchon (2002), the Generalized Cohen Class Janssen (1982), Pseudo Wigner Ville Flandrin & Escudié (1984), Smoothed Pseudo Wigner-Ville Hlawatsch et al. (1995). Due to this freedom, the K-transform is not inherently equivariant to signal translation and/or frequency shift. Nevertheless, we now demonstrate that it is straightforward to impose some standard properties onto $K_{x,\Phi}$ as follows where we denote the time convolution as \star_t and the frequency convolution as \star_f .

Proposition 2. *The K-transform can be equipped with the following properties: (i) translation equivariance $\iff \Phi[t - \tau, f] = \Phi[t, f](\cdot - \tau, \cdot) \iff K_{x,\Phi}(\cdot, f) = WV_x \star_t \Phi[0, f]$, (ii) frequency shift equivariance $\iff \Phi[t, f - \omega] = \Phi[t, f](\cdot, \cdot - \omega) \iff K_{x,\Phi}(t, \cdot) = WV_x \star_f \Phi[t, \cdot]$ (proof in C.4).*

As a result, for the K-transform to be equivariant to translation and frequency shifts the kernel Φ must be such that $\Phi[t, f](\tau, \omega) = \Phi[t', f'](\tau - (t' - t), \omega - (f' - f))$ or equivalently, the representation can be written as the 2D convolution of the Wigner-Ville transform with the 2D kernel $\Phi[0, 0]$ as $K_{x,\Phi} = WV_x \star \Phi[0, 0]$ which falls back to the Cohen class (recall (3)) and in particular the Smoothed Pseudo Wigner Ville Andria & Savino (1996). Notice that if the K-transform is constrained to only be translation equivariant, and that the covariance matrices $\Sigma(t, f)$ are diagonal with elements (constant in time) varying with frequency based on the center frequencies $\mu_f(t, f)$ then the K-transform falls back to the Affine Smoothed Pseudo Wigner Ville Flandrin & Rioul (1990).

A key property of most TFR reside in their time-frequency resolution, that is, how precise will be the representation into reflecting the frequency content at each time step present in the studied signal. Standard TFR have limited resolution while the WV has perfect time and frequency resolution. The K-transform can easily reach super resolution by learning small covariance matrices $\Sigma(t, f)$, however, we now demonstrate that this super-resolution will make the representation more sensitive to input perturbations, and vice-versa. Denote a transformation D applied on the signal x and formally characterize the induced perturbation amount in the K-transform by $\|K_{x,\Phi} - K_{D(x),\Phi}\|_{L^2(\mathbb{R}^2)}$ Mallat (2012), let also denote $\det(\Sigma(t, f)) = \sigma_t \sigma_f - \rho^2$.

Proposition 3. *A transformation of the signal $D(x)$ implies a change in the representation proportional to the inverse of $\sigma_t\sigma_f - \rho^2$ as $\|K_{x,\Phi} - K_{D(x),\Phi}\|_{L^2(\mathbb{R}^2)} \leq \max_{t,f} \frac{\kappa}{\det(\Sigma'(t,f))} \times \|x - D(x)\|_{L^2(\mathbb{R})}$, with κ the Lipschitz constant of the WV which exists and is finite for bounded domain (proof in App. C.5)*

Based on the data and task at hand, the K-transform can adapt its parameters to obtain an invariant representation a la scattering network Mallat (2012) thus providing a robust time-serie representation, while having poor resolution, or, conversely, it can reach a super-resolution representation.

We now demonstrate how the K -transform can be computed efficiently solely from Fast Fourier Transforms, for other properties of the transform such as characterization of the interference based on the covariance please see Appendix D.

4 FAST FOURIER TRANSFORM COMPUTATION

The WV, the seed of any possible K-transform, for a discrete signal of length N is a $N \times N$ matrix obtained by doing N Fourier transforms of length N . Its computational complexity is thus quadratic ($\mathcal{O}(N^2 \log(N))$) making it unsuited for large scale tasks. We now demonstrate how to leverage the Gaussian parametrization to greatly speed-up the K-transform computation.

4.1 THE SHORT-TIME-FOURIER-TRANSFORM TRICK

In order to provide a fast implementation, we will first draw the link between a spectral correlation of the STFT and the WV from which we extend our method. Denote the STFT of a signal with a Gabor apodization window of time spread σ by $\text{STFT}_{x,\sigma}$. We demonstrate the equivalence between (i) doing a spectral autocorrelation of $\text{STFT}_{x,\sigma}$ and (ii) doing a 2D Gaussian with diagonal covariance time-frequency-invariant convolution of the WV in Lemma 2 in the Appendix. By employing this STFT reformulation one can obtain the proposed transformation efficiently in practice for finite sample signals as long as σ_f^{-1} is small as this defines the largeness of the window of the STFT. Similarly, whenever $\sigma_f^{-1} - \sigma_t$ is small, computing (8) is fast as it reduces the amount of frequencies to go through during the frequency integral computation. It is now possible to leverage this already convolved (and fast) transform to obtain the K-transform. Notice that any K-transform can be reached that way, as long as 2D Gaussian used in (8) has smaller effective support than the 2D Gaussians of the desired K-transform.

Theorem 2. *Any K-transform (recall (4)) can be obtained from convolving the WV as follows*

$$K_{x,\theta}(t, f) = ((\text{WV}_x \star \mathcal{N}(\cdot; (0, 0)^T, \text{diag}(\sigma_t', \sigma_f'))) \star \mathcal{N}(\cdot; (0, 0)^T, \Sigma'(t, f))) (\mu(t, f))$$

with $\Sigma'(t, f) + \text{diag}(\sigma_t', \sigma_f') = \Sigma(t, f)$. (Proof in Appendix C.7.)

From the above we see how computing the K-transform for various kernels can be done from a base representation which is not the WV but the already convolved WV (8) allowing fast computation as we informally highlight below.

Proposition 4. *The K-transform time and/or frequency resolution is inversely proportional to its speed of computation.*

The above result demonstrates how the fast transform can be used without any downside as long as the learned parameters θ are not too precise in time and/or frequency. It becomes even better if the parameters are shared across time, leading to an equivariant representation. For details on the Gaussian window truncation please see Appendix E.

4.2 COMPUTATIONAL COMPLEXITY AND PSEUDO-CODE

We provide below the explicit pseudo code that summarizes all the involved steps and their impact of the final representation obtained, for the computational complexity see Appendix G, we describe below the pseudo code from Thm. 2. First, do a Gabor transform of x with a Gaussian window $g_{\sigma_{1/f}}$ for apodization leading to $\text{STFT}_{x,\sigma_{1/f}}$. Selection of the $1/\sigma_f$ parameter will determine the final frequency resolution of the transform given by σ_f (recall Prop. 4 and Fig. 1). Second, do the spectral auto-correlation of $\text{STFT}_{x,\sigma_{1/f}}$ with a gaussian spectral apodization window g_{1/σ_t} with spread $1/\sigma_t$ to obtain (8). This maintains the frequency resolution of $\text{STFT}_{x,\sigma_{1/f}}$ while increasing the original time resolution $\frac{1}{\sigma_f}$ by σ_t . Third, compute the K-transform as per Thm. 2 by applying the 2D kernels $\Phi[t, f]$ onto (8).

5 LEARNING AND EXPERIMENTS

We now propose to briefly describe how learning of the K-transform is done and validate the method on various datasets.

Gaussian Parameter Learning. We propose to learn the parameters $\theta(t, f)$ from (7). Since all the involved operations of the K-transform are differentiable and we will pipe the transform to various DN architectures we leverage gradient based learning. In order to keep an unconstrained optimization problem we instead explicitly constraint the parameters as follows $\sigma_f(t, f) = \text{abs}(\sigma'_f(t, f)) + \epsilon$, $\sigma_t(t, f) = \text{abs}(\sigma'_t(t, f)) + \epsilon$, and $\rho(t, f) = \tanh(\rho'(t, f))\sqrt{\sigma_t(t, f)\sigma_f(t, f)}$. With this parametrization, the unconstrained learnable parameters are $\sigma'_f(t, f)$, σ'_t and ρ' . This also allows to always have a nonzero covariance determinant since $\sigma_{\text{time}} > 0$ and $\sigma_{\text{freq}} > 0$ and $\det(C(t, f)) = \sigma_{\text{time}}\sigma_{\text{freq}} - \rho^2 > 0$. We can now train those unconstrained parameters along with any other model parameters with some flavors of gradient descent for any task at hand. In our case we focus on classification.

5.1 EXPERIMENTAL VALIDATION

We propose to validate our method on various classification tasks we briefly describe below. We briefly describe below the used dataset and DN architectures and provide the accuracy results averaged over 10 runs and over multiple learning rates in Table 1. All dataset are described in Appendix H. For each dataset we experiment composing the TFR with three DN architectures (for detailed description of those architectures and additional hyper parameters please see Appendix A). In short, we experiment with TFR-time pooling-linear classifier (**Linear Scattering**), TFR-time pooling-MLP (**Nonlinear Scattering**) and TFR-Conv 2D - time pooling - linear classifier (**Linear Joint Scattering**). Those terminologies are based upon the Scattering Network as they perform time averaging [Bruna & Mallat \(2013\)](#).

We compare our model that we abbreviate as lwvd to the learned Morlet filter-bank [Khan & Yener \(2018\)](#) denoted as lmorlet, and to the learnable sinc based wavelet [Ravanelli & Bengio \(2018\)](#) denoted as lsinc, which are the current state-of-the-art techniques proposing a learnable TFR. In order to calibrate all the results we also compare with a fixed Morlet filter-bank which is the one that is often seen as the most adapted when dealing with speech and bird signals. Here are external benchmarks, for Audio MNIST using AlexNet on spectrogram leads to 95% [Becker et al. \(2018\)](#) accuracy; for bird Vox, using 3 layers of 2D convolution and 2 fully connected layers on top of a fixed a priori designed TFR leads to 90.48% test accuracy [Lostanlen et al. \(2018\)](#); for the Google command dataset, a DenseNet 121 without pretraining nor data augmentation reaches 80% test accuracy [de Andrade et al. \(2018\)](#). In all cases our model reaches comparable performance while relying on a much simpler DNs models that are the scattering network variants, and outperforms other learnable TFRs and the a priori optimal one across dataset and optimization settings, offering significant performance gains. We propose in Appendix B all the figures of the learned filters and their interpretation.

We provide the results in Tab. 1 and Tab. 2 averaged over 10 runs, for each of the runs, the same data split, DN initialization and parameters are used across all TFRs to allow exact performance comparisons. We also provide the results across various learning rate to perceptually measure the sensitivity of each method to this parameter. The first key observation is that the learnable methods are much more sensitive to the learning rate than when using a fixed TFR. Nevertheless, the proposed K-transform is able to outperform all methods across the datasets and for any DN. This comes from the extreme adaptivity of the produced TFR. Notice that the fixed morlet TFR reaches reasonable accuracy especially on speech data without noise (audio MNIST). This is another key feature of learnable TFR, the ability to learn more robust representations. Another key observation comes from the ability of the proposed method to reach state-of-the-art performances while leveraging a simple (few layer) DN in the Linear Joint Scattering case.

6 CONCLUSIONS

We proposed a novel approach to learn generic WVD based TFR, derived an efficient implementation and demonstrated its ability to outperform standard and other learnable TFR techniques across dataset and architecture settings. In addition of learning any desired TFR, our framework is interpretable, telling what type of TFR is, and access to direct parameters allow to exactly position a learned transform among the standard ones.

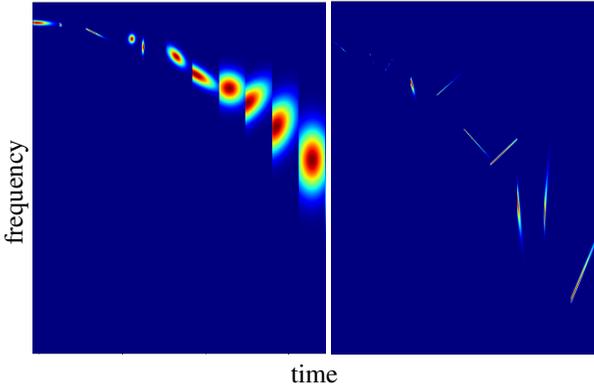


Figure 2: Depiction of the learned filters $\Psi[., f]$ (Def. 1). The full filters banks can be found in Appendix B. **Left, our learned representation on AudioMNIST dataset (Becker et al., 2018)**: This noiseless speech pushes the filters with high frequency resolution (easily seen from the filter covariance as per (7)). For some of these filters the time resolution is also very high (reaching super-resolution), while others favor local translation invariance. For all the medium to low frequency filters, great frequency and time invariance is preferred with large gaussian support, with a slight chirpiness for the medium frequency filters. Thus, the low frequency filters tend to favor time resolution. **Right, our learned representation on Birdvox dataset (Lostanlen et al., 2018)**: We see that the detection of bird songs heavily relies on chirps. Indeed the characteristic sound of birds is increasing or decreasing in frequency over time. Moreover, the learned representation demonstrate filters that reach super-resolution. Thus it represents extreme sensitivity to the time and/or frequency position of the events per Prop. 3). It fits with the extreme time-frequency accuracy of bird audition and acoustic prediction (DOOLING & LOHR, 2006).

Table 1: Average over 10 runs of classification result using three architectures, one layer scattering followed by a linear classifier (Linear Scattering), one layer scattering followed by a two layer neural network (Nonlinear Scattering) and a two layer scattering with joint (2D) convolution for the second layer followed by a linear classifier (Linear Joint Scattering). For each architecture and dataset, we experiment with the baseline, Morlet wavelet filter-bank (morlet), and learnable frameworks being ours (lwvd), learnable sinc based filters (sinc) and learnable Morlet wavelet (lmorlet). As can be seen, across the dataset, architectures and learning rates, the proposed method provides significant performance gains (stds in Tab. 2).

	le. rate	Linear Scattering				Nonlinear Scattering				Linear Joint Scattering			
		<i>morlet</i>	lwvd	sinc	lmorlet	<i>morlet</i>	lwvd	sinc	lmorlet	<i>morlet</i>	lwvd	sinc	lmorlet
DOCC10	0.0002	14.3	63	31.1	29.7	54.1	84.7	74.4	74.9	70.7	83.7	82.4	75.8
	0.001	12.7	65.5	26.0	28.3	50.1	87.9	77.4	77.4	70.1	80.6	80.8	73.2
	0.005	13.0	65.9	17.1	27.0	51.8	87.1	43.3	83.2	65.9	78.0	70.5	80.8
BirdVox	0.0002	63.8	77.9	69.6	65.4	84.7	92.9	88.1	85.8	82.1	90.5	87.2	84.3
	0.001	65.0	80.0	67.2	64.3	85.0	94.2	88.1	86.6	80.3	88.7	86.8	83.1
	0.005	65.2	80.4	67.3	66.9	84.8	94.2	86.0	87.2	78.1	87.5	78.3	82.8
MNIST	0.0002	43.9	68.4	52.2	44.0	82.3	85.3	10.4	83.0	95.3	97.6	22.1	95.4
	0.001	41.5	68.8	43.5	42.2	83.2	89.8	87.1	85.4	89.7	97.8	93.2	90.4
	0.005	34.6	68.8	23.9	36.0	82.7	22.1	68.7	88.1	81.1	12	64.4	80.2
command	0.0002	8.1	24.9	9.5	7.6	33.9	38.2	36.2	33.4	65.8	76.7	3.7	66.8
	0.001	7.5	26.1	8.0	8.2	33.5	42.9	35.5	33.7	53.6	71.8	27.9	51.9
	0.005	7.3	25.7	6.2	6.5	33.0	17.0	28.9	34.8	32.1	35.2	17.2	32.9
fsd	0.0002	9.7	15.3	10.3	9.0	22.9	23.1	2.3	27.9	40.1	38.8	1.6	42.0
	0.001	9.8	16.7	10.4	10.6	24.2	27.4	13.1	31.1	38.9	44.9	2.1	42.3
	0.005	9.0	17.4	5.5	10.2	24.2	28.8	16.9	30.4	25.0	31.5	17.0	33.2

REFERENCES

- Felix Abramovich, Theofanis Sapatinas, and Bernard W Silverman. Wavelet thresholding via a bayesian approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(4):725–749, 1998.
- Jonathan Allen. Short term spectral analysis, synthesis, and modification by discrete fourier transform. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 25(3):235–238, 1977.
- Gregorio Andria and Mario Savino. Interpolated smoothed pseudo wigner-ville distribution for accurate spectrum analysis. *IEEE transactions on instrumentation and measurement*, 45(4):818–823, 1996.
- Randall Balestrierio, Romain Cosentino, Herve Glotin, and Richard Baraniuk. Spline filters for end-to-end deep learning. In *Proc. 35th Int. Conf. on Machine Learning*, volume 80, pp. 364–373, 10–15 Jul 2018. URL <http://proceedings.mlr.press/v80/balestrierio18a.html>.
- Richard G Baraniuk and Douglas L Jones. Wigner-based formulation of the chirplet transform. *IEEE Transactions on signal processing*, 44(12):3129–3135, 1996.
- Sören Becker, Marcel Ackermann, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech Samek. Interpreting and explaining deep neural networks for classification of audio signals. *arXiv preprint arXiv:1807.03418*, 2018.
- Ronald Newbold Bracewell and Ronald N Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.
- Joan Bruna and Stéphane Mallat. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.
- Emre Çakir and Tuomas Virtanen. End-to-end polyphonic sound event detection using convolutional recurrent neural networks with learned time-frequency representation input. In *JCNN*, pp. 1–7. IEEE, 2018.
- Leon Cohen. Time-frequency distributions-a review. *Proceedings of the IEEE*, 77(7):941–981, 1989.
- Leon Cohen. *Time-frequency analysis*, volume 778. Prentice hall, 1995.
- Ronald R Coifman and M Victor Wickerhauser. Entropy-based algorithms for best basis selection. *IEEE Transactions on information theory*, 38(2):713–718, 1992.
- Elena Cordero and Fabio Nicola. Sharp integral bounds for wigner distributions. *International Mathematics Research Notices*, 2018(6):1779–1807, 2018.
- Ingrid Daubechies. The wavelet transform, time-frequency localization and signal analysis. *IEEE transactions on information theory*, 36(5):961–1005, 1990.
- Ingrid Daubechies and Fabrice Planchon. Adaptive gabor transforms. *Applied and Computational Harmonic Analysis*, 13(1):1–21, 2002.
- Douglas Coimbra de Andrade, Sabato Leo, Martin Loesener Da Silva Viana, and Christoph Bernkopf. A neural attention model for speech command recognition. *arXiv preprint arXiv:1808.08929*, 2018.
- Morris H DeGroot and Mark J Schervish. *Probability and statistics*. Pearson Education, 2012.
- David L Donoho. On minimum entropy segmentation. In *Wavelet Analysis and Its Applications*, volume 5, pp. 233–269. Elsevier, 1994.
- Robert J. DOOLING and Bernard LOHR. Auditory temporal resolution in the Zebra Finch (*Taeniopygia guttata*): A model of enhanced temporal acuity. *Ornithological Science*, 5(1):15 – 22, 2006. doi: 10.2326/osj.5.15. URL <https://doi.org/10.2326/osj.5.15>.

- Patrick Flandrin. Temps–fréquence (traité des nouvelles technologies, série traitement du signal). 1993.
- Patrick Flandrin. *Time-frequency/time-scale analysis*. Academic press, 1998.
- Patrick Flandrin and Bernard Escudié. An interpretation of the pseudo-wigner-ville distribution. *Signal Processing*, 6(1):27–36, 1984.
- Patrick Flandrin and Oliver Rioul. Affine smoothing of the wigner-ville distribution. In *International Conference on Acoustics, Speech, and Signal Processing*, pp. 2455–2458. IEEE, 1990.
- Eduardo Fonseca, Manoj Plakal, Frederic Font, Daniel PW Ellis, Xavier Favory, Jordi Pons, and Xavier Serra. General-purpose tagging of freesound audio with audioset labels: Task description, dataset, and baseline. *arXiv preprint arXiv:1807.09902*, 2018.
- Dennis Gabor. Theory of communication. part 1: The analysis of information. *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering*, 93(26):429–441, 1946.
- Hervé Glotin, Julien Ricard, and Randall Balestriero. Fast chirplet transform injects priors in deep learning of animal calls and speech. In *ICLR (Workshop)*, 2017.
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*, volume 1. MIT Press, 2016. <http://www.deeplearningbook.org>.
- W. Heisenberg. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *Zeitschrift für Physik*, 43(3-4):172–198, Mar 1927. doi: 10.1007/BF01397280.
- Franz Hlawatsch, Thulasinath G Manickam, Rüdiger L Urbanke, and William Jones. Smoothed pseudo-wigner distribution, choi-williams distribution, and cone-kernel representation: Ambiguity-domain analysis and experimental comparison. *Signal Processing*, 43(2):149–168, 1995.
- Florent Jaillet and Bruno Torrèsani. Time-frequency jigsaw puzzle: Adaptive multiwindow and multilayered gabor expansions. *Int. J. of Wavelets, Multiresolution and Inf. Proc.*, 5(02):293–315, 2007.
- AJEM Janssen. On the locus and spread of pseudo-density functions in the time-frequency plane. *Philips Journal of Research*, 37(3):79–110, 1982.
- Jechang Jeong and William J Williams. variable-windowed spectrograms: connecting cohen’s class and the wavelet transform. In *Fifth ASSP Workshop on Spectrum Estimation and Modeling*, pp. 270–274. IEEE, 1990.
- Douglas L Jones and Richard G Baraniuk. A simple scheme for adapting time-frequency representations. *IEEE Transactions on Signal Processing*, 42(12):3530–3535, 1994.
- Haidar Khan and Bulent Yener. Learning filter widths of spectral decompositions with wavelets. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Inf. Proc. Sys.* 31, pp. 4601–4612. 2018. URL <http://papers.nips.cc/paper/7711-learning-filter-widths-of-spectral-decompositions-with-wavelets.pdf>.
- Erwin Kreyszig. *Introductory functional analysis with applications*, volume 1. wiley New York, 1978.
- Stefan Lattner, Monika Dörfler, and Andreas Arzt. Learning complex basis functions for invariant representations of audio. *arXiv preprint arXiv:1907.05982*, 2019.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- Caifeng Liu, Lin Feng, Guochao Liu, Huibing Wang, and Shenglan Liu. Bottom-up broadcast neural network for music genre classification. *arXiv preprint arXiv:1901.08928*, 2019.

- Beth Logan et al. Mel frequency cepstral coefficients for music modeling. In *Ismir*, volume 270, pp. 1–11, 2000.
- Vincent Lostanlen, Justin Salamon, Andrew Farnsworth, Steve Kelling, and Juan Pablo Bello. Birdvox-full-night: A dataset and benchmark for avian flight call detection. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 266–270. IEEE, 2018.
- S. Mallat. Group invariant scattering. *Comm. Pure Appl. Math.*, 65(10):1331–1398, July 2012.
- Stephane Mallat. *A wavelet tour of signal processing: the sparse way*. Academic press, 2008.
- Stephane G Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE transactions on pattern analysis and machine intelligence*, 11(7):674–693, 1989.
- Ives Meyer. *Wavelets and applications*, volume 31. Masson Paris, 1992.
- José E Moyal. Quantum mechanics as a statistical theory. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 45, pp. 99–124. Cambridge University Press, 1949.
- Albert H Nuttall. Wigner distribution function: Relation to short-term spectral estimation, smoothing, and performance in noise. Technical report, Naval Underwater sys. center New London Lab, 1988.
- Soo-Chang Pei and Shih-Gu Huang. Stft with adaptive window width based on the chirp rate. *IEEE Transactions on Signal Processing*, 60(8):4065–4080, 2012.
- Hendrik Purwins, Bo Li, Tuomas Virtanen, Jan Schlüter, Shuo-Yiin Chang, and Tara Sainath. Deep learning for audio signal processing. *IEEE J. Sel. Topics in Signal Proc.*, 13(2):206–219, 2019.
- Kannan Ramchandran and Martin Vetterli. Best wavelet packet bases in a rate-distortion sense. *IEEE Transactions on Image Processing*, 2(2):160–175, 1993.
- Mirco Ravanelli and Yoshua Bengio. Interpretable convolutional filters with sincnet. *arXiv preprint arXiv:1811.09725*, 2018.
- D Sen. The uncertainty relations in quantum mechanics. *Current Science*, pp. 203–218, 2014.
- Omid Talakoub, Jie Cui, and Willy Wong. Approximating the time-frequency representation of biosignals with chirplets. *EURASIP Journal on Advances in Signal Processing*, 2010:1–10, 2010.
- Jean Ville. Theorie et application dela notion de signal analytique. *Câbles et transmissions*, 2(1): 61–74, 1948.
- Alex Waibel, Toshiyuki Hanazawa, Geoffrey Hinton, Kiyohiro Shikano, and Kevin J Lang. Phoneme recognition using time-delay neural networks. *IEEE transactions on acoustics, speech, and signal processing*, 37(3):328–339, 1989.
- Pete Warden. Speech commands: A public dataset for single-word speech recognition. *Tensrflow*, 2017.
- E. Wigner. On the Quantum Correction For Thermodynamic Equilibrium. *Physical Review*, 40(5): 749–759, Jun 1932. doi: 10.1103/PhysRev.40.749.
- Neil Zeghidour, Nicolas Usunier, Iasonas Kokkinos, Thomas Schaiz, Gabriel Synnaeve, and Emmanuel Dupoux. Learning filterbanks from raw speech for phone recognition. In *ICASSP*, pp. 5509–5513. IEEE, 2018.

Supplementary Material

This appendix proposes to first review the implementation details and visual results of the paper, studying the learned filters, we conclude with all the proofs of the theoretical results.

A DN TOPOLOGY

We leverage a time translation covariant form of the proposed learnable model as we aim at solving classification tasks based on audio clips. Hence the representation should be translation invariant. We keep the unconstrained frequency dimensions and thus do not impose any frequency shift invariance as in the Cohen class family of representations. The kernels are parametrized as given in the main text and the networks are given as follows:

```

TF: any representation (morlet, lwvd, ...)
the first mean(3) represents time pooling

- onelayer_nonlinear_scattering
  input = T.log(TF.mean(3).reshape([N, -1])+0.1)
  Dropout(0.3)
  Dense(256)
  BatchNormalization([0])
  LeakyReLU
  Dropout(0.1)
  Dense(n_classes)

- onelayer_nonlinear_scattering:
  input = T.log(TF.mean(3).reshape([N, -1])+0.1)
  Dropout(0.1)
  Dense(n_classes)

- joint_linear_scattering:
  feature = T.log(TF.mean(3).reshape([N, -1])+0.1)

  input = T.log(TF+0.1)
  Conv2D(64, (32,16))
  BatchNormalization([0,2,3])
  AbsoluteValue
  Concatenate(AbsoluteValue, feature)
  Dropout(0.1)
  Dense(n_classes)

```

all training are done with the Adam optimizer, same initialization and data splitting.

B ADDITIONAL FIGURES

We represent in this section the leaned filters/kernels Φ applied on the smoothed pseudo Wigner-Ville distribution, for clarity we only depict one every 4 filters, concatenated horizontally. We do so for three dataset and provide analysis in the caption of each figures.

B.1 SAMPLES OF LEARNT FILTERS

We propose in Fig. 3 and Fig. 6 and Fig. 5 the filters after learning for each dataset with their analysis.

C PROOFS

In this section we present in details all the proofs of the main paper results, as well as providing some additional ones.

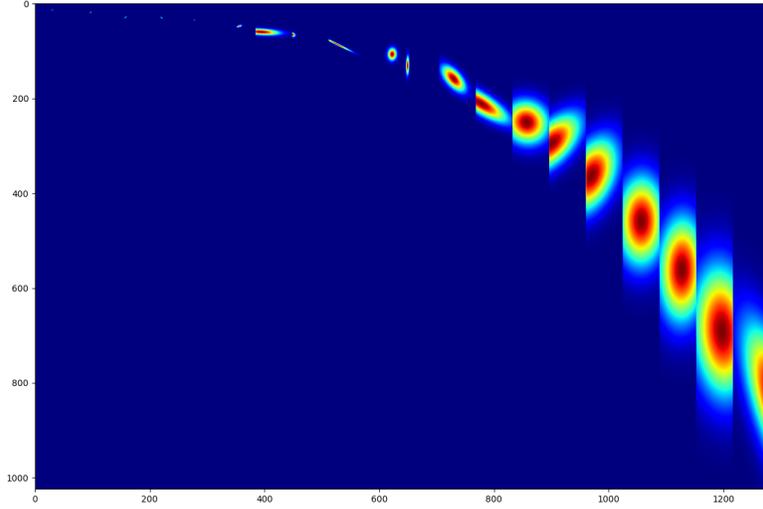


Figure 3: Audio MNIST: this dataset deals with spoken digit classification. A few key observations: the high frequency filters tend to take an horizontal shape greatly favoring frequency resolution, for some of the filters the time resolution is also very high (reaching super-resolution) while others favor local translation invariance. For all the medium to low frequency filters, great frequency and time invariance is preferred (large gaussian support) with a slight chirpness for the medium frequency filters. The low frequency filters tend to favor time resolution.

Lemma 2. *The spectral correlated STFT corresponds to a smoothed Wigner-Ville Distribution as*

$$\begin{aligned} (\mathbf{WV}_x \star \mathcal{N}(\cdot; (0, 0)^T, \text{diag}(\sigma_t, \sigma_f))) (t, f) &= \int_{-\infty}^{\infty} g_{\sigma_f^{-1} - \sigma_t}(\omega) e^{j2\pi\omega t} \text{STFT}_{x, \sigma_f^{-1}}(t, f + \frac{\omega}{2}) \\ &\quad \times \text{STFT}_{x, \sigma_f^{-1}}^*(t, f - \frac{\omega}{2}) d\omega, \quad (8) \end{aligned}$$

where g_σ is a 1-dimensional Gaussian function with spread σ and $\sigma_t \leq \sigma_f^{-1}$. (Proof in App. C.6).

C.1 PROOF OF COR. 1

Let first prove the general case with arbitrary kernels

Lemma 3. *The norm of the difference of two representations obtained from kernel Φ and Φ' is bounded above as $\|\mathbf{K}_{x, \Phi} - \mathbf{K}_{x, \Phi'}\|_{L^2(\mathbb{R}^2)} \leq \|x\|_{L^2(\mathbb{R})}^2 \|\Phi - \Phi'\|_{L^2(\mathbb{R}^4)}$. with $\|\Phi - \Phi'\|_{L^2(\mathbb{R}^4)} = \sqrt{\int_{t, f} \|(\Phi - \Phi')[t, f]\|_{L^2(\mathbb{R}^2)}^2}$.*

Proof. First, one can easily derive $\|\mathbf{WV}_x\|_{L^2(\mathbb{R})} = \|x\|_{L^2(\mathbb{R}^2)}^2$ (see for example Cordero & Nicola (2018)). Given this, and the definition of the K-transform, we obtain that

$$\begin{aligned} \|\mathbf{K}_{x, \Phi} - \mathbf{K}_{x, \Phi'}\|_{L^2(\mathbb{R}^2)} &= \sqrt{\int_{t, f} (\langle \mathbf{WV}_x, \Phi[t, f] \rangle_{L^2(\mathbb{R}^2)} - \langle \mathbf{WV}_x, \Phi'[t, f] \rangle_{L^2(\mathbb{R}^2)})^2} \\ &= \sqrt{\int_{t, f} \langle \mathbf{WV}_x, \Phi[t, f] - \Phi'[t, f] \rangle_{L^2(\mathbb{R}^2)}^2} \\ &\leq \sqrt{\int_{t, f} \|\mathbf{WV}_x\|_{L^2(\mathbb{R}^2)}^2 \|\Phi[t, f] - \Phi'[t, f]\|_{L^2(\mathbb{R}^2)}^2} \quad \text{Cauchy-Schwarz Inequality} \end{aligned}$$

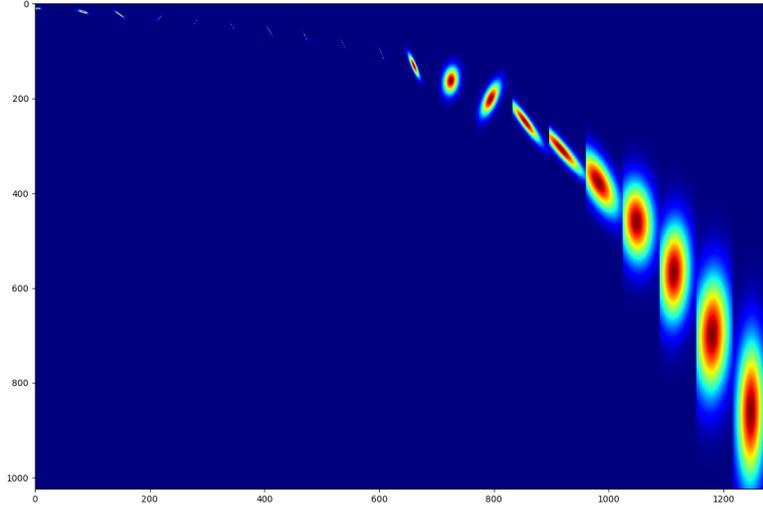


Figure 4: FreeSound: this dataset contains various different classes ranging on different frequencies and without an a priori preferred form of the events in the WV space. As opposed to the AudioMNIST case, we can see that the kernels tend to have smaller covariance (support) hence preferring time and frequency resolution to invariance. This becomes especially true for the high frequency atoms. We also see the clear chirpness for the medium/high frequency kernels with a specific (-30) angle, with decreasing slope. This might be specific to some particular events involving moving objects such as train, cars and so on.

$$\begin{aligned}
&= \|\mathbf{WV}_x\|_{L^2(\mathbb{R}^2)} \sqrt{\int_{t,f} \|\Phi[t, f] - \Phi'[t, f]\|_{L^2(\mathbb{R}^2)}^2} \\
&= \|x\|_{L^2(\mathbb{R})}^2 \sqrt{\int_{t,f} \|\Phi[t, f] - \Phi'[t, f]\|_{L^2(\mathbb{R}^2)}^2} \\
&= \|x\|_{L^2(\mathbb{R})}^2 \|\Phi - \Phi'\|_{L^2(\mathbb{R}^4)}
\end{aligned}$$

□

Now the proof for the special case of a 2D Gaussian kernel follows the exact same procedure as the one above for the Lipschitz constant of the general K-transform.

Proof. In the same way we directly have

$$\|\mathbf{K}_{x,\Phi} - \mathbf{K}_{x,\Phi'}\|_{L^2(\mathbb{R}^2)} \leq \|x\|_{L^2(\mathbb{R})}^2 \sqrt{\int_{t,f} \|\Phi[t, f](\theta) - \Phi[t, f](\theta')\|_{L^2(\mathbb{R}^2)}^2}$$

now the only different is that we need a last inequality to express the distance in term of the θ parameters and not the kernels Φ . To do so, we leverage the fact that the parametric kernel is a 2-dimensional Gaussian with zero mean and unit variance and leverage the following result:

$$\|f(\theta) - f(\theta')\| \leq \max_u \|\nabla_f(u)\| \|\theta - \theta'\|$$

and simply denote by κ the maximum of the Gaussian gradient norm, leading to the desired result by setting f the 2-dimensional Gaussian. □

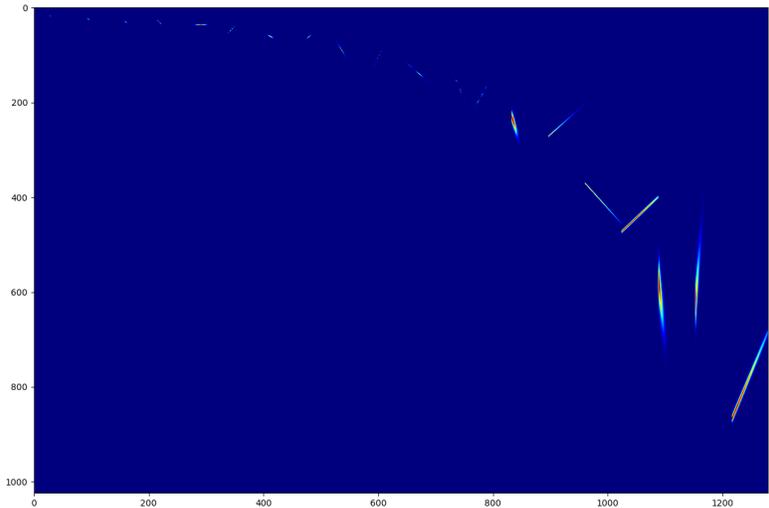


Figure 5: Bird: This dataset proposes to predict the presence or absence of a bird in short audio clips. A priori, detection of such events heavily relies on chirps, the characteristic sound of birds with increasing/decreasing frequency over time. We can see from the learned filters how the learned representation indeed focuses on such patterns, reach super-resolution and thus extreme sensitivity to the time and frequency position of the events, in particular for medium to low frequency kernels. For high frequency kernels, the time and frequency resolution is further increased.

C.2 PROOF OF PROP. 1

First we prove the more general result which follows.

Lemma 4. *For any TFR or signal adapted TFR, there exists a kernel Φ such that $K_{x,\Phi}$ (recall (4)) is equal to it.*

Proof. The proof is a direct application of Moyal Theorem which states that given a signal x and a filter y the application of the filter onto the signal and squaring the result can be expressed as the inner product between the Wigner-Ville transforms of the signal and filter as in

$$\left| \int_{-\infty}^{\infty} x(t)y^*(t)dt \right|^2 = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{WV}_x(\tau, \omega) \mathbf{WV}_y(\tau, \omega) d\tau d\omega.$$

From the above we can see that any time frequency representation (time invariant or not) can be recovered simply by setting $\Phi[t, f] = \mathbf{WV}_y$ of some desired filter y . \square

Now for the special case of the Gabor based transform we can leverage the above result and proof with the following.

Proof. This result is a direct application of the analytically derived kernels from Nuttall (1988); Flandrin & Rioul (1990); Jeong & Williams (1990); Baraniuk & Jones (1996); Talakoub et al. (2010) for various analytical time-frequency representations. For the Gabor transform, Morlet wavelet and Morlet based Chirplet atom, all the analytically derived kernels that are applied onto a Wigner-Ville distribution as a parametric form of the 2-dimensional Gaussian kernel proving the result. \square

C.3 PROOF OF THEOREM 1

Proof. The proof follows directly from the above result. In fact, we obtained above that the Lipschitz constant of the K-transform w.r.t. the Φ parameter is obtained by $\|x\|_{L^2(\mathbb{R})}^2$. Thus, the Lipschitz continuity implies continuity of the transform w.r.t. the Φ parameters which proves the result (see for example Kreyszig (1978)). \square

Table 2: Std over 10 runs of classification result using three architectures, one layer scattering followed by a linear classifier (Linear Scattering), one layer scattering followed by a two layer neural network (Nonlinear Scattering) and a two layer scattering with joint (2D) convolution for the second layer followed by a linear classifier (Linear Joint Scattering). For each architecture and dataset, we experiment with the baseline, Morlet wavelet filter-bank (morlet), and learnable frameworks being ours (lwvd), learnable sinc based filters (sinc) and learnable Morlet wavelet (lmorlet). As can be seen, across the dataset, architectures and learning rates, the proposed method provides significant performance gains.

	le. rate	Linear Scattering				Nonlinear Scattering				Linear Joint Scattering			
		morlet	lwvd	sinc	lmorlet	morlet	lwvd	sinc	lmorlet	morlet	lwvd	sinc	lmorlet
DOCCI0	0.0002	3	0.2	1.1	1.1	0.8	0.2	24.6	0.8	0.2	0.1	40	0.1
	0.001	1.8	0.1	1.5	1.4	1.5	0.2	25.6	0.6	0.8	0.2	26.7	0.3
	0.005	1.4	0.9	1.6	0.6	1.2	28.8	14.3	0.5	0.4	25.8	34.2	0.5
BirdVox	0.0002	0.3	0	1	0.2	0.2	0.2	42.6	0.1	0.3	0.3	0.3	0.2
	0.001	0.7	0.2	1.9	1.6	0.2	0.2	29.1	0.3	0.4	0.4	0.8	0.3
	0.005	0.8	0.9	1.1	0.9	0.2	47.1	28.5	0.3	0.4	37.9	25.9	0.5
MNIST	0.0002	0.8	0.4	0.6	1	0.1	0.1	27.5	0.2	0.2	0.1	38.6	0.2
	0.001	1.9	0.4	1.6	1.5	0.3	0.3	0.6	0.3	0.6	0.1	0.6	1.1
	0.005	2.4	0.9	3.8	2.5	0.3	38.4	4.8	0.2	1.4	31.9	25.7	2.1
command	0.0002	0.3	0.1	0.7	0.5	0.3	0.4	0.2	0.3	0.7	0.2	8.1	0.6
	0.001	0.7	0.2	0.4	0.8	0.2	0.2	0.2	0.1	1.8	0.5	24.8	3
	0.005	0.5	0.4	0.6	0.1	0.1	18.8	0.8	0.3	3.4	30.6	14.9	1.1
fsd	0.0002	0.5	0.2	2.9	0.2	0.4	0.3	7	0.5	1	0.8	4.7	1
	0.001	0.9	0.6	1.1	0.5	0.4	0.6	16.1	0.8	0.9	0.6	6.2	1.3
	0.005	1.1	1.7	1.3	1.1	0.6	0.7	1.2	0.7	0.8	1.2	5.9	1.2

C.4 PROOF OF LEMMA ON INVARIANT

Proof. We provide the proof for each case below:

- Time translation of the signal x by τ to have $y(t) = x(t - \tau)$ leads to

$$\begin{aligned} \mathbf{WV}_y(t, \omega) &= \int_{-\infty}^{\infty} y(t + \frac{\tau}{2}) y^*(t - \frac{\tau}{2}) e^{-i\tau\omega} d\tau \\ &= \int_{-\infty}^{\infty} x(t + \frac{\tau}{2} - \tau) x^*(t - \frac{\tau}{2} - \tau) e^{-i\tau\omega} d\tau \\ &= \mathbf{WV}_x(t - \tau, \omega) \end{aligned}$$

Thus, the time equivariance of the representation defined as $\mathbf{K}_{y, \Phi}(t, f) = \mathbf{K}_{x, \Phi}(t - \tau, f)$ is obtained iff

$$\begin{aligned} \mathbf{K}_{y, \Phi}(t, f) &= \langle \mathbf{WV}_y, \Phi[t, f] \rangle = \langle \mathbf{WV}_x, \Phi[t, f](\cdot + \tau, \cdot) \rangle = \langle \mathbf{WV}_x, \Phi[t - \tau, f] \rangle \\ &\iff \Phi[t, f](\cdot + \tau, \cdot) = \Phi[t - \tau, f] \end{aligned}$$

as a result the above condition demonstrates that filters of different times $\Phi[t - \tau, f], \forall \tau$ are just time translations of each other giving the desired result.

- Frequency modulation/shift with frequency ω_0 to have $y(t) = x(t) e^{i\omega_0 t}$ leads

$$\begin{aligned} \mathbf{WV}_y(t, \omega) &= \int_{-\infty}^{\infty} y(t + \frac{\tau}{2}) y^*(t - \frac{\tau}{2}) e^{-i\tau\omega} d\tau \\ &= \int_{-\infty}^{\infty} x(t + \frac{\tau}{2}) e^{i\omega_0(t + \frac{\tau}{2})} x^*(t - \frac{\tau}{2}) e^{-i\omega_0(t - \frac{\tau}{2})} e^{-i\tau\omega} d\tau \\ &= \int_{-\infty}^{\infty} x(t + \frac{\tau}{2}) e^{i\omega_0\tau} x^*(t - \frac{\tau}{2}) e^{-i\tau\omega} d\tau = \mathbf{WV}_x(t - \tau, \omega) \end{aligned}$$

$$\begin{aligned}
&= \int_{-\infty}^{\infty} x(t + \frac{\tau}{2})x^*(t - \frac{\tau}{2})e^{-i\tau(\omega + \omega_0)}d\tau = W_x(t - \tau, \omega) \\
&= \mathbf{WV}_x(t, \omega + \omega_0)
\end{aligned}$$

now leveraging the same result that for the time equivariance, we obtain the desired result.

- For completeness, we also demonstrate here how the Wigner-Ville behaves under rescaling of the input $y(t) = x(t/a)$

$$\begin{aligned}
\mathbf{WV}_y(t, \omega) &= \int_{-\infty}^{\infty} y(t + \frac{\tau}{2})y^*(t - \frac{\tau}{2})e^{-i\tau\omega}d\tau \\
&= \int_{-\infty}^{\infty} x((t + \frac{\tau}{2})/a)x^*((t - \frac{\tau}{2})/a)e^{-i\tau\omega}d\tau \\
&= \int_{-\infty}^{\infty} x(t/a + \frac{\tau}{2a})x^*(t/a - \frac{\tau}{2a})e^{-i\tau\omega}d\tau \\
&= \int_{-\infty}^{\infty} x(t/a + \frac{\tau}{2a})x^*(t/a - \frac{\tau}{2a})e^{-i\tau\omega}d\tau \\
&= a \int_{-\infty}^{\infty} x(t/a + \frac{\tau}{2})x^*(t/a - \frac{\tau}{2})e^{-i\tau a\omega}d\tau \\
&= a\mathbf{WV}_x(t/a, a\omega)
\end{aligned}$$

□

C.5 PROOF OF PROP 3

Proof. The proof is obtained by using the fact that the maximum of the 2-dimensional Gaussian with given covariance matrix will always be smaller than $\frac{1}{\det(\Sigma)}$ and is always nonnegative, and then simply upper bounding the norm difference by this value times the representation. Now that we have the upper bound in term of $\|WV_x - WV_{D(X)}\|_{L^2(\mathbb{R}^2)}$ simply apply the Lipschitz constant inequality with κ the constant of the WVD which exists as long as x is bounded which is the case as we have $L^2(\mathbb{R})$ leading to the desired result. □

C.6 PROOF OF LEMMA 2

Proof. We will obtain the desired result by unrolling the following equations and see that it coincides with the one of the Theorem with the given kernel applied on the WV representation:

$$\begin{aligned}
&\int_{-\infty}^{\infty} \mathbf{ST}_x(t, \omega + \frac{\eta}{2})\mathbf{ST}_x^*(t, \omega - \frac{\eta}{2})\mathcal{F}_g(\eta)e^{j2\pi\eta t}d\eta \\
&= \int_{-\infty}^{\infty} \mathbf{ST}_x(t, \omega + \frac{\eta}{2})\mathbf{ST}_x^*(t, \omega - \frac{\eta}{2}) \int_{-\infty}^{\infty} g(\xi)e^{-i\xi\eta}d\xi e^{j2\pi\eta t}d\eta \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} w^*(t - \tau)x(\tau)w(t - \theta)x^*(\theta)e^{-i\omega(\tau - \theta)} \int_{-\infty}^{\infty} g(\xi) \int_{-\infty}^{\infty} e^{j2\pi\eta(t - \frac{\tau}{2} - \frac{\theta}{2} - \xi)}d\eta d\xi d\tau d\theta \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} w^*(t - \tau)x(\tau)w(t - \theta)x^*(\theta)e^{-i\omega(\tau - \theta)}g(\xi)\delta(t - \frac{\tau}{2} - \frac{\theta}{2} - \xi)d\xi d\tau d\theta \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} w^*(t - \tau)x(\tau)w(\tau + 2\xi - t)x^*(2t - \tau - 2\xi)g(\xi)e^{-i\omega(2\tau - 2t + 2\xi)}d\xi d\tau \\
&= \int_{-\infty}^{\infty} g(\xi) \int_{-\infty}^{\infty} w^*(\xi - \mu/2)x(t + \mu/2 - \xi)w(\xi + \mu/2)x^*(t - \mu/2 - \xi)e^{-i\omega(2\mu)}d\tau d\xi \\
&\hspace{20em} (\mu = 2(\tau - t + \xi)) \\
&= \int g(\xi)(W_w(\xi, \cdot) \star W_x(t - \xi, \cdot))(\omega)/2d\xi \implies \Pi'(\tau, \omega) = g(\tau)W_w(\tau, \omega)/2
\end{aligned}$$

$$\text{with } \Pi'(t, \omega) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{t^2}{\sigma_t^2/(\sigma_w^2\sigma_t^2+1)} - \frac{\omega^2\sigma_t^2}{2}}.$$

□

C.7 PROOF OF THM. 2

Proof. From the above (lemma) result, it then follows directly that performing convolutions with two Gaussians can be rewritten as a single Gaussian convolution with the given parameters based on both Gaussians. \square

D INTERFERENCES

Finally, we now consider the study of interference that can arise in the K-transform whenever multiple events occur in the signal x Sen (2014). In fact, the Gabor limit also known as the Heisenberg Uncertainty principle Heisenberg (1927), corresponds to the optimal limit one can achieve in time and frequency resolution without introducing interference in the representation. In our case, no constraints are imposed onto $\Phi[t, f]$, and while learning will allow to reach any desired kernel for the task at hand, we propose some conditions that would prevent the presence of interference. In fact, there is a general necessary condition ensuring that the representation does not contain any interference.

Proposition 5. *A sufficient condition to ensure absence of interference in the K-transform is to have a nonnegative representation as in $K_{x, \Phi}(t, f) \geq 0, \forall t, f$.*

Proof. We only consider signals that can be expressed as a linear combination of oscillatory atoms (natural signals) Mallat (1989). Due to the form of the Wigner-Ville distribution (with the quadratic term) the interference are oscillatory terms and thus can not be positive only. In short, there can not be a positive coefficient appearing due to an interference without its negative counterpart. As a result, a nonnegative representation can not have any interference. See (Cohen (1995)) for more details and background on Wigner-Ville distribution and interference. \square

Lastly, thanks to the Gaussian form, we can directly obtain the shape and in particular area of the logon (the inverse of the joint time and frequency resolution) as follows.

Proposition 6. *A sufficient condition to ensure absence of interference in the K-transform is to have the effective area of the 2D Gaussian greater than $\frac{1}{4\pi}$, that is, $\epsilon'_T \epsilon'_F \geq \frac{1}{4\pi}, \forall t, f$ with*

$$\sigma'_T = \sigma_T \cos^2(\theta) + 2\rho \cos(\theta) \sin(\theta) + \sigma_F \sin^2(\theta) \quad (9)$$

$$\sigma'_F = \sigma_T \sin^2(\theta) - 2\rho \cos(\theta) \sin(\theta) + \sigma_F \cos^2(\theta), \quad (10)$$

$$\theta = \frac{\arctan\left(\frac{2\rho}{\sigma_T - \sigma_F}\right)}{2} \quad (11)$$

Proof. The proof applies the Uncertainty principle which provides the minimal area of the Logon to ensure absence of interference (Gabor (1946)), this is then combined with the above result on the area of the Logon in the case of our 2-dimension Gaussian to obtain the desired result. So first, the condition is that $\sigma_T \sigma_F \geq \frac{1}{4\pi}$, however in our case we also have the possible chirpness parameter. But we can obtain the new rotated covariance matrix s.t. the chirpness is removed by rotating the time-frequency place instead. This can be done easily (see for example http://scipp.ucsc.edu/~haber/ph116A/diag2x2_11.pdf) to obtain the new parameter as

$$\sigma_T = \sigma_T \cos^2(\theta) + 2\rho \cos(\theta) \sin(\theta) + \sigma_F \sin^2(\theta)$$

$$\sigma_F = \sigma_T \sin^2(\theta) - 2\rho \cos(\theta) \sin(\theta) + \sigma_F \cos^2(\theta),$$

$$\theta = \frac{\arctan\left(\frac{2\rho}{\sigma_T - \sigma_F}\right)}{2}$$

now the constraint can be expressed as

$$(\sigma_T \cos^2(\theta) + 2\rho \cos(\theta) \sin(\theta) + \sigma_F \sin^2(\theta)) (\sigma_T \sin^2(\theta) - 2\rho \cos(\theta) \sin(\theta) + \sigma_F \cos^2(\theta)) \geq \frac{1}{4\pi}$$

\square

Table 3: Various hand picked K-transform parameters leading to known time invariant TFRs with their respective parameters. For the adaptive versions, such as the wavelet tree with various Gabor mother wavelet parameters σ_0 , then at each time t , the corresponding parameter is any of the optimal one based on the desired criterion.

	$\mu_{\text{time}}(t, f)$	$\mu_{\text{freq}}(t, f)$	$\sigma_{\text{time}}(t, f)$	$\sigma_{\text{freq}}(t, f)$	$\rho(t, f)$ (chirpness)
Spectrogram	t	f	σ_t	σ_t^{-1}	0
Melscale Spectrogram	t	$2^{S(1-f/\pi)}$	σ_t	$2^{S(f/\pi-1)}\sigma_t^{-1}$	0
Scalogram	t	$2^{S(1-f/\pi)}$	$2^{S(1-f/\pi)}\sigma_0$	$2^{S(f/\pi-1)}\sigma_0^{-1}$	0
Scattering Layer	t	$2^{S(1-f/\pi)}$	$s2^{S(1-f/\pi)}\sigma_0$	$2^{S(f/\pi-1)}\sigma_0^{-1}$	0
Chirpogram	t	$2^{S(1-f/\pi)}$	$2^{S(1-f/\pi)}\sigma_0$	$2^{S(f/\pi-1)}\sigma_0^{-1}$	$\rho(t, f) \neq 0$

E GAUSSIAN TRUNCATION

Notice however that the STFT has to be done by padding the signal windows to allow interpolation, this is common when dealing with such transformations. Given some parameters, we convert them into discrete bins to get actual window sizes as follows

$$N_\sigma(\epsilon) = -g_\sigma^{-1}(\epsilon) \times 2Fs$$

with g_σ^{-1} the inverse of the gaussian density distribution with 0 mean and σ standard deviation, taken only on the negative part of its support. As such, N is a function that maps a given tolerance and standard deviation to the window length in bins s.t. the apodization window at the boundaries of this window are of no more than ϵ .

F EXAMPLE OF K-TRANSFORM PARAMETERIZATION

We propose in Table 3 different configurations of the parameter θ corresponding to some standard and known TFRs.

G COMPUTATIONAL COMPLEXITY

The computational complexity of the method varies with the covariance matrix of the kernels, it will range from quadratic when the kernels tend toward a Delta function (as the representation computation bottleneck becomes the computation of the exact WV) to linear for large covariance matrices. In such case, the computation is done with spectrograms of small windows (and with almost no spectral frequency correlation) for the smoothed pseudo WV. In addition to those cases, the use of translation (in time and/or frequency) also greatly reduce computational costs as it allows for the use of convolutions into the K-transform computation. This allows to put all the computation bottleneck into the computation of the smoothed pseudo WV.

H DATA SET DOCC10 : DYNI ODONTOCETE CLICK CLASSIFICATION, 10 SPECIES, ENS DATA CHALLENGE

Audio-MNIST. This dataset consists of multiple (60) speakers of various characteristics enunciation digits from 0 to 9 inclusive [Becker et al. \(2018\)](#). The classification task consists of classifying the spoken digit, 30,000 recordings are given in the dataset. Each recording is from a controlled environment without external noise except breathing and standard recording device artifacts. **BirdVox-70k.** This dataset consists of avian recording obtained during the fall migration season of 2015. The recording were obtained in North American and is made of a balanced binary classification of predicting the presence or not of a bird in a given audio clip [Lostanlen et al. \(2018\)](#). In total 70,000 clips are contained in the dataset. The outside recordings present various types of non stationary noises and sources. **Google Speech Commands.** This dataset consists of 65,000 one-second long utterances of 30 short words such as “yes”, “no”, “right” or “left” [Warden \(2017\)](#). The recordings are obtained from thousands of different people who contributed through the AIY website ¹. The task is thus to classify the spoken command. **FreeSound DCASE.** This task addresses the problem of general-purpose automatic audio tagging [Fonseca et al. \(2018\)](#). The dataset provided for this

¹https://aiyprojects.withgoogle.com/open_speech_recording

task is a reduced subset from FreeSound and is a large-scale, general-purpose audio dataset annotated with labels from the AudioSet Ontology. There are 41 classes and $\approx 95,000$ training samples, classes are diverse such as Bus, Computer keyboard, Flute, Laughter or Bark. The main challenge of the task is the noise present in the training set labels reflecting the expensiveness of having high quality annotations.

DOCC10. The goal of this dataset is to classify recordings obtained from various subsea acoustic stations or autonomous surface vehicles (ASV). The recordings consist of clicks and the task is to perform classification based on the corresponding emitting species among 10 marine mammals such as Grampus griseus- Risso’s dolphin, Globicephala macrorhynchus This task is published by ENS Paris College de France: DOCC10², for additional details please see Appendix ???. This challenge is linked to the largest international bioacoustical challenge from Scripps Institute and LIS CNRS. The goal is to classify transients, clicks emitted as biosonar by 10 cetaceans species (odontocetes). The recordings of the clicks are from a post-processed subset of DCLDE 2018 challenge (9 species, <http://sabiiod.org/DCLDE>), plus recordings of Cachalot (Physeter macrocephalus, Pm) done from ASV Sphyrna (<http://sphyrna-odyssey.com>) near Toulon, France. Each input signal is 8192 bins, at a sampling rate of 200 kHz. They each includes a click centered in the middle of the window in the case of the test set. The clicks in the training set are in various positions and background noises.

Challenge goals The goal is to classify each click according to the corresponding emitting species. The 10 species are : (0) Gg: Grampus griseus- Risso’s dolphin (1) Gma: Globicephala macrorhynchus- Short-finned pilot whale (2) La: Lagenorhynchus acutus- Atlantic white-sided dolphin (3) Mb: Mesoplodon bidens- Sowerby’s beaked whale (4) Me: Mesoplodon europaeus- Gervais beaked whale (5) Pm: Physeter macrocephalus - Sperm whale (6) Ssp: Stenella sp. Stenellid dolphin (7) UDA: Delphinid type A - a group of dolphins (species not yet determined) (8) UDB: Delphinid type B - another group of dolphins (species not yet determined) (9) Zc: Ziphius cavirostris- Cuvier’s beaked whale The metric is the MAP.

Data description The recordings are from various subsea acoustic stations or autonomous surface vehicles (ASV). A part (9 classes) has been provided by the Scripps Institution of Oceanography for the 8th DCLDE Workshop. They consist of acoustic recordings from multiple deployments of high-frequency acoustic recording packages (Wiggins and Hildebrand, 2007) deployed in the Western North Atlantic (US EEZ) and Gulf of Mexico. It has a 100 kHz of bandwidth (200 kHz sample rate). Data were selected to cover multiple seasons and locations while providing high species diversity and call counts as described in <http://sabiiod.org/DCLDE/challenge.html#highFreqData>. The initial data set is 3 To and weak labeled. DYNIS LIS TOULON (Ferrari Glotin 2019) filtered the labels by several detector and clustering, yielding to 90 000 samples (6 Go). Another part (1 class) is from Dyni team recording in Mediterranean sea, from ASV Sphyrna, at 384 kHz Fe, downsampled at 200 kHz Fe. In sum, it yields around 11312 samples per class in the train set, and 2096 samples per class in the test set.

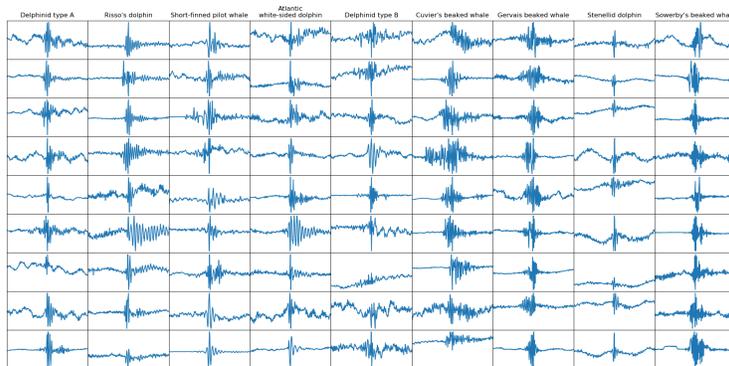


Figure 6: DOCC10 dataset samples, one class per column. The click is centered. Details in ENS DATA CHALLENGE WEB SITE.

²<https://challengedata.ens.fr/challenges/32>