

EXPLORING THE OPTIMALITY OF TIGHT-FRAME SCATTERING NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

The wavelet scattering transform creates geometric invariants and deformation stability. In multiple signal domains, it has been shown to yield more discriminative representations compared to other non-learned representations, and to outperform learned representations in certain tasks, particularly on limited labeled data and highly structured signals. The wavelet filters used in the scattering transform are typically selected to create a tight frame via a parameterized mother wavelet. In this work, we investigate if such a tight frame construction is optimal. Focusing on Morlet wavelets, we propose to learn the scales, orientations, and aspect ratios of the filters to produce problem-specific parameterizations of the scattering transform. We show that our learned versions of the scattering transform yield significant performance gains in small-sample classification settings over the standard scattering transform. Moreover, our empirical results suggest that tight-frames may not always be necessary for scattering transforms to extract effective representations.

1 INTRODUCTION

The scattering transform, proposed in Mallat (2012), is a cascade of wavelets and complex modulus nonlinearities, which can be seen as a convolutional neural network (CNN) with fixed, predetermined filters. This construction can be used to build representations with geometric invariants and is shown to be stable to deformations. It has been demonstrated to yield impressive results on problems involving highly structured signals (Bruna & Mallat, 2013; Oyallon et al., 2013; Andén & Mallat, 2014; Sifre & Mallat, 2014; Hirn et al., 2015; 2017; Eickenberg et al., 2018; Andén et al., 2019; Sinz et al., 2020; Perlmutter et al., 2020), outperforming a number of other classic signal processing techniques. Since scattering transforms are instantiations of CNNs, they have been studied as mathematical models for understanding the impressive success of CNNs in image classification (Bruna & Mallat, 2013; Mallat, 2016). As discussed in Bruna & Mallat (2013), first-order scattering coefficients are similar to SIFT descriptors (Lowe, 2004), and higher-order scattering can provide insight into the information added with depth (Mallat, 2016). Moreover, theoretical and empirical study of information encoded in scattering networks indicates that they often promote linear separability, which in turn leads to effective representations for downstream classification tasks (Bruna & Mallat, 2013; Oyallon et al., 2017; Andén et al., 2015; Eickenberg et al., 2018).

Scattering-based models have been shown to be useful in several applications involving scarcely annotated or limited labeled data (Bruna & Mallat, 2013; Sifre & Mallat, 2013; Oyallon et al., 2018; Eickenberg et al., 2018). Indeed, most breakthroughs in deep learning in general, and CNNs in particular, involve significant effort in collecting massive amounts of well-annotated data to be used when training deep overparameterized networks. While big data is becoming increasingly prevalent, there are numerous applications where the task of annotating more than a small number of samples is infeasible, giving rise to increasing interest in small-sample learning tasks and deep-learning approaches towards them (Brigato et al., 2021; Barz & Denzler, 2020; Bendre et al., 2020; Bruintjes et al., 2021). Recent work has shown that, in image classification, state-of-the-art results can be achieved by hybrid networks that harness the scattering transform as their early layers followed by learned layers based on a wide residual network architecture (Oyallon et al., 2018). Here, we further advance this research avenue by proposing to use the scattering paradigm not only as fixed preprocessing layers in a concatenated architecture, but also as a parametric prior to learn filters in




Figure 1: **Initialized wavelet filters pre and post-training.** Real part of Morlet wavelet filters initialized with *tight-frame* (left) and *random* (right) schemes before (top) and after (bottom) training. The filters were optimized on the entire CIFAR-10 training set with linear model. For the tight-frame filters, we observe substantial changes in both scale and aspect ratio. On the other hand, all random filters undergo major changes in orientation and scale.

a CNN. This allows us to also shed light on whether the tight-frame filterbank construction (Mallat, 1999) is an optimal approach for building filterbanks from a mother-wavelet for discriminative tasks.

Recall that the scattering construction is based on complex wavelets, generated from a mother wavelet via dilations and rotations, aimed to cover the frequency plane while having the capacity to encode informative variability in input signals (Bruna & Mallat, 2013). Further, discrete parameterization and indexing of these operations (i.e., by dilation scaling or rotation angle) have traditionally been carefully constructed to ensure the resulting filter bank forms an efficient tight frame Mallat (1999; 2012) with well-established energy preservation properties. On the other hand, it has been observed that the first layers of convolutional networks resemble wavelets but may not necessarily form a tight frame (Krizhevsky et al., 2012). The question then arises: is it necessary to use a tight frame when constructing wavelet filterbanks? Here, we relax the standard tight frame construction by considering another alternative where a small number of wavelet parameters used to create the wavelet filterbanks are optimized for the task at hand.

To our knowledge, this is the first work that aims to learn the wavelet filters of scattering networks in 2D signals. Related work and the empirical protocol are summarized in Sec. 3. and Sec. 4 respectively. In Sec. 4.1, we explore the different filter construction schemes by comparing the wavelet filter parameterizations they produce when optimized over different datasets. In Sec. 4.2, we evaluate the robustness of our parametric scattering networks to deformation. In Sec. 4.3, we demonstrate the advantages of our approach in limited labeled data settings and study the adaptation of the wavelet parameters toward a supervised task. Finally, in Sec. 4.4, we investigate the adaptation of the parametrized scattering using an unsupervised objective. Further technical details appear in appendices, and code accompanying the work is available on https://github.com/psn-iclr-submission/iclr_anon.

2 RELATED WORK

Learning useful representations from little training data (Bendre et al., 2020) is arduous and a reality in a variety of domains such as in biomedicine and healthcare. Recent works have tried to tackle this problem. Lezama et al. (2018) replace the categorical cross-entropy loss with a geometric loss called Orthogonal Low-rank Embedding (OLÉ) to reduce the intra-class variance and enforce inter-class margins. Barz & Denzler (2020) also propose to replace the categorical cross-entropy loss, but this time with the cosine loss function in order to decrease overfitting in the small-sample classification settings. The cosine loss function, as opposed to the softmax function used with cross-entropy, does not push the logits of the true class to infinity as explained in Szegedy et al. (2016). Other methods show promise by incorporating prior knowledge into the model. Oyallon et al. (2018) introduce hybrid networks where the scattering transform with fixed wavelets was shown to be an effective replacement for early layers of learned convolutional networks on a wide residual network architecture. Cotter & Kingsbury (2019) also propose a hybrid network called a learnable ScatterNet, where learning layers are intermixed between the scattering orders, unlike our work where only a few parameters governing the wavelet construction are modified. Ulicny et al.

Table 1: Canonical Parameters of Morlet wavelet

Param	Role	Param	Role
σ	Gaussian window scale	θ	Global orientation
ξ	Frequency scale	γ	Aspect Ratio

(2019) propose Harmonic Networks (HN), a hybrid network consisting of fixed Discrete Cosine Transform filters combined with learnable weights in CNNs.

Related to our work, adding learnable components to existing wavelet-based representations has been considered in a number of recent works in the context of time-series (Balestriero et al., 2018; Seydoux et al., 2020; Cosentino & Aazhang, 2020; Balestriero et al., 2020). Balestriero et al. (2018); Seydoux et al. (2020) learn a spline parametrized mother wavelet for 1D problems. Similarly, Cosentino & Aazhang (2020) parametrized the group transform in the context of time-series data. Our work, alternatively, focuses on 2D problems and maintains the canonical Morlet wavelet parameterization, but allows deviation from a tight-frame filter bank.

3 PARAMETRIZATION OF SCATTERING NETWORKS

We first revisit the formulation of traditional scattering convolution networks in Sec. 3.1 and introduce our parametric scattering transform in Sec. 3.2 and 3.3. Finally, Sec. 3.4 discusses scattering parameter initialization.

3.1 SCATTERING NETWORKS

For simplicity, we focus here on 2D scattering networks up to their 2nd order. Subsequent orders can be computed by following the same iterative scheme, but have been shown to yield negligible energy (Bruna & Mallat, 2013). Given a signal $x(u)$, where u is the spatial position index, we compute the scattering coefficients S^0x, S^1x, S^2x , of order 0, 1, and 2 respectively. For an integer J , corresponding to the spatial scale of the scattering transform, and assuming an $N \times N$ signal input with one channel, the resulting feature maps are of size $\frac{N}{2^J} \times \frac{N}{2^J}$, with channel sizes varying with the scattering coefficient order (i.e., 1 channel at order 0, JL channels at order 1 and $L^2J(J-1)/2$ channels at order 2).

To calculate 0th-order coefficients, we consider a low pass filter ϕ_J with a spatial window of scale 2^J , such as a Gaussian smoothing function. We then convolve this filter with the signal and down-sample by a factor of 2^J to obtain $S^0x(u) = x * \phi_J(2^J u)$. Due to the low-pass filtering, high-frequency information is discarded here and is recovered in higher-order coefficients via wavelets introduced as in a filter bank.

Morlet wavelets are a typical example of filters used in conjunction with the scattering transform, and are defined as

$$\psi_{\sigma, \theta, \xi, \gamma}(u) = e^{-\|D_\gamma R_\theta(u)\|^2 / (2\sigma^2)} (e^{i\xi u'} - \beta), \quad (1)$$

where β is a normalization constant to ensure wavelets integrate to 0 over the spatial domain, $u' = u_1 \cos \theta + u_2 \sin \theta$, R_θ is the rotation matrix of angle θ and $D_\gamma = \begin{pmatrix} 1 & 0 \\ 0 & \gamma \end{pmatrix}$. The four parameters can be adjusted and are presented in Table 1. From one wavelet $\psi_{\sigma', \theta', \xi', \gamma'}(u)$, a tight frame is obtained by dilating it by factors 2^j , $0 \leq j < J$, and rotating by L angles θ equally spaced over the circle, to get $\{2^{2j} \psi_{\sigma', \theta', \xi', \gamma'}(2^{-j} R_\theta(u))\}$, which is then completed with the lowpass ϕ_J . This can be written in terms of the parameters in Table 1 as $\psi_{2^j \sigma', \theta' - \theta, 2^{-j} \xi', \gamma'}(u) = \psi(2^{-j} R_\theta(u))$. By slight abuse of notations, we use ψ_λ here, $\lambda = (\sigma_j, \theta, \xi_j, \gamma_j)$, to denote such wavelets indexed by θ and j . The resulting set of filters is visualized in the frequency domain in Figure 2.

First-order scattering coefficients are calculated by first convolving the input signal with one of the generated complex wavelets (i.e., indexed by the parameters in Table 1) and downsampling the resulting filtered signal by the scale factor 2^{j_1} of the wavelet chosen. Then, a pointwise complex modulus is used to add nonlinearity, and the resulting real signal is smoothed via a low-pass filter. Finally, another downsampling step is applied, this time by a factor of 2^{J-j_1} , to obtain an optimally compressed output size. Mathematically, we have

$$S^1x(\lambda_1, u) = |x * \psi_{\lambda_1}| * \phi_J(2^J u).$$

The resulting feature map has $J \cdot L$ channels, based on the number of wavelets in the generated family.

Second-order coefficients are generated similarly, with the addition of another cascade of wavelet transform and modulus operator before the low-pass smoothing, i.e.,

$$S^2x(\lambda_1, \lambda_2, u) = ||x * \psi_{\lambda_1}| * \psi_{\lambda_2}| * \phi_J(2^J u).$$

Due to the interaction between the bandwidths and frequency supports of first and second order, only coefficients with $j_1 < j_2$ have significant energy. Hence, the second-order output yields a feature map with $\frac{1}{2}J(J-1)L^2$ channels.

3.2 MORLET CANONICAL WAVELET PARAMETERIZATION

While the wavelet filters are traditionally fixed to approximate a tight frame, we let the network learn the optimal parameters of each wavelet. In other words, we constrain our filters to always be Morlet wavelets by only optimizing the parameters in Table 1. We call this approach the Morlet canonical parameterization of the wavelet. We adapted the Kymatio software package (Andreux et al., 2020) to create the learnable scattering network. To provide such data-driven optimization of scattering parameters, we show, in Appendix A, that it is possible to backpropagate through this construction. We can now learn the parameters jointly with other parameters in an end-to-end differentiable architecture.

3.3 MORLET EQUIVARIANT PARAMETERIZATION

In the Morlet canonical parameterization approach, the canonical parameters of each filter are learned. As an alternative method, we consider the Morlet equivariant parameterization in which the number of learnable parameters is reduced by a factor L compared to the Morlet canonical parameterization. Each filter per scale is constructed using the same four parameters in Table 1: σ , ξ , γ and Θ . However, the global orientation of the L filters for each scale are set to be $[\Theta, \Theta + \frac{\pi}{L}, \Theta + \frac{2\pi}{L}, \dots, \Theta + \frac{(L-1)\pi}{L}]$. By construction, the tight-frame filters are equivariant.

3.4 INITIALIZATION

To evaluate the importance of the tight-frame construction, we consider two initializations and study their impact on resulting performance in both learned and nonlearned settings. First, a tight-frame initialization follows common implementations of the scattering transform by setting $\sigma_{j,\ell} = 0.8 \times 2^j$, $\xi_{j,\ell} = \frac{3\pi}{4} 2^{-j}$, and $\gamma_{j,\ell} = \frac{4}{L}$ for $j = 1, \dots, J$, $\ell = 1, \dots, L$, while for each j , we set $\theta_{j,\ell}$ to be equally spaced on $[0, \pi)$. Second, as an alternative, we consider a random initialization where these parameters are sampled as $\sigma_{j,\ell} \sim \log(U[\exp 1, \exp 5])$, $\xi_{j,\ell} \sim U[0.5, 1]$, $\gamma_{j,\ell} \sim U[0.5, 1.5]$, and $\theta_{j,\ell} \sim U[0, 2\pi]$. That is, orientations are selected uniformly at random on the circle, the filter width σ is selected using an exponential distribution across available scales and the spatial frequency ξ is chosen to be in the interval $[0.5, 1]$, which lies in the center of the feasible range between aliasing ($> \pi$) and the fundamental frequency of the signal size ($2\pi/N$ where N is the number of pixels). Finally, we select the *aspect ratio* variable to vary around the spherical setting of 1.0, with a bias towards stronger orientation selectivity (0.5) compared to lesser orientation selectivity (1.5).

4 EXPERIMENTS

Our empirical evaluations are based on three image datasets: CIFAR-10, COVIDx CRX-2, and KTH-TIPS2. CIFAR-10 and KTH-TIPS2 are natural image and texture recognition datasets, correspondingly. They are often used as general-purpose benchmarks in similar image analysis settings (Azuri & Weinshall, 2021; Sifre & Mallat, 2013). COVIDx CRX-2 is a dataset of X-ray scans for COVID-19 diagnosis; its use here demonstrates the viability of our parametric scattering approach in practice, e.g., in medical imaging applications.

We evaluate the use of the parametrized scattering with two common models. In the first case, we consider the scattering as feeding into a simple linear model (denoted LL). The LL configurations are used to evaluate the linear separability of the obtained scattering representations and have the added benefit of providing a more interpretable model. In the second case, we take the approach

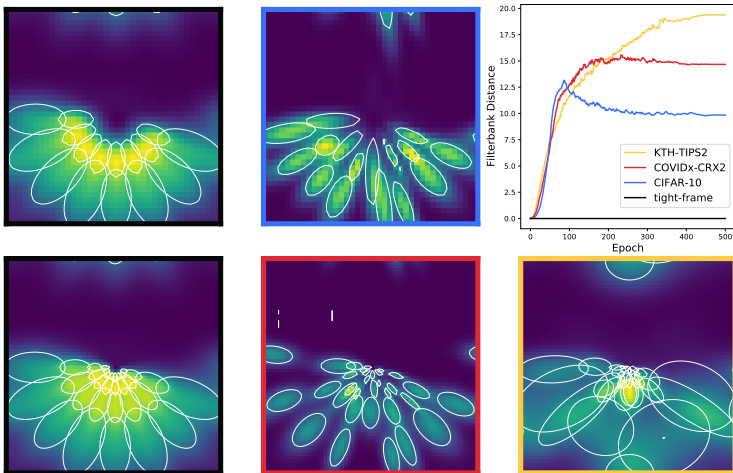


Figure 2: **Parametric scattering network learns dataset specific filters.** The graph (top right) shows the *filterbank distance* over epochs as the filters are trained on different datasets. We visualize dataset specific parameterizations of scattering filterbanks (border colors from the legend) in Fourier space. Scattering filters optimized for natural (CIFAR-10) and medical image (COVIDx CRX2) become more orientation-selective, i.e., thinner in the Fourier domain. On the other hand, filters optimized for texture discrimination (KTH-TIPS2) become less orientation-selective and deviate most from a tight-frame setup.

of Oyallon et al. (2018) and consider the scattering as the first stage of a deeper CNN, specifically a Wide Residual Network (WRN) (Zagoruyko & Komodakis, 2016). The architecture of the WRN hybrid is described in more detail in Appendix C.

For both models (LL and WRN), we compare learned parametric scattering networks (LS) to fixed ones (S). For learned scattering (LS), we consider two scattering parameterization approaches: *Morlet canonical*, described in Sec. 3.2 and *Morlet equivariant*, described in Sec. 3.3. To show the importance of the parametric approach, we also ablate the naive parameterization where all pixels of the wavelets are adapted, which we refer to as a pixel-wise parameterization. For each scattering architecture, we consider both random and tight-frame (TF) initialization. The fixed scattering models determined by the TF construction are equivalent to traditional scattering transforms. Finally, we also compare our approach to a fully learned WRN (with no scattering priors) and ResNet-50 (He et al., 2016) applied directly to input data.

Across all scattering configurations, a batch-normalization layer with learnable affine parameters is added after all scattering layers. Classification is performed via a softmax layer yielding the final output. All models are trained using cross-entropy loss, minimized by stochastic gradient descent with momentum of 0.9. Weight decay is applied to the linear model and to the WRN. The learning rate is scheduled according to the one cycle policy (Smith & Topin, 2019). The scheduler’s div factor is always set to 25. Implementation details specific to each dataset are described in Appendix B. We replicate some of the experiments with learnable scattering networks followed by WRN on CIFAR-10, COVIDx-CRX2, and KTH-TIPS2 using the cosine loss function (Barz & Denzler, 2020). The results are reported in Appendix G.

4.1 EXPLORING DATASET-SPECIFIC PARAMETERIZATIONS

We first compare dataset-specific Morlet wavelet parameterizations and evaluate their similarities to a tight frame. Specifically, we train our parametric scattering networks using the canonical morlet wavelet formulation with a linear classification layer and quantitatively and qualitatively compare the similarities of the learned filter bank to the tight-frame initialization. To facilitate quantitative comparison, we use a distance metric for comparing the sets of Morlet wavelet filters and Morlet wavelet filterbanks (i.e., scattering network instantiations), allowing us to measure deviations from the tight-frame initialization.

We evaluate distances between two individual Morlet wavelets as $\Upsilon(M_1, M_2) = \left\| (\sigma_1, \xi_1, \gamma_1)^T - (\sigma_2, \xi_2, \gamma_2)^T \right\|_2 + \text{arcdist}(\theta_1, \theta_2)$ where $M_i = (\sigma_i, \xi_i, \gamma_i, \theta_i)^T$ denotes the

parameterization of the Morlet wavelet. We use the arc distance on the unit circle to compare values of theta. Since the set of learned scattering filters does not have a canonical order, to compare a learned scattering network to the tight frame scattering network, we use a matching algorithm to match one set of filters to another. Specifically, we first compute Υ between all combinations of filter pairs from both networks, then use a minimum cost bipartite matching algorithm (Kuhn, 1955) to find the minimal distance match between the two sets of filters. The final distance we use as a notion of similarity between two scattering networks is the sum of Υ for all matched pairs in the bipartite graph. Henceforth, we will refer to this distance as the *filterbank distance*.

The graph in Figure 2 leverages the *filterbank distance* to show the evolution of scattering networks initialized from a tight frame and trained on different datasets. Each network is trained on 1188 samples of its respective dataset (the standard size for KTH-TIPS2). All filters deviate quickly from a tight frame, but KTH-TIPS2’s keep changing the longest and ultimately deviate the most. We also observe that, filters initialized with the random initialization of Sec. 3 become more similar to our tight-frame initialization during the course of training (see Figure 9 in Appendix D.4).

On the left-hand side of Figure 2, we visualize the dataset-specific scattering network parameterizations in Fourier space. White contours are drawn around each Morlet wavelet for clarity. The top black border corresponds to tight-frame initialization at $J=2$, shown for comparison to CIFAR-10 in blue (also $J=2$). The bottom black border corresponds to tight-frame initialization at $J=4$, shown for comparison to COVIDX-CRX2 red and KTH-TIPS2 yellow (both $J=4$).

The filters optimized on the KTH-TIPS2 texture dataset (yellow) become less orientation-selective (wider in Fourier space) than the tight-frame initialization, with filters at $J=0$ becoming the least orientation-selective of the whole filter bank. We note that the filters at spatial scales $J= 2$ and 3 seem to change the most from a tight frame as illustrated in the appendix (Fig. 7). In contrast, the filters optimized on COVIDx-CRX2 become more orientation-selective in general i.e., thinner in Fourier space, while changing the most at spatial scale $J=0$ as shown in the appendix (Fig. 5). The filters optimized on CIFAR-10 mirror those optimized on COVIDx-CRX2, also becoming more orientation-selective than their tight-frame counterparts. We suspect that this is due to a reliance on edges for object classification datasets, which seem to require more orientation-selective filters. On the other hand, the morlet wavelets optimized for texture classification seem to discard some edge information in favor of less orientation-specific filters. Each dataset-specific parameterization seems to discard unneeded information from the tight-frame initialization in favor of accentuating problem-specific attributes. In Sec. 4.3, we demonstrate these learned filters are not only interpretable but improve task performance, suggesting the tight frame is not optimal for many problems of interest. Nonetheless, a tight-frame does constitute a good starting point for learning. Indeed, the dataset-specific parameterizations for COVIDX-CRX2 and KTH-TIPS2 are, visually, very different, yet they move similar filterbank distances from the tight-frame initialization (see fig.2), which are small relative to the distances observed for randomly initialized and trained models (see Appendix D.4).

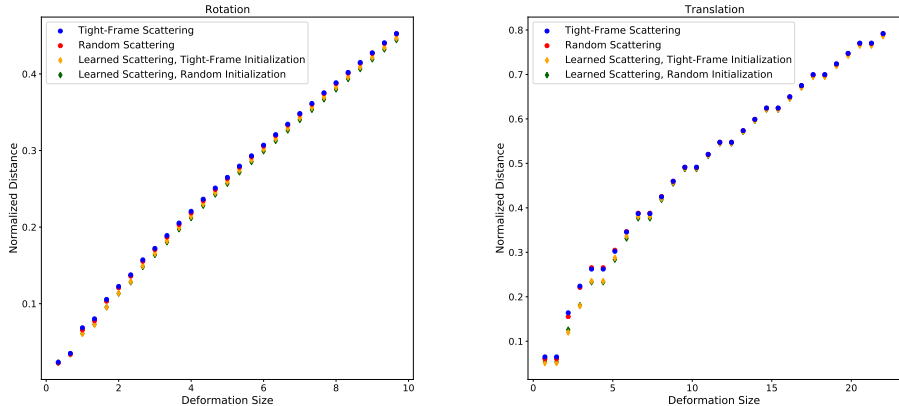


Figure 3: **Normalized distances between scattering representations of an image and its deformation.** Our parametric scattering transform shares similar stability to deformations as the scattering transform.

4.2 ROBUSTNESS TO DEFORMATION

In Mallat (2012), it is shown that the scattering transform is stable to small deformations of the form $x(u - \tau(u))$ where $x(u)$ is a signal and τ a diffeomorphism. Given the substantial changes to the filter composition in the learning process, we ask now whether these seem to significantly deviate from the stability result obtained from the carefully handcrafted construction proposed in Mallat (2012), and extensively used in previous work (e.g., Bruna & Mallat, 2013; Eickenberg et al., 2018). To evaluate the robustness of our parametric scattering network to different geometric distortions, we apply several tractable deformations to a chest X-ray image x with varying deformation strength. The transformed image is denoted by \tilde{x} . For each of the different deformation strengths, we plot the Euclidean distance between the scattering feature constructed from the original image $S(x)$ and the scattering feature constructed from the transformed image $S(\tilde{x})$. We then normalize the obtained distance by $S(x)$ to measure the relative deviation in scattering coefficients (handcrafted or learned). Figure 3 demonstrates representative results for a small rotation and translation on images from the COVIDx datasets, while additional deformations are shown in Appendix H. We observe that the substantial change in the filter construction retains the scattering robustness properties for these simple deformations, thus indicating that the use of learned filters (instead of designed ones) does not necessarily detract from the stability of the resulting transform.

4.3 SMALL DATA REGIME

We evaluate the parametric scattering network in limited labeled data settings. Following the evaluation protocol from Oyallon et al. (2018), we subsample each dataset at various sample sizes to showcase the performance of scattering-based architectures in the small data regime. In our experiments, we train on a small random subset of the training data but always test on the entire test set as done in Oyallon et al. (2018). To obtain comparable and reproducible results, we control for deterministic GPU behavior and assure that each model is initialized the same way for the same seed. Furthermore, we use the same set of seeds for models evaluated on the same number of samples. For instance, the TF learnable hybrid with a linear model would be evaluated on the same ten seeds as the fixed tight-frame hybrid with a linear model when trained on 100 samples of CIFAR-10. Some fluctuation is inevitable when subsampling datasets. Hence all our figures include averages and standard error calculated over different seeds.

CIFAR-10 consists of 60,000 images from ten classes. The train set contains 50,000 class-balanced samples, while the test set contains the remaining images. Table 2 reports the evaluation of our learnable scattering approach on CIFAR-10 with training sample sizes of 100, 500, 1K, and 50K. The training set is augmented with horizontal flipping, random cropping, and pre-specified autoaugmentation (Cubuk et al., 2018) for CIFAR-10. We used autoaugmentation (Cubuk et al., 2018) to showcase the best possible small-sample results and ablate this component in Appendix F. We use a spatial scale of $J = 2$ in the scattering transforms.

As shown in Table 2, the scattering networks with wavelets optimized pixel-wise perform the worst in the small-data regime. It shows that with limited labeled samples, there is not enough data and too many learnable parameters to learn effectively the pixels of the wavelets. Adding more constraints (i.e., constraining the wavelets to be Morlet) is beneficial in this setting. We also observe that the Morlet canonical parameterization yields a similar performance to the Morlet equivariant parameterization (i.e., most standard errors overlap). Thus, adding even more constraints, by reducing the number of learnable parameters in the parametric scattering transform, does not degrade the performance in the small-data regime.

We observe that randomly initialized learnable with canonical parameterization only achieves similar performance to TF learnable canonical when trained on the whole dataset. These results suggest the TF initialization, derived from rigorous signal processing principles, is empirically beneficial as a starting point in the very few sample regime but can be improved upon by learning.

Among the linear models, our TF-initialized learnable scattering networks (i.e., Morlet canonical and equivariant) significantly outperform all others in few sample settings. This demonstrates that learnable scattering networks obtain a more linearly separable representation than their fixed counterparts, perhaps by building greater dataset-specific intra-class invariance. Figure 1 shows the real part of the canonical wavelet filters before and after optimization on the entire training set. In Appendix E, we visualize canonical equivariant wavelet filters.

Table 2: CIFAR-10 mean accuracy and std. error over 10 seeds, with $J = 2$ and multiple training sample sizes. Learnable scattering with TF initialization improves performance for all architectures, while randomly initialized scattering requires more training data to reach similar performance.

Arch.	Init.	Parametrization	100 samples	500 samples	1000 samples	All
LS+LL†	TF	Canonical	37.84 ± 0.57	52.68 ± 0.31	57.43 ± 0.17	69.57
LS+LL†	TF	Equivariant	39.69 ± 0.56	51.98 ± 0.25	57.01 ± 0.16	66.65
LS+LL	TF	Pixel-Wise	32.30 ± 0.69	47.14 ± 0.91	51.87 ± 0.34	64.53
S+LL	TF	-	36.01 ± 0.55	48.12 ± 0.25	53.25 ± 0.24	65.58
LS+LL†	Rand	Canonical	34.81 ± 0.60	49.6 ± 0.39	55.72 ± 0.39	69.39
LS+LL†	Rand	Equivariant	34.67 ± 0.73	46.59 ± 0.60	52.95 ± 0.36	65.64
LS+LL	Rand	Pixel-Wise	29.44 ± 0.41	42.14 ± 0.27	47.44 ± 0.43	62.72
S+LL	Rand	-	29.77 ± 0.47	41.85 ± 0.41	46.3 ± 0.37	57.72
LS+WRN†	TF	Canonical	43.60 ± 0.87	63.13 ± 0.29	70.14 ± 0.26	93.61
LS+WRN†	TF	Equivariant	39.86 ± 1.59	62.85 ± 0.32	69.52 ± 0.23	92.57
LS+WRN	TF	Pixel-Wise	39.20 ± 0.80	54.14 ± 0.68	57.59 ± 0.48	92.97
S+WRN	TF	-	43.16 ± 0.78	61.66 ± 0.32	68.16 ± 0.27	92.27
LS+WRN†	Rand	Canonical	41.42 ± 0.65	59.84 ± 0.40	67.40 ± 0.28	93.36
LS+WRN†	Rand	Equivariant	40.84 ± 1.02	60.81 ± 0.40	68.62 ± 0.31	92.53
LS+WRN	Rand	Pixel-Wise	31.49 ± 0.63	45.85 ± 0.43	50.72 ± 0.28	91.86
S+WRN	Rand	-	32.08 ± 0.46	46.84 ± 0.21	52.76 ± 0.33	85.35
WRN-16-8	-	-	38.78 ± 0.72	62.97 ± 0.41	71.37 ± 0.31	96.84
ResNet-50	-	-	33.17 ± 0.92	52.13 ± 0.74	64.42 ± 0.40	91.23

params : 156k for S+LL; 155k for LS+LL; 22.6M for S+WRN; 22.6M for LS+WRN; 22.3M for WRN; and 22.5M for ResNet
†: ours; TF: Tight-Frame; LS: Learnable Scattering; S: Scattering; Rand: Random

Among the WRN hybrid models, the TF-initialized canonical learnable scattering performs best. Canonical TF learnable still improves over TF fixed when paired with a WRN, indicating some loss of information in the fixed scattering representation is mitigated by data-driven tuning or optimization. Finally, our approach outperforms the fully trained ResNet-50 and outperforms the WRN-16-8 on 100 and 500 training samples, demonstrating the effectiveness of the scattering prior in the small data regime. However, the WRN-16-8 outperforms our model on 1,000 samples and 50,000 samples.

COVIDx CRX-2 is a two-class (positive and negative) dataset of chest X-Ray images of COVID-19 patients (Wang et al., 2020). The train set contains 15,951 unbalanced images, while the test set contains 200 positive and 200 negative images. The spatial scale of the scattering transform is set to $J = 4$. Table 3 reports our evaluation on sample sizes of 100, 500, and 1K images. We use the same protocol as for CIFAR-10. Morlet canonical parameterization yields similar performance to the Morlet equivariant parameterization (i.e., most standard errors overlap), as also observed with CIFAR-10.

When the scattering networks are postpended with a linear layer, TF-initialized learnable (i.e., Morlet canonical and equivariant) performs better than TF fixed, showing the viability of our approach on real-world data. We observe that randomly initialized learnable yields lower performance than TF learnable on 100 and 500 samples. On 1K, it achieves similar performance, demonstrating that random initialization can achieved comparable performance to TF with enough data. WRN-16-8 performs worse than TF-initialized learnable followed with a linear layer. When combined with a CNN, TF-initialized learnable also performs better than TF fixed and outperforms WRN-16-8 and ResNet-50.

KTH-TIPS2 contains 4,752 images from 11 material classes. The images captured the material at scales. Each class is divided into four *samples* (108 images each) of different scales. Using the standard protocol, we train the model on one *sample* ($11 * 108$ images), while the rest are used for testing (Song et al., 2017). In total, each training set contains 1,188 images. Table 3 reports the classification accuracies.

With TF initialization and a linear layer, we observe that the performance is similar for the different architectures. The performance of randomly initialized learnable is also similar to TF. The fixed and randomly initialized model perform the worst, showing that even poorly initialized filters can effectively be optimized. Altogether, these results further corroborate our previous findings, notably that TF initialization acts as a good prior for scattering networks. Out of all the WRN hybrid models, the TF-initialized learnable model using canonical parameterization achieves the highest average accuracy. We note that while WRN increases the performance compared to the linear layer, it also significantly increases the total number of parameters, therefore exhibiting a tradeoff between per-

Table 3: COVIDx CRX-2 and KTH-TIPS2 mean accuracy & std. error with $J = 4$ over 10 seeds and 16 seeds respectively. (COVIDx CRX-2) TF-initialized learnable scattering network performs better than models that do not incorporate scattering priors. (KTH-TIPS2) Similarly, the WRN-16-8 and ResNet-50 perform extremely poorly relative to hybrid models trained on KTH-TIPS2.

Arch.	Init.	Parameterization	C-100 samples	C-500 samples	C-1000 samples	KTH-1188 samples
LS+LL†	TF	Canonical	82.30 ± 1.78	88.50 ± 0.71	89.90 ± 0.40	66.09 ± 1.05
LS+LL†	TF	Equivariant	83.06 ± 1.53	87.56 ± 0.94	89.15 ± 0.60	66.41 ± 1.24
S+LL	TF	-	81.08 ± 1.88	87.20 ± 0.77	89.23 ± 0.69	66.17 ± 1.10
LS+LL†	Rand	Canonical	76.85 ± 1.50	86.45 ± 0.95	89.70 ± 0.65	65.79 ± 0.85
LS+LL†	Rand	Equivariant	76.73 ± 1.57	85.64 ± 1.38	87.98 ± 0.55	65.31 ± 1.42
S+LL	Rand	-	76.08 ± 1.56	84.13 ± 0.91	86.80 ± 0.41	61.37 ± 0.82
LS+WRN†	TF	Canonical	81.20 ± 1.73	90.50 ± 0.70	93.68 ± 0.35	69.23 ± 0.67
LS+WRN†	TF	Equivariant	81.86 ± 2.07	91.56 ± 0.52	93.97 ± 0.34	68.55 ± 0.80
S+WRN	TF	-	80.85 ± 1.85	89.05 ± 0.59	91.90 ± 0.54	68.84 ± 0.71
LS+WRN†	Rand	Canonical	80.95 ± 1.54	88.08 ± 0.70	91.65 ± 0.55	68.30 ± 0.47
LS+WRN†	Rand	Equivariant	80.12 ± 1.76	87.44 ± 1.17	91.40 ± 0.67	67.50 ± 0.72
S+WRN	Rand	-	80.63 ± 1.73	86.68 ± 0.59	90.60 ± 0.50	66.29 ± 0.36
WRN-16-8	-	-	80.50 ± 1.15	85.95 ± 2.04	88.82 ± 1.64	51.24 ± 1.37
ResNet-50	-	-	74.04 ± 1.35	86.45 ± 0.51	90.86 ± 0.57	44.95 ± 0.65

C: COVIDx CRX-2 # params : 493K for LS/S+LL; 23.05M for LS/S+WRN; 22.3M for WRN; 23.5M for ResNet
 KTH: KTH-TIPS2 # params : 883K for LS/S+LL; 23.7M for LS/S+WRN; 22.3M for WRN; 23.5M for ResNet
 †: ours; TF: Tight-Frame; LS: Learnable Scattering; S: Scattering; Rand: Random

Table 4: Scattering and Learned Unsupervised Scattering features evaluated by training a linear classifier on CIFAR-10. We observe the unsupervised learned scattering improves the representation.

Method	100 samples	500 samples	1000 samples	All
Scattering (Fixed)	36.01 ± 0.55	48.12 ± 0.25	53.25 ± 0.24	65.58 ± 0.04
Unsupervised Learnt Scattering	38.05 ± 0.45	52.92 ± 0.28	57.76 ± 0.25	68.47 ± 0.04

formance and model complexity. The WRN-16-8 and ResNet-50 perform extremely poorly relative to hybrid models, showing the effectiveness of the scattering priors for texture discrimination.

4.4 UNSUPERVISED LEARNING OF SCATTERING PARAMETERS

We have studied the adaptation of the wavelet parameters towards a supervised task. We now perform a preliminary investigation to determine if the scattering representation can be improved in a purely unsupervised manner. We consider the recently popularized SimCLR framework (Chen et al., 2020), which encourages representations from two data augmentations of the same input to lie close together. We learn scattering network parameters on CIFAR-10 using this unsupervised objective function and subsequently evaluate the discriminativeness of the features under a standard linear evaluation experiment on the full CIFAR-10 dataset and in the small data regimes comparing them to the standard scattering transform. The results are shown in Table 4. We observe the unsupervised learning of filter parameters can improve the scattering representation under standard unsupervised learning evaluation protocols.

5 CONCLUSION

This work demonstrated the competitive results of adapting a small number of Morlet wavelet filter parameters in the scattering network. We illustrated that filters learned by parametric scattering can be interpreted in relation to the specific task (e.g., becoming thinner in object recognition tasks that require sensitivity to edges). We also empirically demonstrate that our parametric scattering transform shares similar stability to deformations as the traditional scattering transforms. Overall we find that the parametric scattering network provides state-of-the-art results for classification in the low-data regime when combined with a linear layer and as well in a hybrid CNN. These results go towards bridging the gap between the handcrafted filter design in scattering transforms, which provides tractable properties and supports low-parameter models, and the fully (unparametrized) learned ones commonly used in CNN work, especially in computer vision and generally on 2D structured data. In particular, our results can lead to future work investigating the impact of down-sampling on the representations learned by the parametric scattering network, as well as application to uncertainty estimation by leveraging the low parameter CNN in a Bayesian framework.

6 REPRODUCIBILITY STATEMENT

The results presented in the paper are reproducible. We provide an anonymized direct link to the source code. The content of this link contains complete instructions on how to run the code to reproduce the results presented for the paper. In addition, all details for the models and training procedures are described in the main paper under Sec. 4 and the Appendix B.

REFERENCES

- Joakim Andén and Stéphane Mallat. Deep scattering spectrum. *IEEE Transactions on Signal Processing*, 62(16):4114–4128, 2014.
- Joakim Andén, Vincent Lostanlen, and Stéphane Mallat. Joint time-frequency scattering for audio classification. In *2015 IEEE 25th International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6, 2015.
- Joakim Andén, Vincent Lostanlen, and Stéphane Mallat. Joint time–frequency scattering. *IEEE Transactions on Signal Processing*, 67(14):3704–3718, 2019.
- Mathieu Andreux, Tomás Angles, Georgios Exarchakis, Roberto Leonarduzzi, Gaspar Rochette, Louis Thiry, John Zarka, Stéphane Mallat, Joakim Andén, Eugene Belilovsky, et al. Kymatio: Scattering transforms in python. *J. Mach. Learn. Res.*, 21(60):1–6, 2020.
- Idan Azuri and Daphna Weinshall. Generative latent implicit conditional optimization when learning from small sample. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 8584–8591, 2021.
- Randall Balestriero, Romain Cosentino, Hervé Glotin, and Richard Baraniuk. Spline filters for end-to-end deep learning. In *International conference on machine learning*, pp. 364–373. PMLR, 2018.
- Randall Balestriero, Hervé Glotin, and Richard G Baraniuk. Interpretable super-resolution via a learned time-series representation. arXiv:2006.07713, 2020.
- Bjorn Barz and Joachim Denzler. Deep learning on small datasets without pre-training using cosine loss. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1371–1380, 2020.
- Nihar Bendre, Hugo Terashima Marín, and Peyman Najafirad. Learning from few samples: A survey. arXiv:2007.15484, 2020.
- Lorenzo Brigato, Björn Barz, Luca Iocchi, and Joachim Denzler. Tune it or don’t use it: Benchmarking data-efficient image classification. In *2nd Visual Inductive Priors for Data-Efficient Deep Learning Workshop*, 2021.
- Robert-Jan Brintjes, Attila Lengyel, Marcos Baptista Rios, Osman Semih Kayhan, and Jan van Gemert. Vipriors 1: Visual inductive priors for data-efficient deep learning challenges. arXiv:2103.03768, 2021.
- Joan Bruna and Stéphane Mallat. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. *CoRR*, abs/2002.05709, 2020.
- Romain Cosentino and Behnaam Aazhang. Learnable group transform for time-series. In *International Conference on Machine Learning*, pp. 2164–2173. PMLR, 2020.
- Fergal Cotter and Nick Kingsbury. A learnable scatternet: Locally invariant convolutional layers. In *2019 IEEE International Conference on Image Processing (ICIP)*, pp. 350–354. IEEE, 2019.
- Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. arXiv:1805.09501, 2018.

- Michael Eickenberg, Georgios Exarchakis, Matthew Hirn, Stéphane Mallat, and Louis Thiry. Solid harmonic wavelet scattering for predictions of molecule properties. *The Journal of chemical physics*, 148(24):241732, 2018.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Matthew Hirn, Nicolas Poilvert, and Stéphane Mallat. Quantum energy regression using scattering transforms. arXiv:1502.02077, 2015.
- Matthew Hirn, Stéphane Mallat, and Nicolas Poilvert. Wavelet scattering regression of quantum chemical energies. *Multiscale Modeling & Simulation*, 15(2):827–863, 2017.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- José Lezama, Qiang Qiu, Pablo Musé, and Guillermo Sapiro. Ole: Orthogonal low-rank embedding—a plug and play geometric loss for deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8109–8118, 2018.
- David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999.
- Stéphane Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012.
- Stéphane Mallat. Understanding deep convolutional networks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150203, 2016.
- Edouard Oyallon, Stéphane Mallat, and Laurent Sifre. Generic deep networks with wavelet scattering. arXiv:1312.5940, 2013.
- Edouard Oyallon, Eugene Belilovsky, and Sergey Zagoruyko. Scaling the scattering transform: Deep hybrid networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 5618–5627, 2017.
- Edouard Oyallon, Sergey Zagoruyko, Gabriel Huang, Nikos Komodakis, Simon Lacoste-Julien, Matthew Blaschko, and Eugene Belilovsky. Scattering networks for hybrid representation learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2208–2221, 2018.
- Michael Perlmutter, Feng Gao, Guy Wolf, and Matthew Hirn. Geometric wavelet scattering networks on compact riemannian manifolds. In *Mathematical and Scientific Machine Learning*, pp. 570–604. PMLR, 2020.
- Léonard Seydoux, Randall Balestriero, Piero Poli, Maarten De Hoop, Michel Campillo, and Richard Baraniuk. Clustering earthquake signals and background noises in continuous seismic data with unsupervised deep learning. *Nature communications*, 11(1):1–12, 2020.
- Laurent Sifre and Stéphane Mallat. Rotation, scaling and deformation invariant scattering for texture discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1233–1240, 2013.
- Laurent Sifre and Stéphane Mallat. Rigid-motion scattering for texture classification. arXiv:1403.1687, 2014.
- Paul Sinz, Michael W Swift, Xavier Brumwell, Jialin Liu, Kwang Jin Kim, Yue Qi, and Matthew Hirn. Wavelet scattering networks for atomistic systems with extrapolation of material properties. *The Journal of Chemical Physics*, 153(8):084109, 2020.

- Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, pp. 1100612, 2019.
- Yang Song, Fan Zhang, Qing Li, Heng Huang, Lauren J O’Donnell, and Weidong Cai. Locally-transferred fisher vectors for texture classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4912–4920, 2017.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Matej Ulicny, Vladimir A. Krylov, and Rozenn Dahyot. Harmonic networks with limited training samples. In *2019 27th European Signal Processing Conference (EUSIPCO)*, pp. 1–5, 2019. doi: 10.23919/EUSIPCO.2019.8902831.
- Linda Wang, Zhong Qiu Lin, and Alexander Wong. Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific Reports*, 10(1):1–12, 2020.
- Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016.

A BACKPROPAGATION THROUGH THE PARAMETRIC SCATTERING NETWORK

We show here that it is possible to backpropagate through this construction. Namely, we verify the differentiability of this construction by explicitly computing the partial derivatives with respect to these parameters. First, the \mathbb{R} -linear derivative of the complex modulus $f(z) = |z|$ is $f'(z) = \frac{z}{|z|}$. Next, we show the differentiation of convolution with wavelets with respect to their parameters. For simplicity, we focus here on differentiation of the Gabor portion¹ of the filter construction from Eq. 1, written as:

$$\begin{aligned} \varphi(u) = & \exp\left(-\frac{1}{2\sigma^2}(u_1^2(\cos^2(\theta) + \sin^2(\theta)\gamma^2) + u_2^2(\cos^2(\theta)\gamma^2 + \sin^2(\theta)))\right) \\ & + 2\cos(\theta)\sin(\theta)u_1u_2(1 - \gamma^2) + i\xi(\cos(\theta)u_1 + \sin(\theta)u_2). \end{aligned}$$

Its derivatives with respect to the parameters are

$$\begin{aligned} \frac{\partial \varphi}{\partial \theta}(u) &= \frac{1}{\sigma^2}(u_2 \cos \theta - u_1 \sin \theta)(i\xi\sigma^2 + u_1(\gamma^2 - 1)\cos \theta + u_2(\gamma^2 - 1)\sin \theta)\varphi(u); \\ \frac{\partial \varphi}{\partial \sigma}(u) &= \frac{1}{\sigma^3}(u_1^2(\cos^2 \theta + \gamma^2 \sin^2 \theta) + u_2^2(\gamma^2 \cos^2 \theta + \sin^2 \theta) + 2u_1u_2 \cos \theta \sin \theta(1 - \gamma^2))\varphi(u); \\ \frac{\partial \varphi}{\partial \xi}(u) &= i(u_1 \cos \theta + u_2 \sin \theta)\varphi(u); \text{ and} \\ \frac{\partial \varphi}{\partial \gamma}(u) &= -\frac{1}{\sigma^2}(u_1^2\gamma \sin^2 \theta + u_2^2\gamma \cos^2 \theta - 2u_1u_2\gamma \cos \theta \sin \theta)\varphi(u). \end{aligned}$$

Finally, the derivative of the convolution with such filters is given by $\frac{\partial}{\partial \zeta}(f * \varphi)(t) = \int f(t - u) \frac{\partial \varphi}{\partial \zeta}(u) du$ where ζ is any of the filter parameter from Table 1. It is easy to verify that these derivations can be chained together to propagate through the scattering cascades defined in Sec. 3.1.

B IMPLEMENTATION DETAILS

This section describes the implementation details for each dataset.

B.1 CIFAR-10

CIFAR-10 consists of consists of 60,000 images of size $32 \times 32 \times 3$ from ten classes. The linear models were trained using a max learning rate of 0.06 for all parameters on 5K, 1K, 500, and 500 epochs for 100, 500, 1K, 50K samples, respectively. The hybrid WRN models were trained using a max learning rate of 0.1 on 3K, 2K, 1K, and 200 epochs for 100, 500, 1K, and 50K samples respectively. We use batch gradient descent except when the models are trained with 50K samples where we use mini-batch gradient descent of size 1024. On the entire training set, we also train the models on 10 seeds and, in all cases, the standard errors are always inferior to 0.3. All scattering networks use a spatial scale $J = 2$.

B.2 COVIDx CRX-2

COVIDx CRX-2 is a two-class (positive and negative) dataset of $1024 \times 1024 \times 1$ chest X-Ray images of COVID-19 patients (Wang et al., 2020). In our experiments, we always train on a class-balanced subset of the training set. We resize the images to 260×260 and train our network with random crops of 224×224 pixels. The only data augmentation we use is random horizontal flipping. All models were trained on 400 epochs using a max learning rate of 0.01. All hybrid models are trained with a mini-batch size of 128. All scattering networks use a spatial scale $J = 4$.

B.3 KTH-TIPS2

We resize the images to 200×200 and train our network with random crops of 128×128 pixels. The training data is augmented with random horizontal flips and random rotations. All scattering

¹It is not difficult to extend this derivation to Morlet wavelets, but the resulting expressions are rather cumbersome and left out for brevity.

networks use a spatial scale of 4. We set the maximum learning rate of the scattering parameters to 0.1 while it is set to 0.001 for all other parameters. All hybrid models are trained with a mini-batch size of 128. The hybrid linear models are trained for 250 epochs, while the hybrid WRN models are trained for 150 epochs. We evaluate each model, training it with four different seeds on each *sample* of material, amounting to 16 total runs.

C WIDE RESIDUAL NETWORK HYBRID ARCHITECTURE

In the experiments of Sec. 4, the scattering networks are combined with a WRN hybrid described in Oyallon et al. (2017). We follow a similar architecture to the WRN hybrid used in Oyallon et al. (2018). The description of the architecture used for the experiments is given in Table 5. We use the same architecture for all the datasets. The architecture consists of a scattering network that greatly reduces the spatial resolution of the input. The CIFAR-10 scattering stage yields output with 8×8 spatial resolution (Scattering with $J=2$). Similarly, the KTH-TIPS2 data and COVIDx-CRX2 data give outputs with 16×16 and 8×8 spatial resolutions respectively.

Stage	Description
scattering	Learned or Tight Frame
conv1	3×3 , CONV LAYER $128 \rightarrow 256$
conv2	$\begin{bmatrix} 3 \times 3, \text{ CONV LAYER } 256 \\ 3 \times 3, \text{ CONV LAYER } 256 \end{bmatrix} \times 4$
conv3	$\begin{bmatrix} 3 \times 3, \text{ CONV LAYER } 256 \\ 3 \times 3, \text{ CONV LAYER } 256 \end{bmatrix} \times 4$
avg-pool	Avg pooling to a size 1×1

Table 5: Description of WRN hybrid architecture used for the experiments. Each conv layer represents a 2D Convolution followed by a batch normalization and a ReLU non-linearity.

D DATASET SPECIFIC PARAMETERIZATIONS

The following figures show the configuration obtained by our minimal bipartite matching algorithm. Each wavelet titles share a common naming scheme. The first letter of the title is either F (fixed filters) or O (optimized filters). The next character is always a number and corresponds to the ID of the match. For all Oixxxxxx titles, there will be a corresponding Fixxxxxxx title. These filters are matched to each other. The next character is always D (distance). It is superseded by a numerical value, the Morlet wavelet distance between the filter and its match. The next character is σ , followed by a number corresponding to the magnitude of the distances between the σ parameters of both filters. The next character is γ , followed by a number corresponding to the magnitude of the distances between the γ parameters of both filters. The next character is ξ , followed by a number corresponding to the magnitude of the distances between the ξ parameters of both filters.

D.1 COVIDX-CRX2




Figure 4: Filters trained on 1188 samples of COVIDx-CRX2 for 500 epochs, the first, third, fifth, and seventh rows correspond to filters optimized from a tight-frame, while the second, fourth, sixth, and eighth rows correspond to tight-frame initialized filters. The filters are displayed in pairs correspond to the 'closest' (by our distance metric defined above) filters of both types. For instance, the first filter of row one matches the first filter of row 2.




Figure 5: Filters trained on 1188 samples of COVIDx-CRX2 for 500 epochs, the first, third, fifth, and seventh rows correspond to filters optimized from a tight-frame, while the second, fourth, sixth, and eighth rows correspond to tight-frame initialized filters. The filters are displayed in pairs correspond to the 'closest' (by our distance metric defined above) filters of both types. For instance, the first filter of row one matches the first filter of row 2. The filters are displayed in increasing order of their distances. The top left corner corresponds to the filters that changed the least from their initialization, while the filters in the bottom right corner changed the most.

D.2 KTH-TIPS2




Figure 6: Filters trained on 1188 samples of KTH-TIPS2 for 500 epochs, the first, third, fifth, and seventh rows correspond to filters optimized from a tight-frame, while the second, fourth, sixth, and eighth rows correspond to tight-frame initialized filters. The filters are displayed in pairs correspond to the 'closest' (by our distance metric defined above) filters of both types. For instance, the first filter of row one matches the first filter of row 2.




Figure 7: Filters trained on 1188 samples of KTH-TIPS2 for 500 epochs, the first, third, fifth, and seventh rows correspond to filters optimized from a tight-frame, while the second, fourth, sixth, and eighth rows correspond to tight-frame initialized filters. The filters are displayed in pairs correspond to the 'closest' (by our distance metric defined above) filters of both types. For instance, the first filter of row one matches the first filter of row 2. The filters are displayed in increasing order of their distances. The top left corner corresponds to the filters that changed the least from their initialization, while the filters in the bottom right corner changed the most.

D.3 CIFAR-10




Figure 8: Filters trained on 1190 samples of CIFAR-10 for 500 epochs, the first and third, rows correspond to filters optimized from a tight-frame, while the second and fourth rows correspond to tight-frame initialized filters. The filters are displayed in pairs correspond to the 'closest' (by our distance metric defined above) filters of both types. For instance, the first filter of row one matches the first filter of row 2.

D.4 DATASET SPECIFIC INITIALIZATIONS WITH A RANDOM INITIALIZATION

In Figure 9, we show how the filters adapt when initialization begins from a random setting. We note the deviation to tight frame is much greater than in the case where we initialize in a tight frame. However, as per our filterbank distance, we observe the filters do move closer to the tight frame than their initialization.

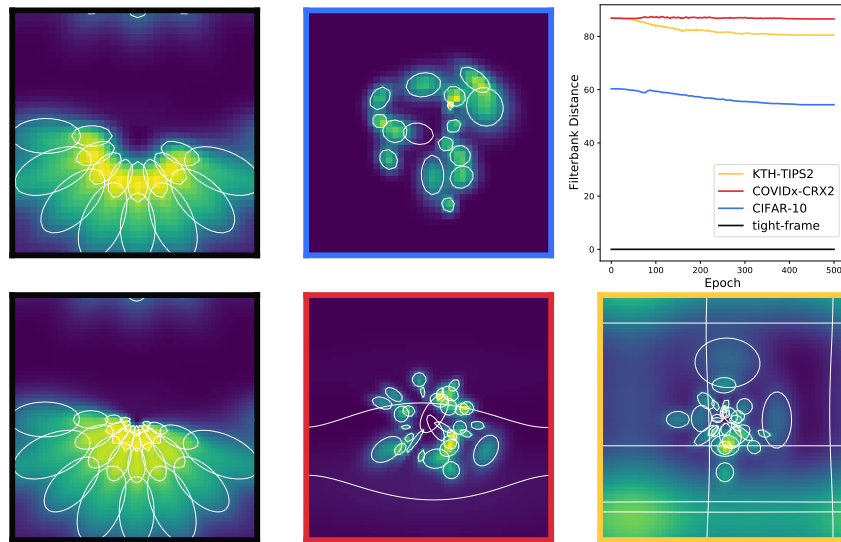


Figure 9: The graph shows the *filterbank distance* over epochs as the filters, initialized from random init, are trained on different datasets. To the left, we visualize dataset specific parameterizations of scattering filterbanks in Fourier space. The graph on the right shows that the randomly initialized filterbanks become more similar to a tight frame during training.

E EQUIVARIANT FILTERS

We observe that, in some cases, using equivariant filters yields better accuracy, as shown in Table 2. Figure 10 illustrates equivariant filters initialized using tight frame construction before and after optimization on 500 training samples of CIFAR-10, where the scattering network is combined with a linear layer. The spatial scale is set to $J = 2$. Similarly, Figure 11 illustrates equivariant filters initialized randomly before and after optimization. In the two figures, each row corresponds to a different scale. Since J is set to 2, we have two rows. We observe that the equivariant filters in each scale/row are the same, except for the global orientation.




Figure 10: Example of equivariant filters initialized using tight frame construction (Top) Real part of wavelet filters before optimization (Bottom) Real part of wavelet filters after optimization.




Figure 11: Example of equivariant filters initialized randomly (Top) Real part of wavelet filters before optimization (Bottom) Real part of wavelet filters after optimization.

F NO AUTOAUGMENT ON CIFAR-10

The training set of CIFAR-10 is augmented with pre-specified autoaugment in Table 2 to demonstrate the best possible results. To understand the effect of autoaugment, we replicate the same experiments except for not augmenting the training set with autoaugment. Table 6 reports the performance of the different architectures on CIFAR-10. We observe that the scattering networks followed by WRN underperform when no autoaugment is used. The difference in performance between using autoaugment and not using it is smaller when the scattering network is followed with a linear layer. Surprisingly, the performance of the scattering networks followed with a linear layer trained on all data increased without autoaugment. It seems that in the case of a scattering network followed by a linear model, autoaugment is not as useful as with a deep model on top and can also be harmful in some cases.

G COSINE LOSS

We replicate the experiments of Sec. 4 with learnable scattering networks followed by WRN on CIFAR-10, COVIDx-CRX2, and KTH-TIPS2 using the same parameters, except for using the cosine

Table 6: CIFAR-10 mean accuracy and std. error over 10 seeds with $J = 2$ and multiple training sample sizes. The table compares the effect of augmenting the training set with pre-specified autoaugment. When the scattering network is followed by a WRN, then using autoaugment is necessary to obtain better performance.

Init.	Arch.	AA	100 samples	500 samples	1000 samples	All
TF	LS+LL†	Yes	37.84 ± 0.57	52.68 ± 0.31	57.43 ± 0.17	69.57 ± 0.1
TF	LS+LL†	No	39.70 ± 0.62	50.74 ± 0.30	54.76 ± 0.22	74.94 ± 0.06
TF	S+LL	Yes	36.01 ± 0.55	48.12 ± 0.25	53.25 ± 0.24	65.58 ± 0.04
TF	S+LL	No	37.55 ± 0.62	49.67 ± 0.33	53.96 ± 0.48	70.71 ± 0.03
Rand	LS+LL†	Yes	34.81 ± 0.60	49.6 ± 0.39	55.72 ± 0.39	69.39 ± 0.41
Rand	LS+LL†	No	32.64 ± 0.38	42.88 ± 0.23	47.40 ± 0.32	74.71 ± 0.08
Rand	S+LL	Yes	29.77 ± 0.47	41.85 ± 0.41	46.3 ± 0.37	57.72 ± 0.1
Rand	S+LL	No	31.71 ± 0.34	40.57 ± 0.32	44.42 ± 0.51	61.79 ± 0.31
TF	LS+WRN†	Yes	43.60 ± 0.87	63.13 ± 0.29	70.14 ± 0.26	93.61 ± 0.12
TF	LS+WRN†	No	34.95 ± 0.96	54.21 ± 0.39	62.17 ± 0.28	90.17 ± 0.34
TF	S+WRN	Yes	43.16 ± 0.78	61.66 ± 0.32	68.16 ± 0.27	92.27 ± 0.05
TF	S+WRN	No	35.15 ± 0.43	52.77 ± 0.35	60.72 ± 0.21	89.05 ± 0.38
Rand	LS+WRN†	Yes	41.42 ± 0.65	59.84 ± 0.40	67.4 ± 0.28	93.36 ± 0.19
Rand	LS+WRN†	No	31.08 ± 1.00	48.37 ± 0.76	55.41 ± 0.49	88.80 ± 0.47
Rand	S+WRN	Yes	32.08 ± 0.46	46.84 ± 0.21	52.76 ± 0.33	85.35 ± 1.06
Rand	S+WRN	No	27.73 ± 0.43	41.05 ± 0.32	47.19 ± 0.37	79.67 ± 0.59

†: ours TF: tight-frame LS: Learnable Scattering AA: Autoaugment

params : 156k for S+LL; 155k for LS+LL; 11M for S+WRN; 22.6M LS+WRN; and 22.3M for WRN only

* From Oyallon et al. (2018)

loss (Barz & Denzler, 2020) function. The wavelet filters were initialized using the tight-frame construction. Table 7 demonstrates the average accuracy on the three datasets. For CIFAR-10, models trained with cosine loss underperform networks trained with cross-entropy loss. We do not observe the same behavior when the models are trained on KTH-TIPS2 where the performance increases slightly by using the cosine loss function. Thus, cosine loss can improve performance on small data regimes for certain datasets.

Table 7: CIFAR-10, COVIDx-CRX2 and KTH-TIPS2 mean accuracy and std. error using cosine loss function.

Init.	Arch.	Dataset	Loss	100 samples	500 samples	1000 samples	1188 samples
TF	LS+WRN	CIFAR-10	CE	43.6 ± 0.87	63.13 ± 0.29	70.14 ± 0.26	-
TF	LS+WRN	CIFAR-10	Cosine	42.94 ± 0.77	61.42 ± 0.26	68.29 ± 0.18	-
TF	LS+WRN	COVIDx	CE	81.20 ± 1.73	90.50 ± 0.70	93.68 ± 0.35	-
TF	LS+WRN	COVIDx	Cosine	80.03 ± 2.16	89.53 ± 0.89	92.75 ± 0.65	-
TF	LS+WRN	KTH-TIPS2	CE	-	-	-	69.23 ± 0.67
TF	LS+WRN	KTH-TIPS2	Cosine	-	-	-	70.86 ± 0.67

TF: tight-frame LS: Learnable Scattering S: Scattering CE: Cross-Entropy Loss

H ROBUSTNESS TO DEFORMATIONS

Models used were pre-trained and evaluated using the COVIDx CRX-2 setup mentioned in Section 4.2 with a linear classifier head. Deformations used are rotation, shear, translation, and several diffeomorphisms (denoted Custom 1 and Custom 2), and strengths for the deformations ranges from 0 to the maximum value for the deformation given in Table 8. Additional results are depicted in Figure 12.

Table 8: Deformations and their maximum value

Deformation	Maximum Value
Rotation	10
Shear	5
Translation	22
Custom1	1
Custom2	1

Custom 1, $\tau_\epsilon^1(u)$, and Custom 2, $\tau_\epsilon^2(u)$, are defined as such:

$$\tau_\epsilon^1(u) = \epsilon \begin{bmatrix} 0.3u_1^2 + 0.2u_2^2 \\ 0.2(0.2u_1) \end{bmatrix}, \tau_\epsilon^2(u) = \epsilon \begin{bmatrix} 0.3(u_1^2 + u_2^2) \\ -0.3(2u_1u_2) \end{bmatrix}.$$

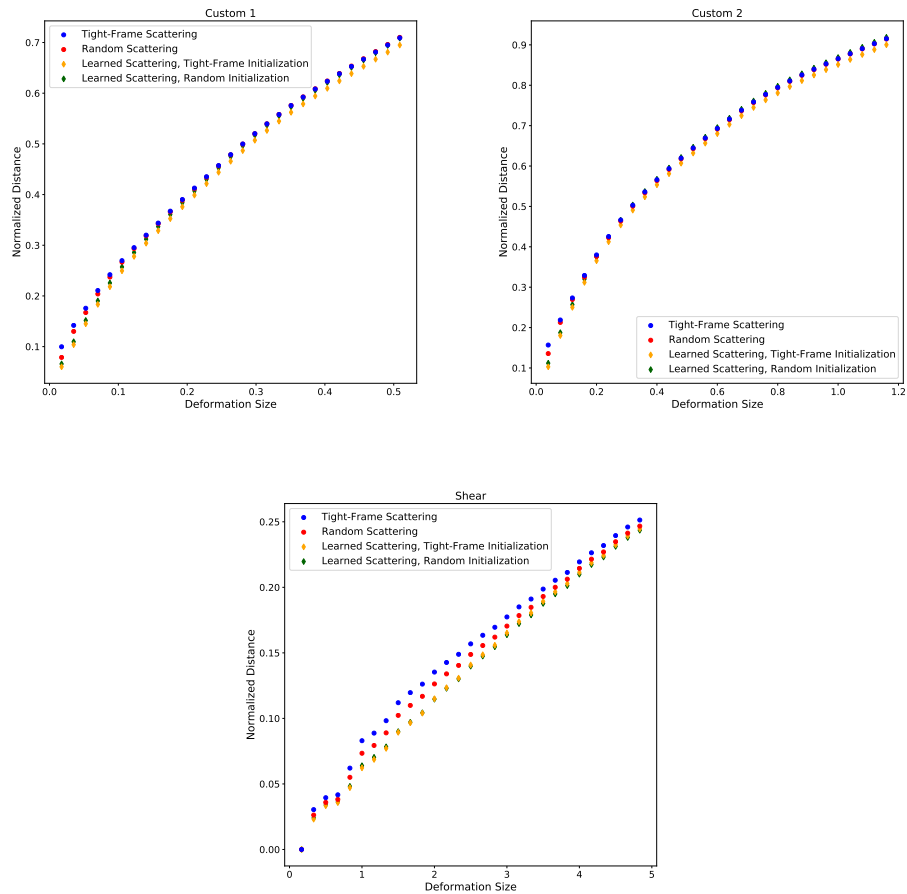


Figure 12: **Normalized distances between scattering representations of an image and its deformation.** (Top Left) Custom 1 Transformation. (Top Right) Custom 2 Transformation. (Bottom) Shear.

I DETAILS OF TRAINING UNSUPERVISED SCATTERING WITH SIMCLR OBJECTIVE

The learnable scattering networks with tight-frame and random initializations were pretrained on CIFAR-10 via the SimCLR method using a temperature of 0.5 and batch size of 128 for 500 epochs. The following basic augmentations were used as part of the SimCLR augmentation pipeline: random crop and resize, random flip, and color distortion. The optimizer used was Adam, with a learning rate of $1e-3$, similar to settings from Chen et al. (2020). Scattering weights were then frozen and used as the backbone for a linear evaluation. For linear evaluation, we used SGD with the same optimization settings as our experiments on CIFAR-10, described in B.1.