

## Parametric Scattering Networks

Shanel Gauthier<sup>1,2</sup> \*   Benjamin Thérien<sup>2,3,4\*</sup>   Laurent Alsène-Racicot<sup>1,2</sup>   Muawiz Chaudhary<sup>3,2</sup>  
Irina Rish<sup>1,2</sup>   Eugene Belilovsky<sup>3,1†</sup>   Michael Eickenberg<sup>4†</sup>   Guy Wolf<sup>1,2†</sup>

<sup>1</sup> Université de Montréal; <sup>3</sup> Mila – Quebec AI Institute; <sup>2</sup> Concordia University, Montreal, QC, Canada  
<sup>4</sup> Waterloo University, Waterloo, ON, Canada; <sup>5</sup> Flatiron Institute, New York, NY, USA

### 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 whether this standard wavelet filterbank 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 traditional filterbank constructions may not always be necessary for scattering transforms to extract effective representations.*

### 1. Introduction

The scattering transform, proposed in [29], 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 [2, 3, 12, 17, 20, 21, 32, 34, 37, 38], 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 [12, 30]. As discussed in [12], first-order scattering coefficients are similar to SIFT descriptors [27], and higher-order scattering can provide insight into the information added with depth [30]. Moreover, theoretical and empirical study of information encoded in scattering networks indicates that they often promote linear separability, which leads to effective representations for downstream classification tasks [1, 12, 17, 31].

Scattering-based models have been shown to be useful in several applications involving scarcely annotated or limited labeled data [12, 17, 33, 36]. 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 over-parameterized 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 [8, 9, 43]. 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 [33]. 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 a CNN. This allows us to also shed light on whether the standard wavelet construction [28] 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 variabil-

\*Equal contributions. †Equal senior author contribution. This research was partially funded by NSERC CGS-M [S.G.,L.A.] and URA [B.T.] scholarships; NSERC Discovery Grant RGPIN-2021-04104. [E.B.]; IVADO PRF Grant [I.R.,E.B.,G.W.]; and CIFAR AI Chairs [I.R.,G.W.]. We acknowledge resources provided by Compute Canada and Calcul Québec. The content is solely the responsibility of the authors and does not necessarily represent the official views of funding agencies. Correspondence to: eugene.belilovsky@concordia.ca, meickenberg@flatironinstitute.org, guy.wolf@umontreal.ca

ity in input signals [12]. 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 [28, 29] 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 [24]. The question then arises: is it necessary to use the standard wavelet filterbank construction? Here, we relax the standard 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 compare scattering parameterizations obtained from optimizing 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. In Sec. 4.4, we investigate the adaptation of the parametrized scattering using an unsupervised objective. Finally, in Sec. 4.5 we evaluate the computational and memory complexity of our hybrid networks. Further technical details appear in appendices, and code accompanying the work is available at <https://github.com/bentherien/parametricScatteringNetworks>.

## 2. Related work

Learning useful representations from little training data [9] 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. [26] 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 and Denzler [8] 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. Other methods show promise by incorporating prior knowledge into the model [10, 11, 18, 22, 23]. Oyallon et al. [33] 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 and Kingsbury [15] also propose a hybrid network called a learnable ScatterNet, where learning layers are intermixed between the scattering or-

ders, unlike our work where only a few parameters governing the wavelet construction are modified. Ulicny et al. [41] 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 [6, 7, 14, 35]. Balestrierio et al. [6] and Seydoux et al. [35] learn a spline-parametrized mother wavelet for 1D problems. Similarly, Cosentino and Aazhang [14] 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. Parameterization 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 3.4. Finally, Sec. 3.3 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 [12]. 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  $\phi_J$  with a spatial window of scale  $2^J$ , such as a Gaussian smoothing function. We then convolve this filter with the signal and downsample 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  $\beta$  is a normalization constant to ensure wavelets integrate to 0 over the spatial domain,  $u' = u_1 \cos \theta + u_2 \sin \theta$ ,  $R_\theta$  is the rotation matrix of angle  $\theta$  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  $\psi_{\sigma', \theta', \xi', \gamma'}(u)$ , the traditional




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. We use the Morlet canonical wavelet parameterization. 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.

Table 1. Canonical Parameters of Morlet wavelet

Param	Role	Param	Role
$\sigma$	Gaussian window scale	$\theta$	Global orientation
$\xi$	Frequency scale	$\gamma$	Aspect Ratio

wavelet filterbank is obtained by dilating it by factors  $2^j$ ,  $0 \leq j < J$ , and rotating by  $L$  angles  $\theta$  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  $\phi_J$ . 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  $\psi_\lambda$  here,  $\lambda = (\sigma_j, \theta, \xi_j, \gamma_j)$ , to denote such wavelets indexed by  $\theta$  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

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

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.,

$$S^2 x(\lambda_1, \lambda_2, u) = ||x * \psi_{\lambda_1}| * \psi_{\lambda_2}| * \phi_J(2^J u). \quad (3)$$

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 Parameterization

While the wavelet filters are traditionally fixed, 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. To provide such data-driven optimization of scattering parameters, we show, in Appendix C, that it is possible to backpropagate through this construction. We adapted the Kymatio software package [4] to create the learnable scattering network.

### 3.3. Initialization

To evaluate the importance of the standard wavelet construction, we consider two initializations and study their impact on resulting performance in both learned and non-learned settings. First, the standard wavelet construction 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  $[0, \pi)$ . The construction ensures the resulting filter bank forms an efficient tight frame. Thus, we call this construction the tight-frame initialization. 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  $\sigma$  is selected using an exponential distribution across available scales and the spatial frequency  $\xi$  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).

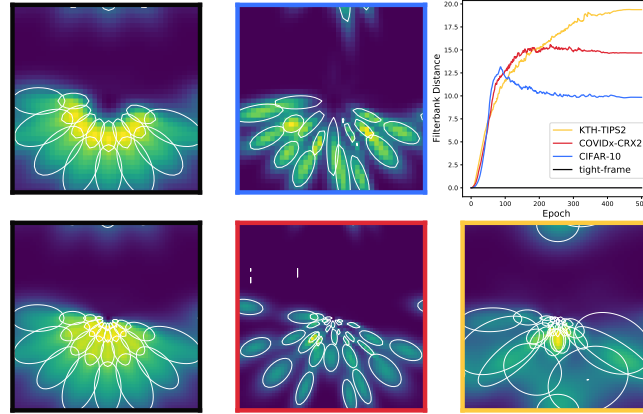


Figure 2. **Parametric scattering network learns dataset specific filters.** The graph (top right) shows the *filterbank distance* over epochs as the filters are trained on different datasets. We visualize dataset specific parameterizations of scattering filterbanks (border colors from the legend) in Fourier space. The x and y axis are the frequency axis. 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.

### 3.4. 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:  $\sigma$ ,  $\xi$ ,  $\gamma$  and  $\Theta$ . However, the global orientation of the  $L$  filters for each scale are set to be  $[\Theta, \Theta + \frac{\pi}{L}, \Theta + \frac{2\pi}{L}, \dots, \Theta + \frac{(L-1)\pi}{L}]$ .

## 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 [5, 36]. 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 of [33] and consider the scattering as the first stage of a deeper CNN, specifically a Wide Residual Network (WRN) [44]. The architecture of the WRN hybrid is described in more detail in Appendix B.

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. 3.4. 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 [19] applied directly to input data. We note that the latter is unmodified from its ImageNet architecture and that we do not initialize it with pre-trained weights.

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 [39]. Implementation details specific to each dataset are described in Appendix A. 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 [8]. The results are reported in Appendix F.

### 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 us-




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.

ing the canonical Morlet wavelet formulation with a linear classification layer 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  $\Upsilon(M_1, M_2) = \|(\sigma_1, \xi_1, \gamma_1)^T - (\sigma_2, \xi_2, \gamma_2)^T\|_2 + \text{arcdist}(\theta_1, \theta_2)$  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  $\Upsilon$  between all combinations of filter pairs from both networks, then use a minimum cost bipartite matching algorithm [25] 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  $\Upsilon$  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 12 in Appendix I.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. 10). 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. 8). 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.

## 4.2. Robustness to Deformation

In [29], 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  $\tau$  a diffeomorphism. Given the substantial changes to the filter composition in the learning process, we

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.	Parameterization	100 samples	500 samples	1000 samples	All
LS+LL†	TF	Canonical	37.84 ± 0.57	<b>52.68</b> ± 0.31	<b>57.43</b> ± 0.17	<b>69.57</b>
LS+LL†	TF	Equivariant	<b>39.69</b> ± 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	<b>43.60</b> ± 0.87	<b>63.13</b> ± 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	-	-	38.78 ± 0.72	62.97 ± 0.41	<b>71.37</b> ± 0.31	<b>96.84</b>
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

ask now whether these seem to significantly deviate from the stability result obtained from the carefully handcrafted construction proposed in [29] and extensively used in previous work e.g., [12, 17]. To evaluate the robustness of our parametric scattering networks to different geometric distortions, we apply several tractable deformations to a chest X-ray image  $x$  with varying deformation strength and encode all images with different (learned and fixed) scattering networks. The learned ones are trained using the Morlet canonical wavelet formulation with a linear classification layer. 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, shear and scale on images from the COVIDx datasets, while additional deformations are shown in Appendix G. 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 [33], 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 [33]. 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 autoaugment [16] for CIFAR-10. We used autoaugment [16] to showcase the best possible small-sample results and ablate this component in Appendix E. 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 con-

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	<b>88.50</b> ± 0.71	<b>89.90</b> ± 0.40	66.09 ± 1.05
LS+LL†	TF	Equivariant	<b>83.06</b> ± 1.53	87.56 ± 0.94	89.15 ± 0.60	<b>66.41</b> ± 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	<b>69.23</b> ± 0.67
LS+WRN†	TF	Equivariant	<b>81.86</b> ± 2.07	<b>91.56</b> ± 0.52	<b>93.97</b> ± 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	-	-	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; 23.7M for LS/S+WRN; 22.3M for WRN; 23.5M for ResNet  
 KTH: KTH-TIPS2 # params : 883K for LS/S+LL; 23.8M for LS/S+WRN; 22.3M for WRN; 23.5M for ResNet  
 †: ours; TF: Tight-Frame; LS: Learnable Scattering; S: Scattering; Rand: Random

straints, 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 D, we visualize canonical equivariant wavelet filters.

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 [42]. 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 [40]. 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 signif-

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	<b>38.05 ± 0.45</b>	<b>52.92 ± 0.28</b>	<b>57.76 ± 0.25</b>	<b>68.47 ± 0.04</b>

icantly increases the total number of parameters, therefore exhibiting a tradeoff between performance 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 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 [13], which encourages representations from two data augmentations of the same input to lie close together. We learn scattering network parameters with the canonical Morlet parameterization 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.

#### 4.5. Computational and memory complexity

The computational complexity of scattering networks and parametric scattering networks is directly related to the FFT (Fast Fourier transform),  $O(N \cdot \log(N))$  for an image of size  $(N \times N)$ . In practice, the computational and memory complexity of our parametric scattering networks varies due to a number of factors. To summarize these factors, we compare runtime (higher is faster), memory, and parameter count per architecture and image size in Table 5. The models were trained using an NVIDIA Tesla T4 GPU. We observe that fixed scattering is two to three times faster than learned scattering for all image sizes and hybrid models. In contrast, WRN-16-8 is faster than LS+WRN at image size  $32^2$ , but slower for larger images. This is due to the scattering transform’s substantial spatial dimension reduction, which leads to speed and memory benefit versus regular CNNs [29]. While gradient computation of Morlet parameters adds compute overhead, learned scattering is still efficient with much fewer parameters than CNNs.

### 5. Limitations

There are two limitations to this study that could be addressed in future research. First, the current implementa-

Table 5. Comparison of training runtime, inference runtime, GPU memory, and parameter count per architecture and image size.

Architecture	Img. size	Train (img <sub>s</sub> /sec)	Infer. (img <sub>s</sub> /sec)	GPU Mem. (GB)	#Params (Million)
LS+L	$32^2$	264	542	0.3	0.2
S+L	$32^2$	650	656	0.1	0.2
LS+WRN	$32^2$	232	430	1.6	22.6
S+WRN	$32^2$	498	588	1.4	22.6
WRN-16-8*	$32^2$	510	1695	4.2	11.0
LS+L	$128^2$	42	102	5.4	0.8
S+L	$128^2$	123	123	0.5	0.8
LS+WRN	$128^2$	42	101	8.5	23.8
S+WRN	$128^2$	115	120	3.3	23.8
WRN-16-8*	$128^2$	31	90	61.1	11.0
LS+L	$224^2$	24	72	13.7	0.5
S+L	$224^2$	78	79	1.3	0.5
LS+WRN	$224^2$	22	67	16.1	23.7
S+WRN	$224^2$	58	71	2.9	23.7
WRN-16-8*	$224^2$	10	36	49.6	11.0

\* from [45]

tion is limited to two-dimensional data. In future work, the implementation could naturally be extended to one-dimensional and three-dimensional data. Second, for popular datasets, such as CIFAR-10, there are pre-trained models available. In the study, to compare performance with our approach, we considered a fully learned WRN-16-8 and ResNet-50, but we did not consider pre-trained models.

### 6. Conclusion

This work showcases the competitive results of adapting a small number of Morlet wavelet filter parameters in scattering convolutional networks [12]. We demonstrate that filters learned by parametric scattering can be interpreted in the context of specific tasks (e.g., becoming thinner in object classification tasks that require sensitivity to edges). We also empirically demonstrate that our parametric scattering transform shares similar deformation stability to the traditional scattering transform. Overall, we find that our hybrid parametric scattering architectures (with LL and WRN) achieve state-of-the-art classification results in the low-data regime. These results verge upon bridging the gap between the handcrafted filter design in traditional scattering transforms, which provides tractable properties and supports low-parameter models, and fully parameterized convolutional neural networks, which lack interpretability but are more flexible.

In the future, our results can lead to work investigating the impact of downsampling 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.



## References

- [1] 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)*, pages 1–6, 2015. 1
- [2] Joakim Andén, Vincent Lostanlen, and Stéphane Mallat. Joint time–frequency scattering. *IEEE Transactions on Signal Processing*, 67(14):3704–3718, 2019. 1
- [3] Joakim Andén and Stéphane Mallat. Deep scattering spectrum. *IEEE Transactions on Signal Processing*, 62(16):4114–4128, 2014. 1
- [4] 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. 3
- [5] Idan Azuri and Daphna Weinshall. Generative latent implicit conditional optimization when learning from small sample. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 8584–8591, 2021. 4
- [6] Randall Balestriero, Romain Cosentino, Hervé Glotin, and Richard Baraniuk. Spline filters for end-to-end deep learning. In *International conference on machine learning*, pages 364–373. PMLR, 2018. 2
- [7] Randall Balestriero, Hervé Glotin, and Richard G Baraniuk. Interpretable super-resolution via a learned time-series representation. arXiv:2006.07713, 2020. 2
- [8] 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*, pages 1371–1380, 2020. 1, 2, 4
- [9] Nihar Bendre, Hugo Terashima Marín, and Peyman Najafirad. Learning from few samples: A survey. arXiv:2007.15484, 2020. 1, 2
- [10] 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. 2
- [11] 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. 2
- [12] Joan Bruna and Stéphane Mallat. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013. 1, 2, 6, 8
- [13] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. *CoRR*, abs/2002.05709, 2020. 8, 5
- [14] Romain Cosentino and Behnaam Aazhang. Learnable group transform for time-series. In *International Conference on Machine Learning*, pages 2164–2173. PMLR, 2020. 2
- [15] Fergal Cotter and Nick Kingsbury. A learnable scatternet: Locally invariant convolutional layers. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 350–354. IEEE, 2019. 2
- [16] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. arXiv:1805.09501, 2018. 6
- [17] 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. 1, 6
- [18] Robert Gens and Pedro M Domingos. Deep symmetry networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. 2
- [19] 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*, pages 770–778, 2016. 4
- [20] Matthew Hirn, Stéphane Mallat, and Nicolas Poilvert. Wavelet scattering regression of quantum chemical energies. *Multiscale Modeling & Simulation*, 15(2):827–863, 2017. 1
- [21] Matthew Hirn, Nicolas Poilvert, and Stéphane Mallat. Quantum energy regression using scattering transforms. arXiv:1502.02077, 2015. 1
- [22] Jorn-Henrik Jacobsen, Jan van Gemert, Zhongyu Lou, and Arnold W. M. Smeulders. Structured receptive fields in cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [23] Osman Semih Kayhan and Jan C. van Gemert. On translation invariance in cnns: Convolutional layers can exploit absolute spatial location. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [24] 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. 2
- [25] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [26] 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*, pages 8109–8118, 2018. 2
- [27] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 1
- [28] Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999. 1, 2
- [29] Stéphane Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012. 1, 2, 5, 6, 8
- [30] Stéphane Mallat. Understanding deep convolutional networks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150203, 2016. 1
- [31] Edouard Oyallon, Eugene Belilovsky, and Sergey Zagoruyko. Scaling the scattering transform: Deep

- hybrid networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5618–5627, 2017. [1](#)
- [32] Edouard Oyallon, Stéphane Mallat, and Laurent Sifre. Generic deep networks with wavelet scattering. arXiv:1312.5940, 2013. [1](#)
- [33] 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. [1](#), [2](#), [4](#), [6](#)
- [34] Michael Perlmutter, Feng Gao, Guy Wolf, and Matthew Hirn. Geometric wavelet scattering networks on compact riemannian manifolds. In *Mathematical and Scientific Machine Learning*, pages 570–604. PMLR, 2020. [1](#)
- [35] 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. [2](#)
- [36] 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*, pages 1233–1240, 2013. [1](#), [4](#)
- [37] Laurent Sifre and Stéphane Mallat. Rigid-motion scattering for texture classification. arXiv:1403.1687, 2014. [1](#)
- [38] 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. [1](#)
- [39] 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, page 1100612, 2019. [4](#)
- [40] 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*, pages 4912–4920, 2017. [7](#)
- [41] Matej Ulicny, Vladimir A. Krylov, and Rozenn Dahyot. Harmonic networks with limited training samples. In *2019 27th European Signal Processing Conference (EUSIPCO)*, pages 1–5, 2019. [2](#)
- [42] 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. [7](#), [1](#)
- [43] Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov, and Gabriel J. Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. [1](#)
- [44] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016. [4](#)
- [45] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. [8](#)