

Kymatio: Scattering Transforms in Python

Mathieu Andreux, Tomás Angles, Georgios Exarchakis, Roberto Leonarduzzi,
Gaspar Rochette, Louis Thiry, John Zarka
École normale supérieure, CNRS, PSL Research University, Paris, France

Stéphane Mallat
École normale supérieure, CNRS, PSL Research University, Paris, France
Collège de France, Paris, France
Flatiron Institute, New York, NY, USA

Joakim Andén
Flatiron Institute
162 5th Avenue, New York, NY, 10010, USA

Eugene Belilovsky
Mila, Université de Montréal, Montréal, Canada

Joan Bruna, Vincent Lostanlen
New York University, New York, NY, USA

Matthew J. Hirn
Michigan State University, East Lansing, MI, USA

Edouard Oyallon
CentraleSupélec, Gif-sur-Yvette, France

Sixin Zhang
Peking University, Beijing, China

Carmine Cella, Michael Eickenberg
University of California, Berkeley, Berkeley, CA, USA

Abstract

The wavelet scattering transform is an invariant signal representation suitable for many signal processing and machine learning applications. We present the *Kymatio* software package, an easy-to-use, high-performance Python implementation of the scattering transform in 1D, 2D, and 3D that is compatible with modern deep learning frameworks. All transforms may be executed on a GPU (in addition to CPU), offering a considerable speedup over CPU implementations. The package also has a small memory footprint. The source code, documentation, and examples are available under a BSD license at <https://www.kymat.io>.

Keywords: Scattering Transform; GPUs; Wavelets; Convolutional Neural Networks; Invariance; Python

1. Introduction

Many classification and regression tasks have a degree of invariance to translations and diffeomorphisms, such as those relating to images, audio recordings, and electronic densities. The scattering transform was introduced in Mallat (2012) to build a signal representation

that is invariant to such transformations while preserving as much signal information as possible. It is defined as a convolutional network whose filters are fixed to be wavelet and lowpass averaging filters coupled with modulus nonlinearities. It has many favorable theoretical properties (Mallat, 2012; Bruna et al., 2015; Waldspurger, 2017) and enjoys considerable success as a powerful tool in modern signal processing (Adel et al., 2017; Bruna and Mallat, 2013; Andén and Mallat, 2014; Chudáček et al., 2014; Sifre and Mallat, 2013; Eickenberg et al., 2017). It is also effective in combination with modern representation learning approaches (Oyallon et al., 2018; Sainath et al., 2014; Zeghidour et al., 2016).

This article presents *Kymatio*, a scattering transform implementation that is user-friendly, well-documented, fast, and compatible with existing automatic differentiation libraries. It brings together scattering transforms in 1D, 2D, and 3D under a *unified* application programming interface (API). The scattering network is also traversed in a depth-first manner to reduce memory requirements, enabling efficient processing in limited-memory settings, such as GPUs.

2. Implementing the Scattering Transform

Definition We consider signals defined on a grid of size $N_1 \times \dots \times N_d$ for $d = 1, 2, 3$. Given two signals $x[n]$ and $y[n]$ on this grid, we denote their periodic convolution by $x \circledast y[n]$. The second-order scattering transform is defined using two wavelet filter banks $\{\psi_{\lambda_1}^{(1)}[n]\}_{\lambda_1 \in \Lambda_1}$ and $\{\psi_{\lambda_2}^{(2)}[n]\}_{\lambda_2 \in \Lambda_2}$, where λ_1 and λ_2 are frequency indices in the sets Λ_1 and Λ_2 . It also includes a lowpass filter $\varphi_J[n]$, where the integer $J > 0$ specifies the averaging scale 2^J of the filter. Together with a non-linearity $\rho(t)$, these filters define the scattering transform.

The zeroth-order scattering coefficient $S_0x[n]$ is the local average given by $S_0x[n] = x \circledast \varphi_J[n]$. Convoluting $x[n]$ with the first-order wavelet filter bank $\{\psi_{\lambda_1}^{(1)}[n]\}_{\lambda_1 \in \Lambda_1}$, applying $\rho(t)$, and convoluting with $\varphi_J[n]$, we obtain the first-order scattering coefficients

$$S_1x[n, \lambda_1] = \rho(x \circledast \psi_{\lambda_1}^{(1)} \circledast \varphi_J[n]) \quad \lambda_1 \in \Lambda_1. \quad (1)$$

More structure is captured by decomposing $\rho(x \circledast \psi_{\lambda_1}^{(1)}[n])$ using the second filter bank, but this is done only for a subset $\Lambda_2(\lambda_1)$ of Λ_2 since $\rho(x \circledast \psi_{\lambda_1}^{(1)}[n])$ is a low-frequency signal.

The result is then passed through $\rho(t)$ and averaged, yielding the second-order coefficients

$$S_2x[n, \lambda_1, \lambda_2] = \rho(\rho(x \circledast \psi_{\lambda_1}^{(1)} \circledast \psi_{\lambda_2}^{(2)} \circledast \varphi_J[n])) \quad \lambda_1 \in \Lambda_1, \lambda_2 \in \Lambda_2(\lambda_1). \quad (2)$$

Implementation Signals obtained by filtering and applying $\rho(t)$ are low-frequency, so intermediate results are downsampled to reduce computational load as in Andén et al. (2014). In 1D and 2D, we use Morlet wavelets which are analytic (*i.e.*, complex-valued with zero energy in the negative frequencies) and the non-linearity is the complex modulus $\rho(t) = |t|$ for $t \in \mathbb{C}$ (Andén and Mallat, 2014; Bruna and Mallat, 2013). The 3D transform is calculated using solid harmonic wavelets $\psi_{\lambda_1} = \psi_{j,\ell,m}$, where j indexes the scale, and ℓ, m are the azimuthal and magnetic quantum numbers. In this case the non-linearity $\rho : \mathbb{C}^{2\ell+1} \rightarrow \mathbb{R}$ is defined, with a slight abuse of notation, as $\rho(x \circledast \psi_{j,\ell}) = \frac{1}{m} |x \circledast \psi_{j,\ell,m}|^2$ (Eickenberg et al., 2017). Following Oyallon et al. (2018), we compute the scattering transform in a depth-first manner, reducing the number of intermediate signals stored at a given time.

	Dimension	GPU	Diff.	Core Devs.	License	Language
ScatNet	1D, 2D			5	Apache 2.0	MATLAB
ScatNetLight	2D			2	GPLv2	MATLAB
PyScatWave	2D	..(3	BSD-3	Python
Scattering.m	1D			1	GPLv3	MATLAB
PyScatHarm	3D	..(1	BSD-3	Python
Wavelet Toolbox	1D			N/A	Proprietary	MATLAB
<i>Kymatio</i>	1D, 2D, 3D	..(..(15	BSD-3	Python

Table 1: Comparison to existing scattering transform packages.

3. Project vision

Code quality Adopting the philosophy of *scikit-learn* (Pedregosa et al., 2011), the goal of *Kymatio* is not to maximize the number of features, but to provide a stable and easy-to-use framework. To this end, we make heavy use of unit tests, minimize the number of dependencies, and strive for intuitive interfaces inspired by modern deep learning paradigms. *Kymatio* also provides an extensive user guide, including an API reference, a tutorial, installation instructions, and fifteen examples, several of which feature real-world applications.

Community and Bug Tracking *Kymatio* is free and open-source software, licensed under the New (3-clause) BSD license. The members of its core development team all have experience implementing scattering transforms as part of other packages. A key goal of *Kymatio* is to combine these disparate efforts and foster a community effort in order to produce high-quality software and maintain a critical mass of contributors for its maintenance. The package was released publicly on GitHub November 17th, 2018. The main communication channels is the GitHub page for questions, bug reports, and feature requests. There is also a dedicated Slack channel for *Kymatio* development.

Relation to previous software Aside from the emphasis on code quality and usability, *Kymatio* provides several improvements over previous scattering implementations:

- *Python* is the *de facto* standard for data science software, but most existing scattering packages are implemented in MATLAB. In contrast, *Kymatio* provides a completely Pythonic implementation, enabling integration with the scientific Python ecosystem.
- *GPU compatibility* is critical to many data science workloads. *Kymatio* offers an easy-to-use GPU implementation for scattering transforms in 1D, 2D, and 3D.
- *PyTorch* integration is achieved by designing scattering transforms to closely mimic PyTorch modules, allowing seamless integration into many deep learning workflows.
- *Differentiability* of the scattering transform simplifies applications in reconstruction and generative modeling, among others.

Table 1 provides a detailed comparison of existing implementations: ScatNet (Andén et al., 2014), ScatNetLight (Oyallon and Mallat, 2015), PyScatWave (Oyallon et al., 2018), Scattering.m (Lostanlen and Mallat, 2015), PyScatHarm (Eickenberg et al., 2018), and the scattering transform implemented in the MATLAB Wavelet Toolbox.

4. User Interface and Documentation

Interface The interface is designed to be flexible and consistent across inputs. We first create a scattering object by specifying the averaging scale J and the input signal shape.

```
S = Scattering1D(J, shape=(length,))
S = Scattering2D(J, shape=(height, width))
S = Scattering3D(J, shape=(height, width, depth))
```

The resulting object S acts like a `nn.Module` object in PyTorch. The scattering transform S is applied through calls of the form

```
x = torch.randn((28, 28))
output = S(x)
```

Switching from GPU or CPU functionality also follows the API of `nn.module`.

```
S.cuda() # Run on GPU
S.cpu()  # Run on CPU
```

Documentation and Examples Several examples are provided with the code, illustrating the power of *Kymatio*. These include image reconstruction and generation from scattering (Angles and Mallat, 2018), hybrid scattering and CNN training on CIFAR and MNIST (Oyallon et al., 2018), regression of molecular properties on QM7/QM9 using solid harmonic scattering (Eickenberg et al., 2017), and classifying recordings of spoken digits.

5. Conclusion

Kymatio provides a well documented, user-friendly, and fast implementation for the scattering transform. It can be used with the PyTorch deep learning framework and supports a variety of applications that have been previously inaccessible to non-experts including hybrid deep learning, generative modeling, and 3-D chemistry applications. Future work includes further optimization for speed, flexibility, and backend support.

Acknowledgments

We thank Laurent Sifre, Sergey Zagoruyko and Gabriel Huang for their helpful comments. The project was supported by ERC InvariantClass 320959. EB is partially supported by a Google Focused Research Award. MJH is partially supported by Alfred P. Sloan Fellowship #FG-2016-6607, DARPA Young Faculty Award #D16AP00117, and NSF grant #1620216.

References

T. Adel, T. Cohen, M. Caan, M. Welling, et al. 3D scattering transforms for disease classification in neuroimaging. *NeuroImage: Clinical*, 14:506–517, 2017.

- J. Andén et al. Scatnet. *Computer Software*, 2014. URL <http://www.di.ens.fr/data/software/scatnet>.
- J. Andén and S. Mallat. Deep scattering spectrum. *IEEE Trans. Signal Process.*, 62(16): 4114–4128, Aug 2014. doi: 10.1109/TSP.2014.2326991.
- T. Angles and S. Mallat. Generative networks as inverse problems with scattering transforms. In *Proc. ICLR*, 2018.
- J. Bruna and S. Mallat. Invariant scattering convolution networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8):1872–1886, 2013.
- J. Bruna, S. Mallat, E. Bacry, and J.-F. Muzy. Intermittent process analysis with scattering moments. *Ann. Statist.*, 43(1):323–351, 02 2015. doi: 10.1214/14-AOS1276.
- V. Chudáček et al. Scattering transform for intrapartum fetal heart rate variability fractal analysis: A case-control study. *IEEE Trans. Biomed. Eng.*, 61(4):1100–1108, 2014.
- M. Eickenberg et al. Solid harmonic wavelet scattering: Predicting quantum molecular energy from invariant descriptors of 3D electronic densities. In *Proc. NIPS*. 2017.
- M. Eickenberg et al. Solid harmonic wavelet scattering for predictions of molecule properties. *The Journal of Chemical Physics*, 148(24):241732, 2018.
- V. Lostanlen and S. Mallat. Wavelet scattering on the pitch spiral. In *Proc. DAFx*, 2015.
- S. Mallat. Group invariant scattering. *Comm. Pure Appl. Math.*, 65(10):1331–1398, 2012.
- E. Oyallon and S. Mallat. Deep roto-translation scattering for object classification. In *Proc. CVPR*, June 2015.
- E. Oyallon et al. Scattering networks for hybrid representation learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2018.
- F. Pedregosa et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.*, 12 (Oct):2825–2830, 2011.
- T. N. Sainath et al. Deep scattering spectra with deep neural networks for LVCSR tasks. In *Proc. Interspeech*, 2014.
- L. Sifre and S. Mallat. Rotation, scaling and deformation invariant scattering for texture discrimination. In *Proc. CVPR*, 2013.
- I. Waldspurger. Exponential decay of scattering coefficients. In *Proc. SampTA*, 2017.
- N. Zeghidour et al. A deep scattering spectrum–deep siamese network pipeline for unsupervised acoustic modeling. In *Proc. ICASSP*, pages 4965–4969. IEEE, 2016.

This figure "algorithm.png" is available in "png" format from:

<http://arxiv.org/ps/1812.11214v2>