

KERNEL COMPUTATIONS FROM LARGE-SCALE RANDOM FEATURES OBTAINED BY OPTICAL PROCESSING UNITS

Ruben Ohana[†], Jonas Wacker^{}, Jonathan Dong^{†,‡}, Sébastien Marmin^{*}
Florent Krzakala^{†,§}, Maurizio Filippone^{*}, Laurent Daudet[§]*

[†] Laboratoire de Physique de l’Ecole Normale Supérieure, ENS, Université PSL, CNRS
Sorbonne Université, Université de Paris, F-75005 Paris, France

^{*} Department of Data Science, EURECOM, 450 Route des Chappes, 06410 Biot, France

[‡] Laboratoire Kastler-Brossel, ENS, Université PSL, CNRS, Sorbonne Université,
Collège de France, Université de Paris, F-75005 Paris, France

[§] LightOn, 2 rue de la Bourse, F-75002 Paris, France

ABSTRACT

Approximating kernel functions with random features (RFs) has been a successful application of random projections for nonparametric estimation. However, performing random projections presents computational challenges for large-scale problems. Recently, a new optical hardware called Optical Processing Unit (OPU) has been developed for fast and energy-efficient computation of large-scale RFs in the analog domain. More specifically, the OPU performs the multiplication of input vectors by a large random matrix with complex-valued i.i.d. Gaussian entries, followed by the application of an element-wise squared absolute value operation – this last nonlinearity being intrinsic to the sensing process. In this paper, we show that this operation results in a dot-product kernel that has connections to the polynomial kernel, and we extend this computation to arbitrary powers of the feature map. Experiments demonstrate that the OPU kernel and its RF approximation achieve competitive performance in applications using kernel ridge regression and transfer learning for image classification. Crucially, thanks to the use of the OPU, these results are obtained with time and energy savings.

Index Terms— Kernel methods, nonparametric estimation, optical computing, random features, kernel ridge regression.

1. INTRODUCTION

Kernel methods represent a successful class of Machine Learning models, achieving state-of-the-art performance on a variety of tasks with theoretical guarantees [1, 2, 3]. Applying kernel methods to large-scale problems, however, poses computational challenges, and this has motivated a variety of contributions to develop them at scale; see, e.g., [2, 4, 5, 6, 7].

Consider a supervised learning task, and let $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ be a set of n inputs with $\mathbf{x}_i \in \mathbb{R}^d$ associated with a set of labels $\{t_1, \dots, t_n\}$. In kernel methods, it is possible to estab-

lish a mapping between inputs and labels by first mapping the inputs to a high-dimensional (possibly infinite dimensional) Hilbert space \mathcal{H} using a nonlinear feature map $\varphi : \mathbb{R}^d \rightarrow \mathcal{H}$, and then to apply the model to the transformed data. What characterizes these methods is that the mapping $\varphi(\cdot)$ does not need to be specified and can be implicitly defined by choosing a kernel function $k(\cdot, \cdot)$. While kernel methods offer a flexible class of models, they do not scale well with the number n of data points in the training set, as one needs to store and perform algebraic operations with the kernel matrix \mathbf{K} , whose entries are $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, and which require $\mathcal{O}(n^2)$ storage and $\mathcal{O}(n^3)$ operations.

In a series of celebrated papers [8, 9], Rahimi and Recht have proposed approximation techniques of the kernel function using random features (RFs), which are based on random projections of the original features followed by the application of a nonlinear transformation. In practice, the kernel function is approximated by means of the scalar product between finite-dimensional random maps $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^D$:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \langle \varphi(\mathbf{x}_i), \varphi(\mathbf{x}_j) \rangle_{\mathcal{H}} \approx \phi(\mathbf{x}_i)^\top \phi(\mathbf{x}_j) \quad (1)$$

The RF-based approximation turns a kernel-based model into a linear model with a new set of nonlinear features $\phi(\mathbf{x})$; as a result, the computational complexity is reduced from $\mathcal{O}(n^3)$ to $\mathcal{O}(ndD)$ to construct the random features and $\mathcal{O}(nD^2 + D^3)$ to optimize the linear model, where D is the RF dimension and n the number of data points. Furthermore, there is no need to allocate the kernel matrix, reducing the storage from $\mathcal{O}(n^2)$ to $\mathcal{O}(nD + D^2)$. Unless approximation strategies to compute random features are used, e.g., [10], computing RFs is one of the main computational bottlenecks.

A completely different approach was pioneered in Saade *et al.* [11], where the random projections are instead made via an analog optical device – the Optical Processing Unit (OPU) – that performs these random projections literally at the speed of light and without having to store the random

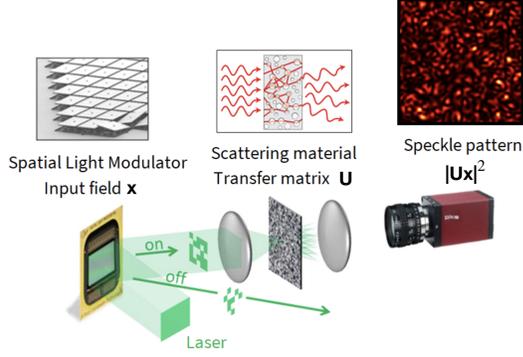


Fig. 1: Experimental setup of the Optical Processing Unit (modified with permission from [11]). The data vector is encoded in the coherent light from a laser using a DMD. Light then goes through a scattering medium and a speckle pattern is measured by a camera.

matrix in memory. Their results demonstrate that the OPU makes a significant contribution towards making kernel methods more practical for large-scale applications with the potential to drastically decrease computation time and memory, as well as power consumption. The OPU has also been applied to other frameworks like reservoir computing [12, 13] and anomaly detection [14].

Building on the milestone work of [11], the goal of the present contribution is threefold: a) we derive in full generality the kernel to which the dot product computed by the OPU RFs converges, generalizing the earlier computation of [11] to a larger class of kernels; b) we present new examples and a benchmark of applications for the kernel of the OPU; and c) we give a detailed comparison of the running time and energy consumption between the OPU and a last generation GPU.

2. THE OPTICAL PROCESSING UNIT

The principle of the random projections performed by the Optical Processing Unit (OPU) is based on the use of a heterogeneous material that scatters the light that goes through it, see Fig. 1 for the experimental setup. The data vector $\mathbf{x} \in \mathbb{R}^d$ is encoded into light using a digital micromirror device (DMD). This encoded light then passes through the heterogeneous medium, performing the random matrix multiplication. As discussed in [15], light going through the scattering medium follows many extremely complex paths, that depend on refractive index inhomogeneities at random positions. For a fixed scattering medium, the resulting process is still linear, deterministic, and reproducible. Reproducibility is important as all our data vectors need to be multiplied by the same realisation of the random matrix.

After going through the "random" medium, we observe a speckle figure on the camera, where the light intensity at each point is modelled by a sum of the components of \mathbf{x} weighted by random coefficients. Measuring the intensity of light induces a non-linear transformation of this sum, leading to:

Proposition 1. Given a data vector $\mathbf{x} \in \mathbb{R}^d$, the random feature map performed by the Optical Processing Unit is:

$$\phi(\mathbf{x}) = \frac{1}{\sqrt{D}} |\mathbf{U}\mathbf{x}|^2 \quad (2)$$

where $\mathbf{U} \in \mathbb{C}^{D \times d}$ is a complex Gaussian random matrix whose elements $U_{i,j} \sim \mathcal{CN}(0, 1)$, the variance being set to one without loss of generality, and depends on a multiplicative factor combining laser power and attenuation of the optical system. We will name these RFs optical random features.

3. COMPUTING THE KERNEL

When we map two data points $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ into a feature space of dimension D using the optical RFs of Eq. 2, we have to compute the following to obtain the associated kernel k_2 :

$$k_2(\mathbf{x}, \mathbf{y}) \approx \phi(\mathbf{x})^\top \phi(\mathbf{y}) = \frac{1}{D} \sum_{j=1}^D |\mathbf{x}^\top \mathbf{u}^{(j)}|^2 |\mathbf{y}^\top \mathbf{u}^{(j)}|^2 \quad (3)$$

$$\stackrel{D \rightarrow +\infty}{\approx} \int |\mathbf{x}^\top \mathbf{u}|^2 |\mathbf{y}^\top \mathbf{u}|^2 \mu(\mathbf{u}) d\mathbf{u} \quad (4)$$

with $\mathbf{U} = [\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(D)}]^\top$ and $\mathbf{u}^{(j)} \in \mathbb{R}^d, \forall j \in \{1, \dots, D\}$.

Theorem 1. The kernel k_2 approximated by the dot product of optical random features of Eq. 2 is given by:

$$k_2(\mathbf{x}, \mathbf{y}) = \|\mathbf{x}\|^2 \|\mathbf{y}\|^2 + (\mathbf{x}^\top \mathbf{y})^2 \quad (5)$$

where the norm is the l_2 norm.

Proof. By rotational invariance of the complex Gaussian distribution, we can fix $\mathbf{x} = \|\mathbf{x}\| \mathbf{e}_1$ and $\mathbf{y} = \|\mathbf{y}\| (\mathbf{e}_1 \cos \theta + \mathbf{e}_2 \sin \theta)$, with θ being the angle between \mathbf{x} and \mathbf{y} , \mathbf{e}_1 and \mathbf{e}_2 being two orthonormal vectors. Letting $\mathbf{e}_i^\top \mathbf{u} = u_i \sim \mathcal{CN}(0, 1)$, $i = 1, 2$ and u_1^* be the complex conjugate of u_1 , we obtain:

$$\begin{aligned} k_2(\mathbf{x}, \mathbf{y}) &= \|\mathbf{x}\|^2 \|\mathbf{y}\|^2 \int |u_1|^2 |u_1 \cos \theta + u_2 \sin \theta|^2 d\mu(\mathbf{u}) \\ &= \|\mathbf{x}\|^2 \|\mathbf{y}\|^2 \int \left(|u_1|^4 \cos^2 \theta + |u_1|^2 |u_2|^2 \sin^2 \theta \right. \\ &\quad \left. + 2|u_1|^2 \operatorname{Re}(u_1^* u_2) \cos \theta \sin \theta \right) d\mu(u_1) d\mu(u_2) \end{aligned}$$

By a parity argument, the third term in the integral vanishes, and the remaining ones can be explicitly computed, yielding:

$$k_2(\mathbf{x}, \mathbf{y}) = \|\mathbf{x}\|^2 \|\mathbf{y}\|^2 (1 + \cos^2 \theta) = \|\mathbf{x}\|^2 \|\mathbf{y}\|^2 + (\mathbf{x}^\top \mathbf{y})^2 \quad \square$$

Numerically, one can change the exponent of the feature map to $m \in \mathbb{R}^+$, which, using notations of Eq. 2, becomes:

$$\phi(\mathbf{x}) = \frac{1}{\sqrt{D}} |\mathbf{U}\mathbf{x}|^m \quad (6)$$

Theorem 2. When the exponent m is even, i.e. $m = 2s$, $\forall s \in \mathbb{N}$, the dot product of feature maps of Eq. 6 tends to the kernel k_{2s} (for $D \rightarrow \infty$):

$$k_{2s}(\mathbf{x}, \mathbf{y}) = \|\mathbf{x}\|^m \|\mathbf{y}\|^m \sum_{i=0}^s (s!)^2 \binom{s}{i}^2 \frac{(\mathbf{x}^\top \mathbf{y})^{2i}}{\|\mathbf{x}\|^{2i} \|\mathbf{y}\|^{2i}} \quad (7)$$

Moreover, a generalization $\forall m \in \mathbb{R}^+$ can be established.

Eq. 7 is connected to the polynomial kernel [1] defined as:

$$(\nu + \mathbf{x}^\top \mathbf{y})^p = \sum_{i=0}^p \binom{p}{i} \nu^{p-i} (\mathbf{x}^\top \mathbf{y})^i \quad (8)$$

with $\nu \geq 0$ and $p \in \mathbb{N}$ the order of the kernel. For $\nu = 0$ the kernel is called homogeneous. For $\nu > 0$ the polynomial kernel consists of a sum of lower order homogeneous polynomial kernels up to order p . It can be seen as having richer feature vectors including all lower-order kernel features. For optical RFs raised to the power of $s \in \mathbb{N}$ we have a sum of homogeneous polynomial kernels taken to even orders up to $m = 2s$.

Since $\mathbf{x}^\top \mathbf{y} = \|\mathbf{x}\| \|\mathbf{y}\| \cos \theta$, the kernel scales with $\|\mathbf{x}\|^i \|\mathbf{y}\|^i$, which is characteristic to any homogeneous polynomial kernel. It is easy to extend this relation to the inhomogeneous polynomial kernel by appending a bias to the input vectors, i.e. $\mathbf{x}'^\top \mathbf{y}' = \nu + \mathbf{x}^\top \mathbf{y}$ when $\mathbf{x}' = (\sqrt{\nu}, \mathbf{x}_1, \dots, \mathbf{x}_d)^\top$ and $\mathbf{y}' = (\sqrt{\nu}, \mathbf{y}_1, \dots, \mathbf{y}_d)^\top$. A practical drawback of this approach is that increasing the power of the optical RFs also increases their variance. Thus, convergence requires higher projection dimensions. Although high dimensional projections can be computed easily using the OPU, solving models on top of them poses other challenges that require special treatment [2] (e.g. Ridge Regression scales cubically with D). Therefore, we did not include these cases in the experiments in the next section and leave them for future research.

4. EXPERIMENTS

In this section, we assess the usefulness of optical RFs for different settings and datasets. The model of our choice in each case is Ridge Regression. OPU experiments were performed remotely on the OPU prototype "Vulcain", running in the LightOn Cloud with library `LightOnOPU v1.0.2`. Since the current version only supports binary input data we decide to binarize inputs for all experiments using a threshold binarizer. The code of the experiments is publicly available¹.

4.1. Optical random features for Fashion MNIST

We compare optical RFs (simulated as well as physical) to an RBF Fourier Features baseline for different projection dimensions D on Fashion MNIST. We use individually optimized hyperparameters for all RFs that are found for $D = 5000$ using an extensive grid search on a held-out validation set.

¹ <https://github.com/joneswack/opu-kernel-experiments>

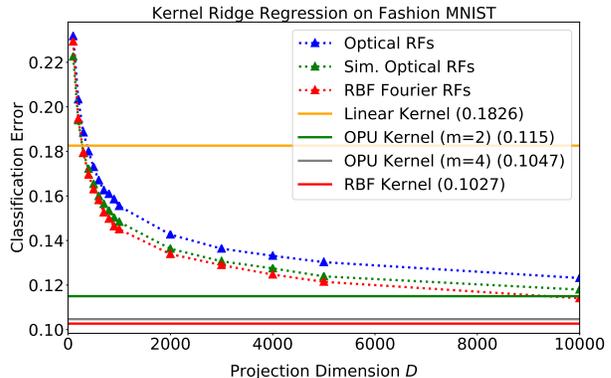


Fig. 2: Test error on Fashion MNIST for different RFs and projection dimensions. Horizontal lines show the test error using the true kernel. Standard deviations for different seeds are negligibly small and not shown in the plot. Optical RFs for $m = 4$ perform slightly worse than the ones in the plot ($m = 2$) due to slower convergence. We did not include them for better readability.

The same hyperparameters are also used for the precise kernel limit. Fig. 2 shows that the overall classification error decreases as D increases. The RBF and synthetic optical RFs perform similarly well although the RBF features have a slight gain (roughly 0.24% for $D = 5000$). This gain is more noticeable when looking at the kernel limit. Using a higher degree OPU kernel ($m = 4$), results can be further improved but they are not shown due to slower convergence.

The real OPU loses around 0.6% accuracy for $D = 5000$, which is due to experimental noise (vibrations, electronic noise) and deviations from the theoretical model (finite pixel size, camera quantization).

4.2. Transfer learning on CIFAR-10

An interesting use case for the OPU is transfer learning for image classification. For this purpose we extract a diverse set of features from the CIFAR-10 image classification dataset using three different convolutional neural networks (ResNet34 [16], AlexNet [17] and VGG16 [18]). The networks were pretrained on the well-known ImageNet classification benchmark [19]. For transfer learning, we can either fine-tune these networks and therefore the convolutional features to the data at hand, or we can directly apply a classifier on them assuming that they generalize well enough to the data. The latter case requires much less computational resources while still producing considerable performance gains over the use of the original features. This light-weight approach can be carried out on a CPU in a short amount of time where the classification error can be improved with RFs.

We compare Optical RFs and Fourier RFs for the RBF kernel to a simple baseline that directly works with the provided convolutional features (no RFs). Table 1 shows the test errors achieved on CIFAR-10. Each column corresponds to

Architecture	ResNet34				AlexNet			VGG16				
Layer	L1	L2	L3	Final	MP1	MP2	Final	MP2	MP3	MP4	MP5	Final
Dimension d	4 096	2 048	1 024	512	576	192	9 216	8 192	4 096	2 048	512	25 088
Sim. Opt. RFs	30.4	24.7	28.9	11.6	38.1	41.9	19.6	28.2	20.5	20.7	29.8	15.2 (12.9)
Optical RFs	31.1	25.7	29.7	12.3	39.2	42.6	20.8	30.9	23.3	21.5	30.2	16.4
RBF Four. RFs	30.1	25.2	30.0	12.3	39.4	41.9	19.1	28.0	20.7	20.7	30.1	14.8 (13.0)
No RFs	31.3	26.7	33.5	14.7	44.6	48.8	19.6	27.1	21.0	22.5	34.8	13.3

Table 1: Test errors (in %) on CIFAR-10 using $D = 10^4$ RFs for each kernel (except linear). Features were extracted from intermediate layers when using the original input size (32x32). Final convolutional layers were used with upscaled inputs (224x224). L(i) refers to the i th ResNet34 layer and MP(i) to the i th MaxPool layer of VGG16/AlexNet. Values for the kernel limit are shown in parenthesis (last column).

convolutional features extracted from a specific layer of one of the three networks.

Since the projection dimension $D = 10^4$ was left constant throughout the experiments, it can be observed that RFs perform particularly well compared to a linear kernel when $D \gg d$ where d is the input dimension. For the opposite case $D \ll d$ the lower dimensional projection leads to an increasing test error. This effect can be observed in particular in the last column where the test error of the RF approximation is higher than without RFs. The contrary can be achieved with large enough D as indicated by the values for the true kernel in large parenthesis. In general, the simulated as well as the physical optical RFs yield similar performances as the RBF Fourier RFs on the provided convolutional data.

4.3. Projection time and energy consumption

The main advantage of the OPU compared to a traditional CPU/GPU setup is that the OPU takes a constant time for computing RFs of arbitrary dimension D (up to $D = 10^6$ on current hardware) for a single input. Moreover, its power consumption stays below 30 W independently of the workload. Fig. 3 shows the computation time and the energy consumption over time for GPU and OPU for different projection dimensions D . In both cases, the time and energy spending do not include matrix building and loading. For the GPU, only the calls to the PyTorch function `torch.matmul` are measured and energy consumption is the integration over time of power values given by the command `nvidia-smi`.

For the OPU, the energy consumption is constant w.r.t. D and equal to 45 Joules (30 W multiplied by 1.5 seconds). The GPU computation time and energy consumption are monotonically increasing except for an irregular energy development between $D = 45\,000$ and $D = 56\,000$. This exact irregularity was observed throughout all simulations we performed and can most likely be attributed to an optimization routine that the GPU carries out internally. The OPU is up to 7 times more energy efficient than the GPU for large random projections. The GPU starts to use more energy than the OPU from $D = 18\,000$ and runs out of memory for $D = 58\,000$. The exact crossover points may change in future hardware versions. The relevant point we make here is that the OPU

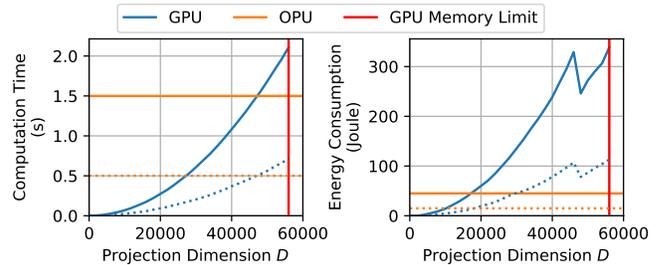


Fig. 3: Time and energy spent for computing a matrix multiplication $(n, D) \times (D, D)$. The batchsize n is 3000 (solid line) or 1000 (dotted). The curves cross each other at the same D independent from n . We verified more precisely that time and energy are linear with n for both OPU and GPU (experiments were run on a NVIDIA P100).

has a better scalability in D with respect to computation time and energy consumption.

Conclusion and perspectives

The increasing size of available data and the benefit of working in high-dimensional spaces led to an emerging need for dedicated hardware. GPUs have been used with great success to accelerate algebraic computations for kernel methods and deep learning. Yet, they rely on finite memory, consume large amounts of energy and are very expensive.

In contrast, the OPU is a scalable memory-less hardware with reduced power consumption. In this paper, we showed that optical RFs are useful in their natural form and can be modified to yield more flexible kernels. In the future, algorithms should be developed to deal with large-scale RFs, and other classes of kernels and applications should be obtained using optical RFs.

Acknowledgements

RO acknowledges support by grants from Région Ile-de-France. MF acknowledges support from the AXA Research Fund and the Agence Nationale de la Recherche (grant ANR-18-CE46-0002). FK acknowledges support from ANR-17-CE23-0023 and the Chaire CFM-ENS. The authors thank LightOn for the access to the OPU and their kind support.

5. REFERENCES

- [1] Bernhard Schölkopf, Alexander J Smola, Francis Bach, et al., *Learning with kernels: support vector machines, regularization, optimization, and beyond*, MIT press, 2002.
- [2] Alessandro Rudi, Luigi Carratino, and Lorenzo Rosasco, “Falkon: An optimal large scale kernel method,” in *Advances in Neural Information Processing Systems*, 2017, pp. 3888–3898.
- [3] Andrea Caponnetto and Ernesto De Vito, “Optimal rates for the regularized least-squares algorithm,” *Foundations of Computational Mathematics*, vol. 7, no. 3, pp. 331–368, 2007.
- [4] Alex J Smola and Bernhard Schölkopf, “Sparse greedy matrix approximation for machine learning,” 2000.
- [5] Yuchen Zhang, John Duchi, and Martin Wainwright, “Divide and conquer kernel ridge regression,” in *Conference on Learning Theory*, 2013, pp. 592–617.
- [6] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco, “Less is more: Nyström computational regularization,” in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., pp. 1657–1665. Curran Associates, Inc., 2015.
- [7] Kurt Cutajar, Edwin Bonilla, Pietro Michiardi, and Maurizio Filippone, “Random feature expansions for deep Gaussian processes,” in *ICML 2017, 34th International Conference on Machine Learning, 6-11 August 2017, Sydney, Australia*, Sydney, AUSTRALIA, 08 2017.
- [8] Ali Rahimi and Benjamin Recht, “Random features for large-scale kernel machines,” in *Advances in neural information processing systems*, 2008, pp. 1177–1184.
- [9] Ali Rahimi and Benjamin Recht, “Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning,” in *Advances in neural information processing systems*, 2009, pp. 1313–1320.
- [10] Quoc Le, Tamás Szepesvári, and Alex Smola, “Fastfood—approximating kernel expansions in loglinear time,” in *Proceedings of the international conference on machine learning*, 2013, vol. 85.
- [11] A. Saade, F. Caltagirone, I. Carron, L. Daudet, A. Drémeau, S. Gigan, and F. Krzakala, “Random projections through multiple optical scattering: Approximating kernels at the speed of light,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 6215–6219.
- [12] Jonathan Dong, Sylvain Gigan, Florent Krzakala, and Gilles Wainrib, “Scaling up echo-state networks with multiple light scattering,” in *2018 IEEE Statistical Signal Processing Workshop (SSP)*. IEEE, 2018, pp. 448–452.
- [13] Jonathan Dong, Mushegh Rafayelyan, Florent Krzakala, and Sylvain Gigan, “Optical reservoir computing using multiple light scattering for chaotic systems prediction,” *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 26, no. 1, pp. 1–12, 2019.
- [14] Nicolas Keriven, Damien Garreau, and Iacopo Poli, “Newma: a new method for scalable model-free online change-point detection,” *arXiv preprint arXiv:1805.08061*, 2018.
- [15] Antoine Liutkus, David Martina, Sébastien Popoff, Gilles Chardon, Ori Katz, Geoffroy Lerosey, Sylvain Gigan, Laurent Daudet, and Igor Carron, “Imaging with nature: Compressive imaging using a multiply scattering medium,” *Scientific reports*, vol. 4, pp. 5552, 2014.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” *arXiv preprint arXiv:1512.03385*, 2015.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., pp. 1097–1105. Curran Associates, Inc., 2012.
- [18] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *CoRR*, vol. abs/1409.1556, 2014.
- [19] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.