

RESEARCH CENTRE

Inria Lyon Centre

IN PARTNERSHIP WITH:

Ecole normale supérieure de Lyon,
Université Claude Bernard (Lyon 1)

2024

ACTIVITY REPORT

Project-Team

OCKHAM

**Optimization, pHysical Knowledge,
Algorithms and Models**

IN COLLABORATION WITH: Laboratoire de l'Informatique du Parallélisme
(LIP)

DOMAIN

**Applied Mathematics, Computation and
Simulation**

THEME

**Optimization, machine learning and
statistical methods**

Inria

Contents

Project-Team OCKHAM	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
3 Research program	3
3.1 Axis 1: Sparsity for high-dimensional learning	4
3.2 Axis 2: Learning on graphs and learning of graphs	5
3.3 Axis 3: Dynamic and frugal learning	6
4 Application domains	7
4.1 Frugal AI on embedded devices	7
4.2 Imaging in physics and medicine	7
4.3 Interactions with computational social sciences	7
5 Social and environmental responsibility	8
6 Highlights of the year	8
6.1 Awards	8
7 New software, platforms, open data	8
7.1 New software	8
7.1.1 FAuST	8
7.1.2 skglm	9
7.1.3 Benchopt	9
7.1.4 Celer	10
7.1.5 TorchDR	10
7.1.6 lazylinop	10
8 New results	11
8.1 Integrating Structured Models in Machine Learning and Signal Processing	11
8.1.1 Optimal Transport and Machine Learning	11
8.1.2 Physics informed neural networks	11
8.1.3 Bilevel and unrolled approaches for the learning of sparse covariance matrices	11
8.1.4 New penalties and proximal operators	12
8.1.5 Inverse problems for medical imaging	12
8.1.6 Interpretable graph neural networks	12
8.2 Deep neural networks : theory and algorithms	13
8.2.1 Mathematics of deep learning: rescaling invariances, generalization bounds, and conservation laws	13
8.2.2 Quantized networks: theory and algorithms	14
8.2.3 Sparse regularization, unfolding, and approximation theory	14
8.2.4 Deep sparsity: from hardness to deformable butterfly algorithms	15
8.2.5 Plug and play methods and generative modelling	15
8.3 Statistical learning, dimension reduction, and privacy preservation	16
8.3.1 Theoretical foundations of compressive learning: sketches, kernels, and optimal transport	16
8.3.2 Practical exploration of sketching and methods with limited resources	16
8.3.3 Dimensionality reduction and optimal transport	17
8.4 Large-scale convex and nonconvex optimization	18
8.4.1 Multilevel schemes for image restoration	18
8.4.2 Subsampling methods for problems involving large datasets	18
8.4.3 Reproducible benchmarking of optimization algorithms	18
8.4.4 Algorithms for large scale sparse linear models	19

9	Bilateral contracts and grants with industry	19
9.1	Bilateral grants with industry	19
10	Partnerships and cooperations	20
10.1	National initiatives	20
10.1.1	PEPR IA project : SHARP	20
10.1.2	ANR IA Chaire : AllegroAssai	20
10.1.3	ANR DataRedux	21
10.1.4	ANR Darling	21
10.1.5	ANR JCJC MASSILIA	21
10.1.6	ANR JCJC Multisc-In	22
10.1.7	ANR JCJC EROSION	22
10.1.8	ANR JCJC MEPHISTO	22
10.1.9	DI2A - Subvention Simone et Cino del Duca, Institut de France.	23
10.1.10	GDR ISIS project PROSSIMO	23
10.1.11	ANR TSIA BenchArk	24
10.1.12	ANR SEIZURE	24
11	Dissemination	24
11.1	Promoting scientific activities	24
11.1.1	Scientific events: organisation	24
11.1.2	Scientific events: selection	25
11.1.3	Journal	25
11.1.4	Invited talks	25
11.1.5	Leadership within the scientific community	25
11.1.6	Scientific expertise	26
11.1.7	Research administration	26
11.2	Teaching - Supervision - Juries	26
11.2.1	Teaching	26
11.2.2	Supervision	26
11.2.3	Juries	27
11.3	Popularization	27
11.3.1	Productions (articles, videos, podcasts, serious games, ...)	27
11.3.2	Participation in Live events	27
12	Scientific production	27
12.1	Major publications	27
12.2	Publications of the year	28
12.3	Cited publications	30

Project-Team OCKHAM

Creation of the Project-Team: 2023 March 01

Keywords

Computer sciences and digital sciences

- A3.4.1. – Supervised learning
- A3.4.4. – Optimization and learning
- A3.4.6. – Neural networks
- A3.4.7. – Kernel methods
- A3.4.8. – Deep learning
- A3.5. – Social networks
- A3.5.1. – Analysis of large graphs
- A5.3.2. – Sparse modeling and image representation
- A5.8. – Natural language processing
- A5.9. – Signal processing
- A5.9.4. – Signal processing over graphs
- A5.9.5. – Sparsity-aware processing
- A5.9.6. – Optimization tools
- A6.3.1. – Inverse problems
- A8.2. – Optimization
- A8.6. – Information theory
- A8.12. – Optimal transport

Other research topics and application domains

- B2.6. – Biological and medical imaging
- B6.6. – Embedded systems
- B7.2.1. – Smart vehicles
- B9.5.1. – Computer science
- B9.5.2. – Mathematics
- B9.5.6. – Data science
- B9.10. – Privacy

1 Team members, visitors, external collaborators

Research Scientists

- Remi Gribonval [Team leader, INRIA, Senior Researcher]
- Paulo Goncalves [INRIA, Senior Researcher]
- Mathurin Massias [INRIA, Researcher]
- Titouan Vayer [INRIA, Researcher]

Faculty Members

- Marion Foare [CPE LYON, Associate Professor]
- Elisa Riccietti [ENS DE LYON]

Post-Doctoral Fellows

- Etienne Lasalle [ENS DE LYON, Post-Doctoral Fellow]
- Rémi Vaudaine [ENS DE LYON, Post-Doctoral Fellow, until Jul 2024]

PhD Students

- Alice Brenon [ENS DE LYON, ATER, from Sep 2024]
- Mael Chaumette [INRIA, from Nov 2024]
- Edgar Desainte-Mareville [ENS DE LYON, from Nov 2024]
- Anne Gagneux [UNIV LYON I]
- Antoine Gonon [ENS DE LYON, until Nov 2024]
- Guillaume Lauga [ENS DE LYON, from Nov 2024]
- Guillaume Lauga [INRIA, until Oct 2024]
- Arthur Lebeurrier [ENS DE LYON, from Oct 2024]
- Sibylle Marcotte [ENS PARIS]
- Can Pouliquen [ENS DE LYON]
- Léon Zheng [VALEO AI, until May 2024]

Technical Staff

- Pascal Carrivain [INRIA, Engineer]

Interns and Apprentices

- Mael Chaumette [INRIA, Intern, from Apr 2024 until Oct 2024]
- Arthur Lebeurrier [ENSIMAG, Intern, from Mar 2024 until Sep 2024]

Administrative Assistant

- Emilie Gatignol [INRIA, from Sep 2024]

External Collaborator

- Márton Karsai [UNIV CEU, until Sep 2024]

2 Overall objectives

Building on a culture at the interface of signal modeling, mathematical optimization and statistical machine learning, the global objective of OCKHAM is to develop **computationally efficient and mathematically founded methods and models to process high-dimensional data**. Our ambition is to develop **frugal signal processing and machine learning methods** able to **exploit structured models**, intrinsically associated to **resource-efficient implementations**, and endowed with **solid statistical guarantees**.

Challenge 1: Developing frugal methods with robust expressivity. The idea of frugal approaches means algorithms relying on a controlled use of computing resources, but also methods whose expressivity and flexibility provably relies on the versatile notion of sparsity. This is expected to avoid the current pitfalls of costly over-parameterizations and to robustify the approaches with respect to adversarial examples and overfitting. More specifically, it is essential to contribute to the understanding of methods based on neural networks, in order to improve their performance and most of all, their efficiency in resource-limited environments.

Challenge 2: Integrating models in learning algorithms. To make statistical machine learning both more frugal and more interpretable, it is important to develop techniques able to exploit not only high-dimensional data but also models in various forms when available. When some partial knowledge is available about some phenomena related to the processed data, e.g. under the form of a physical model such as a partial differential equation, or as a graph capturing local or non-local correlations, the goal is to use this knowledge as an inspiration to adapt machine learning algorithms. The main challenge is to flexibly articulate a priori knowledge and data-driven information, in order to achieve a controlled extrapolation of predicted phenomena much beyond the particular type of data on which they were observed, and even in applications where training data is scarce.

Challenge 3: Guarantees on interpretability, explainability, and privacy. The notion of sparsity and its structured avatars –notably via graphs– is known to play a fundamental role in ensuring the identifiability of decompositions in latent spaces, for example for high-dimensional inverse problems in signal processing. The team’s ambition is to deploy these ideas to ensure not only frugality but also some level of explainability of decisions and an interpretability of learned parameters, which is an important societal stake for the acceptability of “algorithmic decisions”. Learning in small-dimensional latent spaces is also a way to spare computing resources and, by limiting the public exposure of data, it is expected to enable tunable and quantifiable tradeoffs between the utility of the developed methods and their ability to preserve privacy.

3 Research program

This project is resolutely at the interface of signal modeling, mathematical optimization and statistical machine learning, and concentrates on scientific objectives that are both ambitious –as they are difficult and subject to a strong international competition– and realistic thanks to the richness and complementarity of skills they mobilize in the team.

Sparsity constitutes a backbone for this project, not only as a target to ensure resource-efficiency and privacy, but also as prior knowledge to be exploited to ensure the identifiability of parameters and the interpretability of results. **Graphs** are its necessary *alter ego*, to flexibly model and exploit relations between variables, signals, and phenomena, whether these relations are known a priori or to be inferred from data. Lastly, **advanced large-scale optimization** is a key tool to handle in a statistically controlled and algorithmically efficient way the dynamic and incremental aspects of learning in varying environments.

The scientific activity of the project is articulated around the three axes described below. A common endeavor to these three axes consists in designing structured low-dimensional models, algorithms of bounded complexity to adjust these models to data through learning mechanisms, and a control of the performance of these algorithms to exploit these models on tasks ranging from low-level signal processing to the extraction of high-level information.

3.1 Axis 1: Sparsity for high-dimensional learning.

As now widely documented, the fact that a signal admits a sparse representation in some signal dictionary [50] is an enabling factor not only to address a variety of inverse problems with high-dimensional signals and images, such as denoising, deconvolution, or declipping, but also to speedup or decrease the cost of the acquisition of analog signals in certain scenarios compatible with compressive sensing [51, 43]. The flexibility of the models, which can incorporate learned dictionaries [81], as well as structured and/or low-rank variants of the now-classical sparse modeling paradigm [60], has been a key factor of the success of these approaches. Another important factor is the existence of algorithms of bounded complexity with provable performance, often associated to convex regularization and proximal strategies [41, 47], allowing to identify latent sparse signal representations from low-dimensional indirect observations.

While being now well-mastered (and in the core field of expertise of the team), these tools are typically constrained to relatively rigid settings where the unknown is described either as a sparse vector or a low-rank matrix or tensor in high (but finite) dimension. Moreover, the algorithms hardly scale to the dimensions needed to handle inverse problems arising from the discretization of physical models (e.g., for 3D wavefield reconstruction). A major challenge is to establish a comprehensive algorithmic and theoretical toolset to handle continuous notions of sparsity [45], which have been identified as a way to potentially circumvent these bottlenecks. The other main challenge is to extend the sparse modeling paradigm to resource-efficient and interpretable statistical machine learning. The methodological and conceptual output of this axis provides tools for Axes 2 and 3, which in return fuel the questions investigated in this axis.

- **1.1 Versatile and efficient sparse modeling.** The goal is to propose flexible and resource-efficient sparse models, possibly leveraging classical notions of dictionaries and structured factorization, but also the notion of sparsity in continuous domains (e.g. for sketched clustering, mixture model estimation, or image super-resolution), low-rank tensor representations, and neural networks with sparse connection patterns.

Besides the empirical validation of these models and of the related algorithms on a diversity of targeted applications, the challenge is to determine conditions under which their success can be mathematically controlled, and to determine the fundamental tradeoffs between the expressivity of these models and their complexity.
- **1.2 Sparse optimization.** The main objectives are: a) to define cost functions and regularization penalties that integrate not only the targeted learning tasks, but also a priori knowledge, for example under the form of conservation laws or as relation graphs, cf Axis 2; b) to design efficient and scalable algorithms [3, 62] to optimize these cost functions in a controlled manner in a large-scale setting. To ensure the resource-efficiency of these algorithms, while avoiding pitfalls related to the discretization of high-dimensional problems (aka curse of dimensionality), we investigate the notion of “continuous” sparsity (i.e., with sparse measures), of hierarchies (along the ideas of multilevel methods), and of reduced precision (cf also Axis 3). The nonconvexity and non-smoothness of the problems are key challenges, and the exploitation of proximal algorithms and/or convexifications in the space of Borelian measures are privileged approaches.
- **1.3 Identifiability of latent sparse representations.** To provide solid guarantees on the interpretability of sparse models obtained via learning, one needs to ensure the identifiability of the latent variables associated to their parameters. This is particularly important when these parameters bear some meaning due to the underlying physics. Vice-versa, physical knowledge can guide the choice of which latent parameters to estimate. By leveraging the team’s know-how obtained in the field of inverse problems, compressive sensing and source separation in signal processing, we aim at establishing theoretical guarantees on the uniqueness (modulo some equivalence classes to be

characterized) of the solutions of the considered optimization problems, on their stability in the presence of random or adversarial noise, and on the convergence and stability of the algorithms.

3.2 Axis 2: Learning on graphs and learning of graphs.

Graphs provide synthetic and sparse representations of the interactions between potentially high-dimensional data, whether in terms of proximity, statistical correlation, functional similarity, or simple affinities. One central task in this domain is how to infer such discrete structures, from the observations, in a way that best accounts for the ties between data, without becoming too complex due to spurious relationships. The graphical lasso [52] is among the most popular and successful algorithm to build a sparse representation of the relations between time series (observed at each node) and that unveils relevant patterns of the data. Recent works (e.g. [61]) strived to emphasize the clustered structure of the data by imposing spectral constraints to the Laplacian of the sought graphs, with the aim to improve the performance of spectral approaches to unsupervised classification. In this direction, several challenges remain, such as for instance the transposition of the framework to graph-based semi-supervised learning [1], where natural models are stochastic block models rather than strictly multi-component graphs (e.g. Gaussian mixtures models). As it is done in [88], the standard l_1 -norm penalization term of graphical lasso could be questioned in this case. On another level, when low-rank (precision) matrices and / or when preservation of privacy are important stakes, one could be inspired by the sketching techniques developed in [56] and [46] to work out a *sketched graphical lasso*. There exists other situations where the graph is known a priori and does not need to be inferred from the data. This is for instance the case when the data naturally lie on a graph (e.g. social networks or geographical graphs) and so, one has to combine this data structure with the attributes (or measures) carried by the nodes or the edges of these graphs. Graph signal processing (GSP) [78] [10], which underwent methodological developments at a very rapid pace in recent years, is precisely an approach to jointly exploit algebraically these structures and attributes, either by filtering them, by re-organizing them, or by reducing them to principal components. However, as it tends to be more and more the case, data collection processes yield very large data sets with high dimensional graphs. In contrast to standard digital signal processing that relies on regular graph structures (cycle graph or cartesian grid) treating complex structured data in a global form is not an easily scalable task [53]. Hence, the notion of distributed GSP [48, 49] has naturally emerged. Yet, very little has been done on graph signals supported on dynamical graphs that undergo vertices/edges editions.

- **2.1 Learning of graphs.** When the graphical structure of the data is not known a priori, one needs to explore how to build it or to infer it. In the case of partially known graphs, this raises several questions in terms of relevance with respect to sparse learning. For example, a challenge is to determine which edges should be kept, whether they should be oriented, and how attributes on the graph could be taken into account (in particular when considering time-series on graphs) to better infer the nature and structure of the un-observed interactions. We strive to adapt known approaches such as the *graphical lasso* to estimate the covariance under a sparsity constraint (integrating also temporal priors), and investigate diffusion approaches to study the identifiability of the graphs. In connection with Axis 1.2, a particular challenge is to incorporate a priori knowledge coming from physical models that offer concise and interpretable descriptions of the data and their interactions.
- **2.2 Distributed and adaptive learning on graphs.** The availability of a known graph structure underlying training data offers many opportunities to develop distributed approaches, open perspectives where graph signal processing and machine learning can mutually fertilize each other.

Some classifiers can be formalized as solutions of a constrained optimization problem, and an important objective is then to reduce their global complexity by developing distributed versions of these algorithms. Compared to costly centralized solutions, distributing the operations by restricting them to local node neighborhoods will enable solutions that are both more frugal and more privacy-friendly. In the case of dynamic graphs, the idea is to get inspiration from adaptive processing techniques to make the algorithms able to track the temporal evolution of data, either in terms of structural evolution or of temporal variations of the attributes. This aspect finds a natural continuation in the objectives of Axis 3.

3.3 Axis 3: Dynamic and frugal learning.

With the resurgence of neural networks approaches in machine learning, training times of the order of days, weeks, or even months are common. Mainstream research in deep learning somehow applies it to an increasingly large class of problems and uses the general wisdom to improve the models prediction accuracy by “stacking more layers”, making the approach ever more resource-hungry. Underpinning theory on which resources are needed for a network architecture to achieve a given accuracy is still in its infancy. Efficient scaling of such techniques to massive sample sizes or dimensions in a resource-restricted environment remains a challenge and is a particularly active field of academic and industrial R&D, with recent interest in techniques such as sketching, dimension reduction, and approximate optimization.

A central challenge is to develop novel approximate techniques with reduced computational and memory imprint. For certain unsupervised learning tasks such as PCA, unsupervised clustering, or parametric density estimation, random features (e.g. random Fourier features [76]) allow to compute aggregated *sketches* guaranteed to preserve the information needed to learn, and no more: this has led to the *compressive learning* framework, which is endowed with statistical learning guarantees [56] as well as privacy preservation guarantees [46]. A sketch can be seen as an embedding of the empirical probability distribution of the dataset with a particular form of kernel mean embedding [79]. Yet, designing random features given a learning task remains something of an art, and a major challenge is to design provably good end-to-end sketching pipelines with controlled complexity for supervised classification, structured matrix factorization, and deep learning.

Another crucial direction is the use of dynamical learning methods, capable of exploiting wisely multiple representations at different scales of the problem at hand. For instance, many low and mixed-precision variants of gradient-based methods have been recently proposed [86, 85], which are however based on a static reduced precision policy, while a dynamic approach can lead to much improved energy-efficiency. Also, despite their massive success, gradient-based training methods still possess many weaknesses (low convergence rate, dependence on the tuning of the learning parameters, vanishing and exploding gradients) and the use of dynamical information promises to allow for the development of alternative methods, such as second-order or multilevel methods, which are as scalable as first-order methods but with faster convergence guarantees [77, 87].

The overall objective in this axis is to adapt in a controlled manner the information that is extracted from datasets or data streams and to dynamically use such information in learning, in order to optimize the tradeoffs between statistical significance, resource-efficiency, privacy-preservation and integration of a priori knowledge.

- **3.1 Compressive and privacy-preserving learning.** The goal is to compress training datasets as soon as possible in the processing workflow, before even starting to learn. In the spirit of compressive sensing, this is desirable not only to ensure the frugal use of resources (memory and computation), but also to preserve privacy by limiting the diffusion of raw datasets and controlling the information that could actually be extracted from the targeted compressed representations, called *sketches*, obtained by well-chosen nonlinear random projections. We aim to build on a *compressive learning* framework developed by the team with the viewpoint that sketches provide an embedding of the data distribution, which should preserve some metrics, either associated to the specific learning task or to more generic optimal transport formulations. Besides ensuring the identifiability of the task-specific information from a sketch (cf Axis 1.3), an objective is to efficiently extract this information from a sketch, for example via algorithms related to avatars of continuous sparsity as studied in Axis 1.2. A particular challenge, connected with Axis 2.1 when inferring dynamic graphs from correlation of non-stationary times series, and with Axis 3.2 below, is to dynamically adapt the sketching mechanism to the analyzed data stream.
- **3.2 Sequential sparse learning.** Whether aiming at dynamically learning on data streams (cf. Axes 2.1 and 2.2), at integrating *a priori* physical knowledge when learning, or at ensuring domain adaptation for transfer learning, the objective is to achieve a statistically near-optimal update of a model from a sequence of observations whose content can also dynamically vary. When considering time-series on graphs, to preserve resource-efficiency and increase robustness, the algorithms further need to update the current models by dynamically integrating the data stream.

- **3.3 Dynamic-precision learning.** The goal is to propose new optimization algorithms to overcome the cost of solving large scale problems in learning, by dynamically adapting the precision of the data. The main idea is to exploit multiple representations at different scales of the problem at hand. We explore in particular two different directions to build the scales of problems: a) exploiting ideas coming from multilevel optimization to propose dynamical hierarchical approaches exploiting representations of the problem of progressively reduced dimension; b) leveraging the recent advances in hardware and the possibility of representing data at multiple precision levels provided by them. We aim at improving over state-of-the-art training strategies by investigating the design of scalable multilevel and mixed-precision second-order optimization and quantization methods, possibly derivative-free.

4 Application domains

The primary objectives of this project, which is rooted in Signal Processing and Machine Learning methodology, are to develop flexible methods, endowed with solid mathematical foundations and efficient algorithmic implementations, that can be adapted to numerous application domains. We are nevertheless convinced that such methods are best developed in strong and regular connection with concrete applications, which are not only necessary to validate the approaches but also to fuel the methodological investigations with relevant and fruitful ideas. The following application domains are primarily investigated in partnership with research groups with the relevant expertise.

4.1 Frugal AI on embedded devices

There is a strong need to drastically compress signal processing and machine learning models (typically, but not only, deep neural networks) to fit them on embedded devices. For example, on autonomous vehicles, due to strong constraints (reliability, energy consumption, production costs), the memory and computing resources of dedicated high-end image-analysis hardware are two orders of magnitude more limited than what is typically required to run state-of-the-art deep network models in real-time. The research conducted in the OCKHAM project finds direct applications in these areas, including: compressing deep neural networks to obtain low-bandwidth video-codecs that can run on smartphones with limited memory resources; sketched learning and sparse networks for autonomous vehicles; or sketching algorithms tailored to exploit optical processing units for energy efficient large-scale learning.

4.2 Imaging in physics and medicine

Many problems in imaging involve the reconstruction of large scale data from limited and noise-corrupted measurements. In this context, the research conducted in OCKHAM pays a special attention to modeling domain knowledge such as physical constraints or prior medical knowledge. This finds applications from physics to medical imaging, including: multiphase flow image characterization; near infrared polarization imaging in circumstellar imaging; compressive sensing for joint segmentation and high-resolution 3D MRI imaging; or graph signal processing for radio astronomy imaging with the Square Kilometer Array (SKA).

4.3 Interactions with computational social sciences

Based on collaborations with the relevant experts the team also regularly investigates applications in computational social science. For example, modeling infection disease epidemics requires efficient methods to reduce the complexity of large networked datasets while preserving the ability to feed effective and realistic data-driven models of spreading phenomena. In another area, estimating the vote transfer matrices between two elections is an ill-posed problem that requires the design of adapted regularization schemes together with the associated optimization algorithms.

5 Social and environmental responsibility

Machine learning methods achieve remarkable performance across various domains. However, the training of underlying models typically relies on significant computational resources, and consequently, energy resources. For the most high-performing models, these resources are far from negligible. Therefore, it becomes crucial to move towards more "frugality" and be capable of constructing learning models under resource constraints. We organized a workshop day titled "Frugality and Machine Learning", in partnership with IXXI, with the aim of discussing the feasibility of this objective from both technical and societal perspectives: how can we build models with minimal resources while maintaining good performance? Is it possible to surpass certain limits, such as those imposed by the rebound effect? This day brought together around fifty people, and the exchanges were fruitful. The various presentations enabled rich discussions that will contribute to reflections on the role of AI in society.

6 Highlights of the year

6.1 Awards

Sibylle Marcotte, Ph.D. student at ENS Paris - PSL and member of the Ockham team, **is among the 35 recipients of the 18th *Prix Jeunes Talents France L'Oréal-UNESCO Pour Les Femmes et la Science***.

Anne Gagneux, Ph.D. student in the Ockham team, received the best poster award at the SMAI MODE days 2024, for her work conducted during her M2 internship on "Automatic and unbiased coefficients clustering with non-convex SLOPE".

The paper titled "A path-norm toolkit for modern networks: consequences, promises and challenges" [19] was accepted as a *spotlight presentation* (5% of the accepted papers) at the ICLR 2024 conference (overall acceptance rate around 31%).

7 New software, platforms, open data

7.1 New software

7.1.1 FAuST

Keywords: Matrix calculation, Multilayer sparse factorisation

Scientific Description: FAuST allows to approximate a given dense matrix by a product of sparse matrices, with considerable potential gains in terms of storage and speedup for matrix-vector multiplications.

Functional Description: FAUST is a C++ toolbox designed to decompose a given dense matrix into a product of sparse matrices in order to reduce its computational complexity (both for storage and manipulation).

Faust includes Matlab and Python wrappers and scripts to reproduce the experimental results of the following papers: - Le Magoarou L. and Gribonval R., "Flexible multi-layer sparse approximations of matrices and applications", Journal of Selected Topics in Signal Processing, 2016. - Le Magoarou L., Gribonval R., Tremblay N. "Approximate fast graph Fourier transforms via multi-layer sparse", IEEE Transactions on Signal and Information Processing over Networks, 2018 - Quoc-Tung Le, Rémi Gribonval. Structured Support Exploration For Multilayer Sparse Matrix Factorization. ICASSP 2021 – IEEE International Conference on Acoustics, Speech and Signal Processing, Jun 2021, Toronto, Ontario, Canada. pp.1-5. - Sibylle Marcotte, Amélie Barbe, Rémi Gribonval, Titouan Vayer, Marc Sebban, et al.. Fast Multiscale Diffusion on Graphs. 2021.

Release Contributions: Faust 1.x contains Matlab routines to reproduce experiments of the PANAMA team on learned fast transforms.

Faust 2.x contains a C++ implementation with preliminary Matlab / Python wrappers.

Faust 3.x includes Python and Matlab wrappers around a C++ core with GPU acceleration, new algorithms.

URL: <https://faust.inria.fr/>

Publications: [hal-03212764](#), [hal-01416110](#), [hal-01627434](#), [hal-01167948](#), [hal-01254108](#), [tel-01412558](#), [hal-01156478](#), [hal-01104696](#), [hal-01158057](#), [hal-03132013](#)

Contact: Remi Gribonval

Participants: Luc Le Magoarou, Nicolas Tremblay, Remi Gribonval, Nicolas Bellot, Adrien Leman, Hakim Hadj-Djilani

7.1.2 skglm

Keywords: Optimization, Machine learning, Sparsity

Functional Description: skglm is a Python package that offers fast estimators for Generalized Linear Models (GLMs) that are compatible with scikit-learn. It is highly flexible and supports a wide range of GLMs. Its main feature is flexibility: you can implement virtually any estimator as a combination of datafit and penalty.

Thanks to this flexible design, skglm supports many missing models in scikit-learn while ensuring high performance. There are several reasons to opt for skglm:

- Support for many fast solvers able to tackle large datasets, either dense or sparse, with millions of features up to 100 times faster than scikit-learn
- User-friendly API than enables composing custom estimators with any combination of existing datafits and penalties
- Flexible design that makes it simple and easy to implement new datafits and penalties, a matter of few lines of code
- Estimators fully compatible with the scikit-learn API and drop-in replacements of its GLM estimators

skglm is integrated into scikit-learn via the scikit-learn-contrib organization.

URL: <https://contrib.scikit-learn.org/skglm/>

Publication: [hal-03819082](#)

Contact: Mathurin Massias

Participants: Mathurin Massias, Badr Moufad

7.1.3 Benchopt

Keywords: Benchmarking, Machine learning, Optimization

Functional Description: BenchOpt is a package to simplify, make more transparent and more reproducible the comparisons of optimization algorithms. It is written in Python but it is available with many programming languages. So far it has been tested with Python, R, Julia and compiled binaries written in C/C++ available via a terminal command. If it can be installed via conda, it should just work!

BenchOpt is used through a simple command line and ultimately running and replicating an optimization benchmark should be as easy as cloning a repo and launching the computation with a single command line. For now, BenchOpt features benchmarks for around 10 convex optimization problems and we are working on expanding this to feature more complex optimization problems. We are also developing a website to display the benchmark results easily.

Release Contributions: <https://github.com/benchopt/benchopt/releases/tag/1.5.1>

Publication: [hal-03830604](#)

Contact: Thomas Moreau

Participants: Thomas Moreau, Alexandre Gramfort, Mathurin Massias, Badr Moufad

7.1.4 Celer

Keywords: Mathematical Optimization, Machine learning, Sparsity

Functional Description: celer is a Python package that solves Lasso-like problems and provides estimators that under the popular scikit-learn API. Thanks to a tailored implementation, celer provides a fast solver that tackles large-scale datasets with millions of features up to 100 times faster than scikit-learn. It handles Lasso, ElasticNet, Group Lasso, Multitask Lasso and Sparse Logistic regression, and comes with - automated parallel cross-validation - support of sparse and dense data - optional feature centering and normalization - unpenalized intercept fitting

celer also provides easy-to-use estimators as it is designed under the scikit-learn API.

URL: <http://mathurinm.github.io/celer>

Publications: [hal-02263500](#), [hal-01833398](#)

Contact: Mathurin Massias

Participants: Badr Moufad, Alexandre Gramfort

7.1.5 TorchDR

Keywords: Optimal transportation, Machine learning, Dimensionality reduction, High Dimensional Data

Scientific Description: TorchDR is an open-source dimensionality reduction (DR) library using PyTorch. Its goal is to accelerate the development of new DR methods by providing a common simplified framework.

Functional Description: TorchDR is an open-source dimensionality reduction (DR) library using PyTorch. Its goal is to accelerate the development of new DR methods by providing a common simplified framework.

URL: <https://torchdr.github.io/>

Contact: Titouan Vayer

Participants: Hugues Van Assel, Mathurin Massias, Nicolas Courty, Remi Flamary, Cédric Vincent-Cuaz

7.1.6 lazylinop

Name: lazylinop

Keywords: Signal processing, Numerical algorithm, Scientific computing

Scientific Description: lazylinop is an easy way to combine existing operators into more complex operators with direct access to its adjoint.

Functional Description: Lazy evaluation of linear operators applied to vectors or matrices. lazylinop aims at providing an easy way to combine existing operators into more complex operators with direct access to its adjoint. Thanks to the lazy computation paradigm, lazylinop offers potential performances gains and memory sparing.

Release Contributions: - Basic linear operators: Kronecker product, addition, diagonal, block-diagonal, concatenation ... - Polynomial of linear operators. - Usual signal processing linear operators.

Work-In-Progress: - Usual image processing linear operators. - Butterfly linear operators.

URL: <https://faustgrp.gitlabpages.inria.fr/lazylinop/#>

Contact: Pascal Carrivain

Participants: Pascal Carrivain, Hakim Hadj-Djilani, Simon Delamare, Remi Gribonval

8 New results

8.1 Integrating Structured Models in Machine Learning and Signal Processing

8.1.1 Optimal Transport and Machine Learning

Participants: Titouan Vayer, Rémi Gribonval, Arthur Lebeurrier.

Optimal transport (OT) is a tool that plays a central role in various machine learning applications today, whether it's for generative models, domain adaptation, analysis of cellular dynamics, neural networks, or graph models as illustrated in our past works [12, 13]. In its original formulation, OT faces a scalability issue: solving the underlying optimization problem has cubic complexity with respect to the number of points. The introduction of regularized transport marked a groundbreaking advancement in the field by achieving a mere quadratic complexity. However, for truly large-scale applications, this quadratic complexity remains prohibitive. In the context of Arthur Lebeurrier's internship, we investigated a hierarchical factorization scheme of the kernel matrix involved in the regularized OT problem and we developed a new linear algorithm for approximating the OT problem. This work led to a poster presentation at the LORAINNE'24 workshop in Nancy.

8.1.2 Physics informed neural networks

Participants: Elisa Riccietti.

Collaboration with Serge Gratton, Valentin Mercier (IRIT, Toulouse), Philippe Toint (U. Namur, Belgium), Stefania Bellavia and Mahsa Yousefi (UNIFI, Florence, Italy).

Physics informed neural networks (PINNs) are specialized network architectures designed for the solution of partial differential equations (PDEs) that take into account the underlying physics of the problem. We investigated their use both for direct and inverse problems involving PDEs.

In the context of the postdoc of Mahsa Yousefi, we pursued the work started last year on the investigation of their ability to deal with ill-posed inverse problems, focusing especially on parameter identification problems. We investigated the regularizing properties of PINNs and the use of regularising training procedures to correctly fit noisy data in such a context.

In the context of the Ph.D. work of Valentin Mercier, we published our work [55] on the integration of a multigrid approach in the training of PINNs, a large scale optimization problem involving complex solutions with multiple frequency components. The proposed training scheme leverages the structure of modern Mscale networks and through a block coordinate strategy ensures not only to reduce the training time, but also to improve the quality of the approximated solutions.

8.1.3 Bilevel and unrolled approaches for the learning of sparse covariance matrices

Participants: Can Pouliquen, Paulo Goncalves, Mathurin Massias, Titouan Vayer.

The PhD of Can Pouliquen is devoted to the dynamic inference of brain connectivity graphs for epileptic patients. We have adopted the mathematical framework of the Graphical Lasso, and pursue two directions. First, we have developed a bilevel optimization framework, that eases the tuning of individual correlation strengths in the Graphical Lasso penalty [75]. Second, we have introduced a new deep neural network architecture for sparse covariance matrix estimation, which guarantees a simultaneously sparse and symmetric positive definite output. This highly desirable property was so far a missing feature of

existing architectures, and has many potential applications in graph learning beyond neurosciences [36]. This work was submitted to ICLR 2025.

8.1.4 New penalties and proximal operators

Participants: Anne Gagneux, Remi Gribonval, Mathurin Massias.

Collaboration with Emmanuel Soubies (CNRS, IRIT, Toulouse).

Finishing the internship work of Anne Gagneux, we have studied the properties of sorted non convex penalties. Convex sorted penalties such as SLOPE are known to automatically cluster coefficients associated to correlated variables; non convex penalties on the other hand mitigate the well-known amplitude bias of the L1 norm. Combining non-convexity with automatic grouping is therefore a promising venue. However the technical difficulties raised by such new penalties are many (non convexity, non smoothness). We have derived an algorithm based on the Pool Adjacent Violators Algorithm (PAVA) that computes the exact proximal operator of a first kind of sorted penalties (sorted MCP, sorted Log-sum). We have also extended it to compute the proximal operators of the sorted ℓ_q ($0 < q > 1$) penalties, which presented more difficulties due to non Lipschitzianity.

8.1.5 Inverse problems for medical imaging

Participants: Marion Foare.

Collaboration with Luis Enrique Amador Arya (Creatis, Villeurbanne), H el ene Ratiney (Creatis, Villeurbanne),  Eric Van Reeth (Creatis, Villeurbanne), and Siemens Healthcare, Saint Denis

It is of particular interest in the field of medical imaging to quickly acquire low-resolution volumes (compromise between acquisition time, SNR and spatial resolution), and enhance their resolution as a post-processing step. In particular, isotropic super-resolution (ISR) techniques consist in reconstructing an isotropic volume from the combination of several anisotropic volumes acquired with different orientations.

In the context of the PhD work of Luis Enrique Amador Araya, we pursued the development of specialized piecewise-smooth variational methods combining data fitting terms with geometric priors (e.g. the Discrete Mumford-Shah model) to build faithful super-resolution images in 3D Magnetic Resonance Imaging (MRI). Preliminary work has been submitted to ISBI 2025.

Second, we explore new data fidelity terms to extend this approach to multi-contrasts ISR, that is, to reconstruct isotropic and multi-contrasts high resolution images from multi-contrasts anisotropic acquisitions.

8.1.6 Interpretable graph neural networks

Participants: Titouan Vayer.

Collaborations with Pierre Borgnat (Physics Lab, ENS de Lyon).

As part of Thomas Bobille's internship, we explore the supervised classification of graph signals on a common graph, aiming to identify compelling models and scenarios for different types of graph signals.

The study also seeks to evaluate the advantages of incorporating graph structures into learning models by comparing graph-agnostic approaches, such as logistic regression, with graph-based methods, like Graph Convolutional Networks (GCN). This work serves as an initial step toward a broader objective of developing more interpretable graph neural networks.

8.2 Deep neural networks : theory and algorithms

8.2.1 Mathematics of deep learning: rescaling invariances, generalization bounds, and conservation laws

Participants: Rémi Gribonval, Antoine Gonon, Elisa Riccietti, Sibylle Marcotte.

Collaborations with Nicolas Brisebarre (ARIC team, ENS de Lyon), and with Gabriel Peyré (DMA, ENS, Paris)

Rescaling invariance in ReLU networks. Neural networks with the ReLU activation function are described by weights and bias parameters, and implemented into a piecewise linear continuous function. Natural *scalings* and permutations operations on the parameters leave the realization unchanged, leading to equivalence classes of parameters that yield the same realization.

Path-embedding and path-norm based generalization bounds. The path-embedding of parameters that we introduced in [80] was invariant to such scalings but limited to strictly layered ReLU architectures. In the context of the PhD of Antoine Gonon [25], we extended it [19] to fully encompass general DAG ReLU networks with biases, skip connections and any operation based on the extraction of order statistics: max pooling, GroupSort etc. The norm of the resulting embedding is called a path-norm, and we established a general toolkit to obtain statistical generalization bounds for such modern neural networks. The resulting bounds are not only the most widely applicable path-norm based ones, but also recover or beat the sharpest known bounds of this type. These extended path-norms further enjoy the usual benefits of path-norms: ease of computation, invariance under the symmetries of the network, and improved sharpness on feed-forward networks compared to the product of operators' norms, another complexity measure most commonly used. The versatility of the toolkit and its ease of implementation allowed us to challenge the concrete promises of path-norm-based generalization bounds, by numerically evaluating the sharpest known bounds for ResNets on ImageNet. Building on this toolkit, we more recently investigated a rescaling-invariant Lipschitz bound on the mapping from parameter space to function space and illustrated its potential for neural network pruning and quantization [28].

Conservation laws. In the thesis of Sibylle Marcotte, the above path-embedding also served as a key enabler for the analysis of conservation laws in gradient descent dynamics of ReLU networks [72]. Understanding the geometric properties of gradient descent dynamics is indeed a key ingredient in deciphering the recent success of very large machine learning models. A striking observation is that trained over-parameterized models retain some properties of the optimization initialization. This "implicit bias" is believed to be responsible for some favorable properties of the trained models and could explain their good generalization properties.

Last year our work on this topic [72] was conducted with a motivation that was threefold. First, we rigorously exposed the definition and basic properties of "conservation laws", which are maximal sets of independent quantities conserved during gradient flows of a given model (e.g. of a ReLU network with a given architecture) with any training data and any loss. Then we explained how to find the exact number of these quantities by performing finite-dimensional algebraic manipulations on the Lie algebra generated by the Jacobian of the model. Finally, we provided algorithms (implemented in SageMath) to: a) compute a family of polynomial laws; b) compute the number of (not necessarily polynomial) conservation laws. We provided showcase examples that we fully work out theoretically. Besides, applying the two algorithms confirmed for a number of ReLU network architectures that all known laws are recovered by the algorithm, and that there are no other laws. Such computational tools paved the way to understanding desirable properties of optimization initialization in large machine learning models. This year[20] we studied the notion of conservation law and the corresponding algorithms for optimization

flows associated to non-Euclidean geometries and momentum-based dynamics. We characterized "all" conservation laws in this general setting. In stark contrast to the case of gradient flows, we proved that the conservation laws for momentum-based dynamics exhibit *temporal dependence*. Additionally, we often observed a "conservation loss" when transitioning from gradient flow to momentum dynamics. Specifically, for linear networks, our framework allowed us to identify all momentum conservation laws, which are less numerous than in the gradient flow case except in sufficiently over-parameterized regimes. With ReLU networks, no conservation law remains. This phenomenon also manifests in non-Euclidean metrics, used e.g. for Nonnegative Matrix Factorization (NMF): all conservation laws can be determined in the gradient flow context, yet none persists in the momentum case. Work in progress includes extending the analysis to general DAG ReLU network architectures, using the associated extended path-embedding, to cover notably ResNets but also attention layers.

8.2.2 Quantized networks: theory and algorithms

Participants: Rémi Gribonval, Elisa Riccietti, Antoine Gonon.

Collaboration with Nicolas Brisebarre (ARIC team, ENS de Lyon), with Silviu Filip and El-Mehdi El arar (IRISA, Rennes), and with Theo Mary (LIP6, Paris)

Quantization of neural networks: theory Motivated by the importance of quantizing networks besides pruning them to achieve sparsity, in the context of the thesis of Antoine Gonon (defended on the 12/11/2024, [25]), we studied the expressivity of *quantized* deep networks from an approximation theoretic perspective [54]. Our objective was to define and compare the corresponding approximation classes [6] with the unquantized ones. We also characterized the error of nearest-neighbour uniform quantization of ReLU networks and we investigated when ReLU networks can be expected, or not, to have better approximation properties than other classical approximation families.

Quantization of neural networks: algorithms From a more computational perspective, and as a first step towards a better understanding of nonlinear quantized networks, we studied the simpler linear case. Particularly, we investigated the problem of optimally quantizing low rank matrices by exploiting scaling invariances inherent to the optimization problem. We proposed [58, 59] an optimal solution algorithm with polynomial complexity in the dimension of the problem and exponential complexity in the number of bits. We showed that it provides much more accurate quantizations than the simple round to nearest strategy. Particularly we used this algorithm in combination with the hierarchical procedure in [71], to design an heuristic strategy to efficiently quantize the family of butterfly matrices, which very often occur in machine learning applications, for instance to sparsify dense neural networks. Our work may help to improve the compression rate in this context by coupling sparsification and quantization.

In order to further exploit the benefits of quantization in neural networks and the multiple reduced numerical formats made available by modern computer architectures, we studied the introduction of mixed precision in the inference of neural networks. We proposed an analysis on the propagation of the error in the forward pass of neural networks, which suggests a good rule to choose the numerical format of each line of the weight matrices, yielding a mixed-precision procedure that provides the same accuracy of classical inference but with a lower energy consumption. This work makes the object of a preprint that is in preparation.

8.2.3 Sparse regularization, unfolding, and approximation theory

Participants: Marion Foare.

Collaborations with Nelly Pustelnik (Physics lab, ENS de Lyon).

In the PhD work of Hoang Trieu Vy Le, we investigated several unfolding strategies of standard proximal algorithms and their associated accelerated version in the context of image denoising, deconvolution. The goal was to study the impact of accelerated schemes on learning performance and robustness. Currently, we are studying various unrolling approaches to tackle the joint task of image restoration and edge detection. First, we proposed a two-step procedure mimicking the Blake-Zisserman minimization strategy, and relying on a smoothing Proximal Neural Network, followed by an edge detection layer ([67]). On the other hand, we are working on the unrolling procedure of the Mumford-Shah model.

8.2.4 Deep sparsity: from hardness to deformable butterfly algorithms

Participants: Rémi Gribonval, Elisa Riccietti, Pascal Carrivain, Léon Zheng.

Collaboration with Patrick Perez and Gilles Puy (Valeo AI, Paris), Quoc-Tung Le (TSE, Toulouse)

Matrix factorization with sparsity constraints plays an important role in many machine learning and signal processing problems such as dictionary learning, data visualization, dimension reduction.

We have deeply investigated this subject in the last years in the context of the thesis of Quoc-Tung Le [68] and Léon Zheng [26].

Building on this series of work on the hardness, tractability, and uniqueness properties of sparse matrix factorizations under various sparsity constraints [90, 70, 71], we have proposed a white paper for the signal processing magazine (SPM) Special Issue "Mathematics of Deep Learning" that has been accepted to be published in an extended version. In this paper we will propose an overview on the role of sparsity in a deep learning context.

Notably this work will include our previous results on the subject.

First of all, this will include the extension of the tractable algorithm for so-called butterfly sparsity patterns (which somehow factorizes a given matrix essentially at the cost of a single matrix-vector multiplication, with exact recovery guarantees) to so-called *deformable butterflies*. We have studied its performance guarantees beyond the case of matrices admitting an exact factorization and this is the object of a paper submitted to SIMAX. The corresponding algorithm has been incorporated in the FA μ ST software library (see Section 7).

Second, we will include also our study on the understanding on how to fully exploit the specific structure of butterfly factors and translate it into practical time gains. Specifically, we have studied how to optimize memory access to the matrix elements and we implemented a CUDA kernel to multiply on GPU a dense matrix with a deformable butterfly factor. This is also available in FA μ ST. We are currently working to improve the paper associated to the algorithm for submission to an international conference in 2025. In the paper we benchmark our implementation against existing multiplication algorithms to select the optimal one based on the matrix multiplication settings.

Going beyond the linear case, we will include in the white paper also our results on neural networks. We have indeed shown that the pitfalls that we had identified for certain sparse matrix factorization problems [71] also hold for certain sparse ReLU neural network training problems [69]. In particular, there exist settings where the optimization is necessarily unstable, in the sense that minimizing the loss function can only be achieved by letting some coefficients diverge to infinity.

Finally, we will mention also our developed heuristics to handle butterfly approximations for matrices under unknown permutations of rows and/or columns [89].

8.2.5 Plug and play methods and generative modelling

Participants: Anne Gagneux.

Collaboration with Emmanuel Soubies (CNRS, IRIT), Ségolène Martin, Paul Hagemann, Gabriele Steidl (TU Berlin).

In the PhD work of Anne Gagneux, we are investigating use of neural networks to implement convex functions. Learning convex functions has many applications in imaging (notably in Plug and Play methods) and in optimal transport. In [27] we have investigated the expressive power of Input Convex Neural Networks (ICNNs), a special architectural constraint. In particular, we have shown that ICNNs are restrictive, and may require more neurons than unconstrained networks to implement a given convex function.

In imaging tasks, Plug and Play (PnP) methods leverage the strength of pre-trained denoisers, often deep neural networks, by integrating them in optimization schemes. While they achieve state-of-the-art performance on various inverse problems in imaging, PnP approaches face inherent limitations on more generative tasks like inpainting. On the other hand, generative models such as Flow Matching pushed the boundary in image sampling yet lack a clear method for efficient use in image restoration. In [35], we have proposed to combine the PnP framework with Flow Matching (FM) by defining a time-dependent denoiser using a pre-trained FM model. Our algorithm alternates between gradient descent steps on the data-fidelity term, reprojections onto the learned FM path, and denoising. On tasks such as denoising, super-resolution, deblurring, and inpainting, our algorithm demonstrates superior results compared to existing PnP algorithms and Flow Matching based state-of-the-art methods.

8.3 Statistical learning, dimension reduction, and privacy preservation

8.3.1 Theoretical foundations of compressive learning: sketches, kernels, and optimal transport

Participants: Rémi Gribonval, Titouan Vayer, Paulo Goncalves, Etienne Lassalle.

Collaboration with Ayoub Belhadji (MIT)

The compressive learning framework proposes to deal with the large scale of datasets by compressing them into a single vector of generalized random moments, called a *sketch*, from which the learning task is then performed. In past works we established statistical guarantees on the generalization error of this procedure, first in a general abstract setting illustrated on PCA [4], then for the specific case of compressive k -means and compressive Gaussian Mixture Modeling [57]. The overall framework is described in a tutorial paper [5].

Theoretical guarantees in compressive learning fundamentally rely on comparing certain metrics between probability distributions as explored in a paper from last year [11]. Compressive learning also exploits the ability to approximate certain kernels by finite dimensional quadratures. We revisited in [14] existing proofs of the Restricted Isometry Property of sketching operators with respect to certain mixtures models. We proposed an alternative analysis that circumvents the need to assume importance sampling when drawing random Fourier features to build random sketching operators. Our analysis is based on new deterministic bounds on the restricted isometry constant that depend solely on the set of frequencies used to define the sketching operator. This analysis opens the door to theoretical guarantees for structured sketching with frequencies associated to fast random linear operators.

Finally, we demonstrated last year in [84] how sketching techniques can be employed to estimate the precision matrix used in the Graphical Lasso algorithm. The central advantage lies in providing a graph estimation method with a limited amount of data compared to standard methods. This year's work focused on exploring efficient algorithms for the decoding problem and establishing theoretical guarantees for structured operators.

8.3.2 Practical exploration of sketching and methods with limited resources

Participants: Etienne Lassalle, Rémi Gribonval, Titouan Vayer, Paulo Goncalves.

Collaborations with Rémi Vaudaine (previously postdoctoral researcher), Ayoub Belhadji (MIT), Marton Karsai (CEU, Vienne, Austria) and Pierre Borgnat (Physics Lab, ENS deLyon)

From a more empirical perspective, we pursued our efforts to make sketching for compressive learning and sketching more versatile and efficient. This notably involved investigating improved algorithms to learn from a sketch [15]. In the context of compressive clustering, the standard heuristic is CL-OMPR, a variant of sliding Frank-Wolfe. We showed how this algorithm can fail to recover clusters even in advantageous scenarios, and showed how its deficiencies can be attributed to optimization difficulties related to the structure of a correlation function appearing at core steps of the algorithm. To address these limitations, we propose an alternative decoder offering substantial improvements over CL-OMPR. Its design was notably inspired from the mean shift algorithm, a classic approach to detect the local maxima of kernel density estimators. The proposed algorithm can extract clustering information from a sketch of the MNIST dataset that is 10 times smaller than previously.

Sketching was also explored for temporal network compression [18]. In the context of temporal networks, which can model spreading processes such as epidemics, the out-component of a source node is the set of nodes reachable from this node, and the distribution of the size of out-components is an important characteristics which computation can be demanding for large networks. We proposed both an exact online matrix algorithm with controlled complexity footprint to compute this distribution, and a sketching-based framework to estimate it from a highly compressed representation of the temporal network.

Finally, we explore the sketching approach in the context of graph clustering, a key task in graph analysis. Many methods, like spectral clustering, are impractical for large graphs due to computational constraints. To address this, we introduced PASCO in [31], a sketching-based overlay that accelerates clustering algorithms. PASCO involves: 1- generating small, structure-preserving coarse graphs from the input graph, 2- running clustering algorithms in parallel on these graphs to produce partitions, and 3- aligning and merging these partitions using optimal transport. The PASCO framework is based on two key contributions: a novel global algorithm structure designed to enable parallelization and a fast, empirically validated graph coarsening algorithm that preserves structural properties. This work was submitted to ECML-PKDD 2025.

8.3.3 Dimensionality reduction and optimal transport

Participants: Titouan Vayer, Mathurin Massias.

Collaborations with Hugues Van Assel (PhD student, ENS Lyon), Cédric Vincent-Cuaz (post-doctoral researcher, EPFL), Rémi Flamary (CMAP, Ecole Polytechnique), Nicolas Courty (IRISA, Université Bretagne Sud), Pascal Frossard (EPFL).

Exploring and analyzing high-dimensional data is a core problem of data science that requires building low-dimensional and interpretable representations of the data through dimensionality reduction (DR). In a series of work we provide new methods an analysis for DR, inspired from optimal transport (OT). A key requirement for dimensionality reduction is to incorporate global dependencies among original and embedded samples while preserving clusters in the embedding space. In a previous work [82], we introduced and explored an innovative nonlinear dimensionality reduction method by utilizing the optimal transport framework and entropic affinities.

Building on these results, we extended our work to generalize dimension reduction, as detailed in [83]. Our approach leverages OT, specifically the Gromov-Wasserstein distance (GW), to propose a framework that simultaneously reduces both the dimensionality and the number of points in a dataset, enabling significant data compression. Notably, when the number of points is preserved, we demonstrated strong connections between our method and traditional dimensionality reduction techniques, such as spectral methods and t-SNE. We refer to our framework as "Distributional Dimension Reduction" which can be interpreted as projecting a distribution, and a geometry encoding the relationships among

data points in high-dimensional space, into a lower-dimensional space using the GW perspective. The corresponding article was submitted to AISTATS 2025. Based on these principles, we developed a library for dimensionality reduction in Pytorch [7.1.5](#).

8.4 Large-scale convex and nonconvex optimization

8.4.1 Multilevel schemes for image restoration

Participants: Elisa Riccietti, Paulo Gonçalves, Guillaume Lauga, Edgar Desainte-Mareville

Collaboration with Nelly Pustelnik (CNRS, ENS de Lyon), Nils Laurent (ENS de Lyon)

In the context of the Ph.D. work of Guillaume Lauga (defended on the 18/12/2024, we concluded our the study on the combination of proximal methods and multiresolution analysis in large-scale image denoising problems. In the spirit of multilevel gradient methods [\[44\]](#) we developed a family of multilevel inexact inertial proximal methods, tailored for problems arising in imaging, which exploit wavelets-based transfer operators [\[17\]](#). Their ability to accelerate proximal algorithms was shown in several large dimensional problems [\[65\]](#) [\[66\]](#), and particularly in real-word problems arising in radio-interferometry [\[63\]](#) and involving hyperspectral images [\[64\]](#). We also studied the link between multilevel and block coordinate methods and their convergence analysis. This work is the object of a preprint that is in preparation. The work of Guillaume Lauga is pursued by Edgar Desainte-Mareville that in the context of his thesis has started to investigate how to unroll such multilevel strategies in order to learn important ingredients such as the transfer operators.

In the context of the postdoc of Nils Laurent, we pursued the investigation on the use of multilevel schemes in conjunction with plug and play (PnP) methods, with the aim of reducing their computational cost. As these methods involve neural networks, the strategy to integrate multilevel schemes is naturally different from the one used so far in classical image denoising problems. In particular we have proposed a multilevel initialisation strategy that allows to treat inpainting problems efficiently, while classical PnP methods fail to do that. This work makes the object of a preprint that is in preparation.

8.4.2 Subsampling methods for problems involving large datasets

Participants: Elisa Riccietti.

Collaboration with Margherita Porcelli (UNIFI, Firenze, Italy) and Filippo Marini (UNIBO, Bologna, Italy)

Training problems usually involve a large number of data and are thus typically solved by stochastic optimization methods. The choice of the batch size for such methods affects their convergence and their efficiency. In particular variance reduction techniques aim at making stochastic methods more stable and less dependent on this choice. The choice of the learning rate remains however difficult. We concluded the study started last year on a variant of stochastic variance reduction methods inspired by multilevel schemes based on an automatic choice of the learning rate [\[73\]](#). Our method turns out to be competitive with standard variance reduction techniques on convex problems and to greatly outperform them on non-convex problems.

8.4.3 Reproducible benchmarking of optimization algorithms

Participants: Mathurin Massias.

Collaboration with Thomas Moreau (MIND, Inria Saclay), Badr Moufad (Ecole Polytechnique), Nelly Pustelnik (CNRS, ENS de Lyon).

The team continues working on reproducible optimisation benchmarks, with Benchopt [74], a collaborative framework to automate, reproduce and publish optimization benchmarks in machine learning across programming languages and hardware architectures. We continued to publish open source implementations of state-of-the-art solvers on major ML problems, and a detailed comparison of the regimes in which they succeed and fail respectively. We have organized a yearly sprint in June, and are working on new applications of the framework to imaging tasks in collaboration with Nelly Pustelnik.

8.4.4 Algorithms for large scale sparse linear models

Participants: Mathurin Massias.

Collaboration with Quentin Bertrand (INRIA MALICE), Badr Moufad (Ecole Polytechnique)

Based on our seminal works in [8] and [2], we continued to develop and implement new state-of-the-art solvers for optimization problems with millions of variables in the context of sparse linear models [42], implemented in the `skglm` package (see Section 7.1.2), that was integrated into the ecosystem of the `scikit-learn` package.

9 Bilateral contracts and grants with industry

9.1 Bilateral grants with industry

- CIFRE contract with CNES, Paris on "Optimized on-board decision with fast energy-efficient neural networks". This PhD thesis is in collaboration with Stéphane May, engineer at CNES.

Participants: Rémi Gribonval, Titouan Vayer, Arthur Lebeurier.

Duration: 3 years (2024-2027)

Partners: CNES, Paris; ENS de Lyon

Funding: CNES, Paris; PEPR IA SHARP

Context: ANR Chaire IA AllegroAssai 10.1.2

This thesis aims to develop compact, high-performance neural networks tailored to on-board constraints, enabling optimized decision-making on low-energy platforms. It includes an exploration of parsimony structures suited for deep networks and a comprehensive study of quantization and optimization techniques for neural networks.

- Funding from Facebook Artificial Intelligence Research, Paris

Participants: Rémi Gribonval.

Duration: 4 years (2021-2024)

Partners: Facebook Artificial Intelligence Research, Paris; ENS de Lyon

Funding: Facebook Artificial Intelligence Research, Paris

Context: Chaire IA AllegroAssai [10.1.2](#)

This is supporting the research conducted in the framework of the Chaire IA AllegroAssai.

10 Partnerships and cooperations

10.1 National initiatives

10.1.1 PEPR IA project : SHARP

Participants: Rémi Gribonval (*correspondant*), Paulo Gonçalves, Elisa Riccietti, Marion Foare, Mathurin Massias, Titouan Vayer, Arthur Lebourrier, Mael Chaumette.

Partnership with LAMSADE (PSL); LIGM (ENPC); GENESIS (Inria London & University College London); IRISA; CEA List; ISIR (Sorbonne Université)

Duration of the project: **2023 - 2029**.

The vision of the SHARP proposal is that the resources required to train ML models can be decreased by several orders of magnitude, with negligible performance loss compared to the state of the art. This means significantly reducing the dimensionality of predictors (to reduce inference costs) and of their gradients (to reduce training and bandwidth costs in distributed settings), the amount of data needed to learn (to address data scarce settings up to zero-shot learning, and incremental learning scenarios), and compressing datasets before learning (to reduce storage and compute requirements, and address privacy concerns).

10.1.2 ANR IA Chaire : AllegroAssai

Participants: Rémi Gribonval (*correspondant*), Paulo Gonçalves, Elisa Riccietti, Marion Foare, Mathurin Massias, Léon Zheng, Quoc-Tung Le, Antoine Gonon, Titouan Vayer, Ayoub Belhadji, Clement Lalanne, Can Pouliquen.

Past members: Luc Giffon.

Duration of the project: **2020 - 2025**.

AllegroAssai focuses on the design of machine learning techniques endowed both with statistical guarantees (to ensure their performance, fairness, privacy, etc.) and provable resource-efficiency (e.g. in terms of bytes and flops, which impact energy consumption and hardware costs), robustness in adversarial conditions for secure performance, and ability to leverage domain-specific models and expert knowledge. The vision of AllegroAssai is that the versatile notion of sparsity, together with sketching techniques using random features, are key in harnessing these fundamental tradeoffs. The first pillar of the project is to investigate sparsely connected deep networks, to understand the tradeoffs between the approximation capacity of a network architecture (ResNet, U-net, etc.) and its “trainability” with provably-good algorithms. A major endeavor is to design efficient regularizers promoting sparsely connected networks with provable robustness in adversarial settings. The second pillar revolves around the design and analysis of provably-good end-to-end sketching pipelines for versatile and resource-efficient large-scale learning, with controlled complexity driven by the structure of the data and that of the task rather than the dataset size.

10.1.3 ANR DataRedux

Participants: Paulo Gonçalves (*correspondant*), Rémi Gribonval, Marion Foare, Rémi Vaudaine.

Duration of the project: **February 2020 - January 2024 prolonged to June 2025.**

DataRedux puts forward an innovative framework to reduce networked data complexity while preserving its richness, by working at intermediate scales (“mesoscales”). Our objective is to reach a fundamental breakthrough in the theoretical understanding and representation of rich and complex networked datasets for use in predictive data-driven models. Our main novelty is to define network reduction techniques in relation with the dynamical processes occurring on the networks. To this aim, we will develop methods to go from data to information and knowledge at different scales in a human-accessible way by extracting structures from high-resolution, diverse and heterogeneous data. Our methodology will involve the identification of the most relevant subparts of time-resolved datasets while remapping the remaining parts of the system, the simultaneous structural-temporal representations of time-varying networks, the development of parsimonious data representations extracting meaningful structures at mesoscales (“mesostructures”), and the building of models of interactions that include mesostructures of various types. Our aim is to identify data aggregation methods at intermediate scales and new types of data representations in relation with dynamical processes, that carry the richness of information of the original data, while keeping their most relevant patterns for their manageable integration in data-driven numerical models for decision making and actionable insights.

10.1.4 ANR Darling

Participants: Paulo Gonçalves (*correspondant*), Rémi Gribonval, Marion Foare.

Duration of the project: **February 2020 - January 2024.**

This project meets the compelling demand of developing a unified framework for distributed knowledge extraction and learning from graph data streaming using in-network adaptive processing, and adjoining powerful recent mathematical tools to analyze and improve performances. The project draws on three major parallel directions of research: network diffusion, signal processing on graphs, and random matrix theory which DARLING aims at unifying into a holistic dynamic network processing framework. Signal processing on graphs has recently provided a comprehensive set of basic instruments allowing for signal on graph filtering or sampling, but it is limited to static signal models. Network diffusion on the opposite inherently assumes models of time varying graphs and signals, and has pursued the path of proposing and understanding the performance of distributed dynamic inference on graphs. Both areas are however limited by their assuming either deterministic graph or signal models, thereby entailing often inflexible and difficult-to-grasp theoretical results. Random matrix theory for random graph inference has taken a parallel road in explicitly studying the performance, thereby drawing limitations and providing directions of improvement, of graph-based algorithms (e.g., spectral clustering methods). The ambition of DARLING lies in the development of network diffusion-type algorithms anchored in the graph signal processing lore, rather than heuristics, which shall systematically be analyzed and improved through random matrix analysis on elementary graph models. We believe that this original communion of as yet remote areas has the potential to path the pave to the emergence of the critically needed future field of dynamical network signal processing.

10.1.5 ANR JCJC MASSILIA

Participants: Titouan Vayer.

Duration of the project: **December 2021 - December 2025.**

Collaboration with Arnaud Breloy (PI of the project, Univ. Paris Nanterre), Florent Bouchard (CentraleSupélec), Cédric Richard (Univ. Côte d'Azur), Rémi Flamary (Ecole Polytechnique) and Ammar Mian (Univ. Savoie Mont Blanc)

This project aims at tackling current problems related to graph learning and its applications in a unified way centered around the spectral decomposition of the graph Laplacian and/or adjacency matrices. The central objective of this project is to model graph structures (distributions on spectral parameters) and leverage this formalism in to two main directions 1) improve graph learning processes by directly learning structured spectral decompositions from the data 2) handle collections of graphs in order to compute structured graphs barycenters, compress graphs representations, and classify/cluster data using their graph as the main feature.

10.1.6 ANR JCJC Multisc-In

Participants: Marion Foare, Elisa Riccietti.

Collaboration with Nelly Pustelnik (PI of the project, ENS de Lyon), Laurent Condat (KAUST, Saudi Arabia), Luis Briceño-Arias (Univ. Técnica Federico Santa Maria, Chili)

Duration of the project: **October 2019 - March 2024.**

Interface detection is a challenging question in image processing, and more generally in graph processing, leading to a large panel of applications going from geophysics research to societal studies. The common point to these applications is the willingness to have an interface detection at a fine scale, in order to extract physical or societal parameters, from high resolution data. This project is devoted to original image processing tools relying both on optimization and multiresolution techniques in order to provide a new paradigm for the interface detection on large scale data.

10.1.7 ANR JCJC EROSION

Participants: Mathurin Massias.

Duration of the project: **December 2023 - December 2026.**

Collaboration with Emmanuel Soubies (PI of the project, CNRS, IRIT), Paul Escande (CR CNRS, I2M), Cédric Févotte (DR CNRS, IRIT), Henrique Goulart (MdC INP, IRIT) and Joseph Salmon (Prof. Université de Montpellier, IMAG)

The promise of EROSION is to push the frontiers of sparse and low-rank optimization by combining the strengths of exact relaxations and local optimization. More precisely, we propose to move away from the appealing convex relaxation requiring too strong assumptions to ensure the equivalence with the original problem. Instead, EROSION will address the following two research objectives. 1 : Deriving exact relaxations of ℓ_0 regression (= same global minimizers) which, although still non-convex, are more amenable to non-convex local optimization (e.g., less local minimizers, wider basins of attraction). 2 : Developing new local optimization strategies that exploit the nice properties of such exact relaxations so as to improve both the quality of reached local extrema and the convergence speed over existing solvers.

10.1.8 ANR JCJC MEPHISTO

Participants: Elisa Riccietti (*correspondant*).

Duration of the project: **November 2024 - November 2028.**

This project focuses on large scale optimization problems in signal processing and imaging. We consider a special class of such problem: those that admit a hierarchical structure. The aim of the project is to develop parsimonious methods for their solution by exploiting such underlying structure. We will focus on four different kinds of hierarchical structures: those arising from the geometry or physics of the problem (such as multiple resolutions in images or discretization of infinite dimensional problems); those that can be built by exploiting the analytical structure of some problems (training of neural networks, data-fitting problems); those that can be built exploiting the intrinsic structure of the algebraic tools involved (matrix, tensors, such as in matrix factorization problems); those that can be built exploiting multiple numerical formats (floating point numbers with reduced number of bits) .

The ambition of this project is thus to develop a large family of parsimonious multiresolution, multi-level and multiprecision algorithms that are not only efficient but that can also rely on solid mathematical foundations.

10.1.9 DI2A - Subvention Simone et Cino del Duca, Institut de France.

Participants: Elisa Riccietti, Marion Foare, Paulo Gonçalves.

Duration of the project: **December 2023 - December 2025.**

This project focuses on the physics-informed design of architectures and multiresolution deep learning techniques for large scale image restoration and data analysis for astronomy. With the term physics-informed design we refer to all the deep learning strategies in which the choice of the architecture, biases and activation functions of neural networks is guided by the underlying physics of data acquisition and/or from the optimization proximal schemes employed for the solution. From an application point of view, the project targets problems in astronomy and specifically the study of circumstellar environments through the instrument SPHERE/IRDIS. We aim to propose innovative reconstruction approaches partially supervised or even non supervised.

10.1.10 GDR ISIS project PROSSIMO

Participants: Mathurin Massias (*correspondant*), Rémi Gribonval, Anne Gagneux, Emmanuel Soubies.

Duration of the project: **September 2023 - September 2025.**

Composite optimisation problems are ubiquitous in machine learning, signal, and image processing. With the proximal algorithms used to solve them, they have met with great success in applications and have been extensively studied. More recently, so-called 'plug-and-play' (PNP) methods, inspired by proximal algorithms, propose new iterative algorithms in which the application of the proximal operator of the regulariser is replaced by a pre-existing denoiser or a learned operator. Their flexibility, however, complicates their theoretical analysis, because in the general case the operator does not have the interesting properties of proximal operators. In the PROSSIMO project, we propose to implement and study PNP operators via neural networks, while guaranteeing that these operators have the same properties as proximal operators. We aim at combining the flexibility of PNP methods with the rigorous theoretical guarantees of model-based methods. In addition to implementing such networks, we propose to study their approximation capacity: what classes of function can they approximate, and at what speed?

10.1.11 ANR TSIA BenchArk

Participants: Mathurin Massias (*correspondant*).

Duration of the project: **October 2024 - October 2028**.

Collaboration with Thomas Moreau, Gaël Varoquaux (INRIA Saclay) and Joseph Salmon (INRIA Montpellier).

Numerical evaluation of novel methods, a.k.a. benchmarking, is a pillar of the scientific method in machine learning. However, due to practical and statistical obstacles, the reproducibility of published results is currently insufficient: many details can invalidate numerical comparisons, from insufficient uncertainty quantification to improper methodology. In 2022, the **Benchopt** initiative provided an open source Python package together with a framework to seamlessly run, reuse, share and publish benchmarks in numerical optimization. The BenchArk project aims at bringing Benchopt to the whole machine learning community, making it a new standard in benchmarking by empowering researchers and practitioners with efficient and valid benchmarking methods. Our goal is to ensure reproducibility and consistency in model evaluation. We will federate the machine learning community to develop informative and statistically valid benchmarks, while providing methods to reduce identified hurdles in implementing such practices.

10.1.12 ANR SEIZURE

Participants: Paulo Gonçalves (*correspondant*), Can Pouliquen.

Duration of the project: **September 2024 - August 2028**

Collaboration with Carole Lartizien (PI of the project, CNRS, Insa de Lyon, CREATIS), Julien Jung (MD-PhD, Hospices Civils de Lyon, CRNL), Pierre Borgnat (CNRS, ENS de Lyon, Physics Lab).

“Seeing the Epileptogenic Zone through machine Learning on structural, functional and clinical neurological data”

This project deals with the multimodal detection and the characterisation of epileptic zones in neuroimaging and intracranial EEG (iEEG). Ockham is mainly involved in WP3 (P Borgnat leader) that aims at analysing the propagation of biomarkers within the brain as an indicator of the dynamic interictal epileptogenic network. A detailed understanding of the brain network and its key hubs provides invaluable insights into surgical outcomes. In a previous PhD work (G. Frusque, 2017-2020) we derived graphical lasso techniques on iEEG data to infer graphs times series, as relevant connectivity networks. In Seizure, we envision to enrich our previous approaches with deep learning based models and more specifically with graph recurrent neural networks.

11 Dissemination

Participants: Rémi Gribonval, Paulo Gonçalves, Marion Foare, Mathurin Massias, Elisa Riccietti, Titouan Vayer.

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair

- Organization of the COLT **2025 conference**, 30/06-04/07 2025

Member of the organizing committees

- Organization of [Journées SMAI-MODE](#), 27–29 Mar 2024, Lyon
- Organization of [Learning and Optimization in Luminy](#), 17-21 Jun 2024, CIRM, Marseille

Organization of the weekly "Machine Learning and Signal Processing (MLSP)" seminar (about thirty presentations in 2024 ([link to abstracts](#)))

11.1.2 Scientific events: selection

Member of the conference program committees

- Mathurin Massias: Area Chair for NeurIPS, ICML.
- Rémi Gribonval: Member of the [MIA'25 Program Committee](#)
- Rémi Gribonval: Member of the scientific board of [JRAF](#) (Journées de recherche en apprentissage frugal), Grenoble, Nov 21-22 2024

11.1.3 Journal

Member of the editorial boards

- Mathurin Massias: Associate editor for TMLR, Associate editor for *Computo* (French Statistical Society)
- Rémi Gribonval: Associate Editor for *Constructive Approximation* (Springer), Senior Area Editor for the *IEEE Signal Processing Magazine*

11.1.4 Invited talks

- Elisa Riccietti, [Journée POPILS](#), INSA Lyon, Feb 16th 2024
- Rémi Gribonval, [Workshop Frugalité en IA et en statistiques](#), Sorbonne Université, Paris, Oct 4th 2024
- Rémi Gribonval, [PolSys Seminar](#), LIP6, Paris, Sep 27th 2024
- Rémi Gribonval, [Workshop Learning and Optimization in Luminy](#), CIRM, Marseille, 17-21 June 2024
- Rémi Gribonval, [IEM Distinguished Lecturers Seminar](#), EPFL, May 24th 2024
- Rémi Gribonval, [Colloquium du MAP5](#), Paris, May 17th 2024
- Rémi Gribonval, [Math Machine Learning seminar MPI MIS + UCLA](#), online, May 16, 2024
- Rémi Gribonval, [Workshop on Sparsity and Singular Structures](#), Aachen, 26-29 February 2024
- Mathurin Massias, [TraDE-OPT final conference](#), Sestri Levante (Italy), 25-28 October 2024

11.1.5 Leadership within the scientific community

- Rémi Gribonval: member of the Scientific Committee of [RT MAIAGES](#) (formerly RT/GDR MIA)
- Rémi Gribonval: member of the Comité de Liaison SIGMA-SMAI
- Rémi Gribonval: member of the Cellule ERC of Inria, mentoring for ERC candidates in computer science and applied mathematics at the national Inria level

11.1.6 Scientific expertise

- Rémi Gribonval: member of the Scientific Advisory Board (vice-president) of the Acoustics Research Institute of the Austrian Academy of Sciences, and a member of the Commission Prospective of Institut de Mathématiques de Marseille
- Elisa Riccietti: member of the "Conseil Scientifique de la FIL", member of the "Commission formation Milyon" and member of the jury for the selection of a candidate for the position ATER CPES, ENS Lyon.

11.1.7 Research administration

- P. Gonçalves is member of the steering committee for the ShapeMed@Lyon consortium's Data for Health workshop
- Paulo Gonçalves was Deputy Scientific Director and is now Scientific Director of Inria Lyon, and member of the Inria Evaluation Committee.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Master :
 - Rémi Gribonval: Inverse problems and high dimension; Mathematical foundations of deep neural networks; Concentration of measure in probability and high-dimensional statistical learning; M2, ENS Lyon
 - Mathurin Massias: Python for Datascience (M1, Ecole Polytechnique/HEC); Optimisation (M1, ENS Lyon), Computational Optimal Transport for Machine and Deep Learning (M2, ENS Lyon), Fundamentals of Machine Learning (M1, ENS Lyon)
 - Titouan Vayer: Fundamentals of Machine Learning (M1, ENS Lyon), Machine Learning for Graphs and on Graphs (M2, ENS Lyon), Computational Optimal Transport for Machine and Deep Learning (M2, ENS Lyon)
 - Elisa Riccietti: Fundamentals of Machine Learning (M1, ENS Lyon), Optimization and approximation (M1, ENS Lyon), Harnessing inexactness in scientific computing (M2, ENS Lyon), tutor responsibility at ENS Lyon
- Engineer cycle (Bac+3 to Bac+5):
 - Paulo Gonçalves: Traitement du Signal (déterministe, aléatoire, numérique), Estimation statistique. 80 heures Eq. TD. CPE Lyon, France
 - Marion Foare: Traitement du Signal (déterministe, numérique, aléatoire), Traitement et analyse d'images, Optimisation, Compression, Projets. 280 heures Eq. TD. CPE Lyon, France
- Other formations: "Fondements et pratique du machine learning et du deep learning", CNRS formation for Dassault Systèmes, 3 x 3 days (18h) with Mathurin Massias, Titouan Vayer and Aurélien Garivier.

11.2.2 Supervision

All PhD students of the team are co-supervised by at least one team member. In addition, some team members are involved in the co-supervision of students hosted in other labs.

- Titouan Vayer: co-supervision of the PhD thesis of Hugues Van Assel with Aurélien Garivier (UMPA, ENS Lyon), co-supervision the the M2 internship of Thomas Bobille with Pierre Borgnat (CNRS, ENS de Lyon).

- Mathurin Massias: co-supervision of the M1 internship of Wassim Mazouz (EC Lyon) with Nelly Pustelnik (CNRS, ENS de Lyon)
- Elisa Riccietti: co-supervision of the PhD thesis of Filippo Marini with Margherita Porcelli (UNIBO, Italy)
- Rémi Gribonval: co-supervision of the Ph.D. of Sibylle Marcotte with Gabriel Peyré since 2022 (Center for Data Science, ENS Paris).

PhD defenses in OCKHAM in 2023:

- Léon Zheng
- Antoine Gonon
- Guillaume Lauga

11.2.3 Juries

Members of the OCKHAM team participated to the following juries:

- P. Gonçalves: HDR Rémi Emonet (Reviewer), PhD Victor Léger (examiner)

11.3 Popularization

11.3.1 Productions (articles, videos, podcasts, serious games, ...)

Mathurin Massias, Anne Gagneux (team members), Ségolène Martin (TU Berlin), Quentin Bertrand (Inria MALICE) and Rémi Emonet (Laboratoire Hubert Curien) wrote a blog post on flow matching for generative modelling that was accepted at the "blog post" track of the ICLR conference. The blog post contains an in-depth introduction to modern generative modelling for images, with a focus on normalizing flows and flow matching techniques. It is accompanied by high quality interactive visualizations and is visible [online](#) before being published on the ICLR website in April.

11.3.2 Participation in Live events

- Science Festival (*Fête de la Science*) at ENS Lyon: Rémi Gribonval participated to a workshop presenting the basics of time-frequency analysis to the general public, in cooperation with the Physics Lab of ENS de Lyon

12 Scientific production

12.1 Major publications

- [1] E. Bautista, P. Abry and P. Gonçalves. 'L^Y -PageRank for Semi-Supervised Learning'. In: *Applied Network Science* 4.57 (2019), pp. 1–20. DOI: [10.1007/s41109-019-0172-x](https://doi.org/10.1007/s41109-019-0172-x). URL: <https://inria.hal.science/hal-02063780> (cit. on p. 5).
- [2] Q. Bertrand, Q. Klopfenstein, M. Massias, M. Blondel, S. Vaiter, A. Gramfort and J. Salmon. 'Implicit differentiation for fast hyperparameter selection in non-smooth convex learning'. In: *Journal of Machine Learning Research* 23.149 (Apr. 2022), pp. 1–48. URL: <https://hal.science/hal-03228663> (cit. on p. 19).
- [3] M. Foare, N. Pustelnik and L. Condat. 'Semi-Linearized Proximal Alternating Minimization for a Discrete Mumford-Shah Model'. In: *IEEE Transactions on Image Processing* (7th Oct. 2019), pp. 1–13. DOI: [10.1109/TIP.2019.2944561](https://doi.org/10.1109/TIP.2019.2944561). URL: <https://hal.science/hal-01782346> (cit. on p. 4).
- [4] R. Gribonval, G. Blanchard, N. Keriven and Y. Traonmilin. 'Compressive Statistical Learning with Random Feature Moments'. In: *Mathematical Statistics and Learning* 3.2 (21st Aug. 2021), pp. 113–164. DOI: [10.4171/msl/20](https://doi.org/10.4171/msl/20). URL: <https://inria.hal.science/hal-01544609> (cit. on p. 16).

- [5] R. Gribonval, A. Chatalic, N. Keriven, V. Schellekens, L. Jacques and P. Schniter. ‘Sketching Data Sets for Large-Scale Learning: Keeping only what you need’. In: *IEEE Signal Processing Magazine* 38.5 (Sept. 2021), pp. 12–36. DOI: [10.1109/MSP.2021.3092574](https://doi.org/10.1109/MSP.2021.3092574). URL: <https://inria.hal.science/hal-03350599> (cit. on p. 16).
- [6] R. Gribonval, G. Kutyniok, M. Nielsen and F. Voigtlaender. ‘Approximation spaces of deep neural networks’. In: *Constructive Approximation*. special issue on "Deep Networks in Approximation Theory" (2021). DOI: [10.1007/s00365-021-09543-4](https://doi.org/10.1007/s00365-021-09543-4). URL: <https://inria.hal.science/hal-02117139> (cit. on p. 14).
- [7] G. Lauga, E. Riccietti, N. Pustelnik and P. Gonçalves. ‘IML FISTA: A Multilevel Framework for Inexact and Inertial Forward-Backward. Application to Image Restoration’. In: *SIAM Journal on Imaging Sciences* (1st June 2024). DOI: [10.1137/23M1582345](https://doi.org/10.1137/23M1582345). URL: <https://inria.hal.science/hal-04075814>.
- [8] M. Massias, S. Vaïter, A. Gramfort and J. Salmon. ‘Dual Extrapolation for Sparse Generalized Linear Models’. In: *Journal of Machine Learning Research* 21.234 (Oct. 2020), pp. 1–33. URL: <https://hal.science/hal-02263500> (cit. on p. 19).
- [9] T. Moreau, M. Massias, A. Gramfort, P. Ablin, P.-A. Bannier, B. Charlier, M. Dagr eou, T. Dupr e La Tour, G. Durif, C. F. Dantas, Q. Kloppenstein, J. Larsson, E. Lai, T. Lefort, B. Mal ezieux, B. Moufad, B. T. Nguyen, A. Rakotomamonjy, Z. Ramzi, J. Salmon and S. Vaïter. ‘Benchopt: Reproducible, efficient and collaborative optimization benchmarks’. In: *NeurIPS 2022 - 36th Conference on Neural Information Processing Systems*. New Orleans, United States, 28th Nov. 2022. URL: <https://hal.science/hal-03830604>.
- [10] B. Ricaud, P. Borgnat, N. Tremblay, P. Gonalves and P. Vandergheynst. ‘Fourier could be a Data Scientist: from Graph Fourier Transform to Signal Processing on Graphs’. In: *Comptes Rendus. Physique* (19th Sept. 2019), pp. 474–488. DOI: [10.1016/j.crhy.2019.08.003](https://doi.org/10.1016/j.crhy.2019.08.003). URL: <https://inria.hal.science/hal-02304584> (cit. on p. 5).
- [11] T. Vayer and R. Gribonval. ‘Controlling Wasserstein Distances by Kernel Norms with Application to Compressive Statistical Learning’. In: *Journal of Machine Learning Research* 24.149 (1st Apr. 2023), pp. 1–51. URL: <https://hal.science/hal-03461492> (cit. on p. 16).
- [12] C. Vincent-Cuaz, R. Flamary, M. Corneli, T. Vayer and N. Courty. ‘Semi-relaxed Gromov Wasserstein divergence with applications on graphs’. In: *ICLR 2022 - 10th International Conference on Learning Representations*. Virtual, France, 25th Apr. 2022, pp. 1–28. URL: <https://hal.science/hal-03832589> (cit. on p. 11).
- [13] C. Vincent-Cuaz, R. Flamary, M. Corneli, T. Vayer and N. Courty. ‘Template based Graph Neural Network with Optimal Transport Distances’. In: *NeurIPS 2022 – 36th Conference on Neural Information Processing Systems*. New Orleans, United States, 2022. URL: <https://hal.science/hal-03839517> (cit. on p. 11).

12.2 Publications of the year

International journals

- [14] A. Belhadji and R. Gribonval. ‘Revisiting RIP guarantees for sketching operators on mixture models’. In: *Journal of Machine Learning Research* 25.55 (2024), pp. 1–68. URL: <https://hal.science/hal-03872878> (cit. on p. 16).
- [15] A. Belhadji and R. Gribonval. ‘Sketch and shift: a robust decoder for compressive clustering’. In: *Transactions on Machine Learning Research Journal* (24th Apr. 2024). URL: <https://inria.hal.science/hal-04356977>. In press (cit. on p. 17).
- [16] S. Gratton, V. Mercier, E. Riccietti and P. Toint. ‘A block-coordinate approach of multi-level optimization with an application to physics-informed neural networks’. In: *Computational Optimization and Applications* 89.2 (6th Aug. 2024), pp. 385–417. DOI: [10.1007/s10589-024-00597-1](https://doi.org/10.1007/s10589-024-00597-1). URL: <https://hal.science/hal-04735477>.

- [17] G. Lauga, E. Riccietti, N. Pustelnik and P. Gonçalves. ‘IML FISTA: A Multilevel Framework for Inexact and Inertial Forward-Backward. Application to Image Restoration’. In: *SIAM Journal on Imaging Sciences* (1st June 2024). DOI: [10.1137/23M1582345](https://doi.org/10.1137/23M1582345). URL: <https://inria.hal.science/hal-04075814> (cit. on p. 18).
- [18] R. Vaudaine, P. Borgnat, P. Gonçalves, R. Gribonval and M. Karsai. ‘Temporal network compression via network hashing’. In: *Applied Network Science* 9.1 (17th Jan. 2024), p. 3. DOI: [10.1007/s10618-020-00683-y](https://doi.org/10.1007/s10618-020-00683-y). URL: <https://hal.science/hal-04228037> (cit. on p. 17).

International peer-reviewed conferences

- [19] A. Gonon, N. Brisebarre, E. Riccietti and R. Gribonval. ‘A path-norm toolkit for modern networks: consequences, promises and challenges’. In: *International Conference on Learning Representations*. International Conference on Learning Representations. Wien, Austria, 2024. URL: <https://hal.science/hal-04225201> (cit. on pp. 8, 13).
- [20] S. Marcotte, R. Gribonval and G. Peyré. ‘Keep the Momentum: Conservation Laws beyond Euclidean Gradient Flows’. In: *Forty-first International Conference on Machine Learning*. 41st International Conference on Machine Learning. Vienna, Austria, 21st May 2024. URL: <https://hal.science/hal-04590932> (cit. on p. 13).

Conferences without proceedings

- [21] L. Amador, M. Foare, O. Beuf, H. Ratiney and E. V. Reeth. ‘SUPER-RESOLUTION ISOTROPE POUR L’IRM : UNE APPROCHE VARIATIONNELLE’. In: Journées Recherche en Imagerie et Technologies pour la Santé (RITS). Aubière, France, 5th June 2024. URL: <https://hal.science/hal-04628336>.
- [22] P. Borgnat, R. Vaudaine, E. Lassale, P. Gonçalves, R. Gribonval and M. Karsai. ‘Coarsened Spectral Clustering’. In: NetSci 2024. Québec (CA), Canada, 16th June 2024. URL: <https://hal.science/hal-04732048>.
- [23] G. Lauga, A. Repetti, E. Riccietti, N. Pustelnik, P. Gonçalves and Y. Wiaux. ‘A multilevel framework for accelerating uSARA in radio-interferometric imaging’. In: European Signal Processing Conference (EUSIPCO). Lyon, France, 26th Aug. 2024. DOI: [10.23919/EUSIPCO63174.2024.10715263](https://doi.org/10.23919/EUSIPCO63174.2024.10715263). URL: <https://inria.hal.science/hal-04509285>.

Scientific book chapters

- [24] A. Brenon. ‘Encoding the Specificities of Encyclopedias’. In: *Structuring Lexical Data and Digitising Dictionaries*. Brill, 7th Nov. 2024, pp. 36–62. URL: <https://hal.science/hal-04806947>.

Doctoral dissertations and habilitation theses

- [25] A. Gonon. ‘Harnessing symmetries for modern deep learning challenges : a path-lifting perspective’. Ecole normale supérieure de lyon - ENS LYON, 12th Nov. 2024. URL: <https://theses.hal.science/tel-04784426> (cit. on pp. 13, 14).
- [26] L. Zheng. ‘Data frugality and computational efficiency in deep learning’. Ecole normale supérieure de lyon - ENS LYON, 29th May 2024. URL: <https://theses.hal.science/tel-04680306> (cit. on p. 15).

Reports & preprints

- [27] A. Gagneux, M. Massias, E. Soubies and R. Gribonval. *Convexity in ReLU Neural Networks: beyond ICNNs?* 9th Jan. 2025. URL: <https://hal.science/hal-04877619> (cit. on p. 16).
- [28] A. Gonon, N. Brisebarre, E. Riccietti and R. Gribonval. *A rescaling-invariant Lipschitz bound based on path-metrics for modern ReLU network parameterizations*. 11th Jan. 2025. URL: <https://hal.science/hal-04584311> (cit. on p. 13).

- [29] A. Gonon, L. Zheng, P. Carrivain and Q.-T. Le. *Fast inference with Kronecker-sparse matrices*. 3rd Nov. 2024. URL: <https://hal.science/hal-04584450>.
- [30] C. Lage, N. Pustelnik, J.-M. Arbona, B. Audit and R. Gribonval. *Identifying a piecewise affine signal from its nonlinear observation - application to DNA replication analysis*. 1st Apr. 2024. URL: <https://hal.science/hal-04528634>.
- [31] E. Lasalle, R. Vaudaine, T. Vayer, P. Borgnat, R. Gribonval, P. Gonçalves and M. Karsai. *PASCO (PArallel Structured COarsening): an overlay to speed up graph clustering algorithms*. 17th Dec. 2024. URL: <https://hal.science/hal-04837207> (cit. on p. 17).
- [32] H. T. V. Le, M. Foare, A. Repetti and N. Pustelnik. *Embedding Blake-Zisserman Regularization in Unfolded Proximal Neural Networks for Enhanced Edge Detection*. 7th Nov. 2024. URL: <https://hal.science/hal-04771534>.
- [33] Q.-T. Le, R. Gribonval, E. Riccietti and L. Zheng. *Butterfly factorization with error guarantees*. 2024. URL: <https://hal.science/hal-04763712>.
- [34] F. Marini, M. Porcelli and E. Riccietti. *A multilevel stochastic regularized first-order method with application to training*. 13th Jan. 2025. URL: <https://inria.hal.science/hal-04884049>.
- [35] S. Martin, A. Gagneux, P. Hagemann and G. Steidl. *PNP-FLOW: Plug-And-Play Image Restoration with Flow Matching*. 3rd Oct. 2024. URL: <https://hal.science/hal-04881033> (cit. on p. 16).
- [36] C. Pouliquen, M. Massias and T. Vayer. *Schur's Positive-Definite Network: Deep Learning in the SPD cone with structure*. 2024. URL: <https://hal.science/hal-04726325> (cit. on p. 12).
- [37] H. Van Assel, C. Vincent-Cuaz, N. Courty, R. Flamary, P. Frossard and T. Vayer. *Distributional Reduction: Unifying Dimensionality Reduction and Clustering with Gromov-Wasserstein*. 2024. DOI: 10.48550/arXiv.2402.02239. URL: <https://hal.science/hal-04726195>.

Software

- [38] [SW] A. Gonon, N. Brisebarre, E. Riccietti and R. Gribonval, *Code for reproducible research - A pathnorm toolkit for modern networks: consequences, promises and challenges* version 1.0.0, 11th Mar. 2024. LIC: BSD 3-Clause "New" or "Revised" License. HAL: [hal-04498597](https://hal.science/hal-04498597), URL: <https://hal.science/hal-04498597>, VCS: https://github.com/agonon/pathnorm_toolkit, SWHID: <https://hal.archives-ouvertes.fr/hal-04498597;visit=swh:1:snp:3b7c23b687511f3d2e4673d222c3ba96195bb004;anchor=swh:1:rel:7c123216ebb2018ea3290cbdb4cf4f4b8ddea964;path=/>.
- [39] [SW] Q.-T. Le, L. Zheng, R. Gribonval and E. Riccietti, *Code for reproducible research - "Butterfly factorization with error guarantees"* 4th Oct. 2024. LIC: BSD 3-Clause "New" or "Revised" License. HAL: [hal-04720713](https://inria.hal.science/hal-04720713), URL: <https://inria.hal.science/hal-04720713>, SWHID: <https://hal.archives-ouvertes.fr/hal-04720713;visit=swh:1:snp:21befd2becea13037f5829e23457593a1d0ba16f;anchor=swh:1:rel:40482acf97bf213f4ae48549e39b3450ee631392;path=/>.
- [40] [SW] L. Zheng, G. Puy, E. Riccietti, P. Pérez and R. Gribonval, *Code for reproducible research - Butterfly factorization by algorithmic identification of rank-one blocks* 15th May 2024. LIC: BSD 3-Clause "New" or "Revised" License. HAL: [hal-04576156](https://inria.hal.science/hal-04576156), URL: <https://inria.hal.science/hal-04576156>, VCS: <https://github.com/leonzheng2/permutation-butterfly>, SWHID: <https://hal.archives-ouvertes.fr/hal-04576156;visit=swh:1:snp:baaf5c1e6241789b5d3986b85175366cc976774a;anchor=swh:1:rel:1c64b94d969137e80d21e29e49499a583a26d5e5;path=/>.

12.3 Cited publications

- [41] H. H. Bauschke, P. L. Combettes et al. *Convex analysis and monotone operator theory in Hilbert spaces*. Vol. 408. Springer, 2011 (cit. on p. 4).

- [42] Q. Bertrand, Q. Klopfenstein, P.-A. Bannier, G. Gidel and M. Massias. ‘Beyond l1: Faster and better sparse models with skglm’. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 38950–38965 (cit. on p. 19).
- [43] H. Boche, R. Calderbank, G. Kutyniok and J. Vybiral. *Compressed Sensing and its Applications*. English. Ed. by H. Boche, R. Calderbank, G. Kutyniok and J. Vybiral. Vol. Series: Applied and Numerical Harmonic Analysis. MATHEON Workshop 2013. ISSN: 2296-5009. Cham: Birkhäuser, Cham, 2015. DOI: 10.1007/978-3-319-16042-9. URL: http://books.google.cz/books?id=6KoYCGAAQBAJ%5C&pg=PA340%5C&dq=intitle:Compressed+Sensing+and+its+Applications%5C&hl=%5C&cd=1%5C&source=gbs%5C_api (cit. on p. 4).
- [44] H. Calandra, S. Gratton, E. Riccietti and X. Vasseur. ‘On a multilevel Levenberg–Marquardt method for the training of artificial neural networks and its application to the solution of partial differential equations’. In: *Optimization Methods and Software* 0.0 (2020), pp. 1–26. DOI: 10.1080/10556788.2020.1775828. eprint: [\url{https://doi.org/10.1080/10556788.2020.1775828}](https://doi.org/10.1080/10556788.2020.1775828). URL: <https://doi.org/10.1080/10556788.2020.1775828> (cit. on p. 18).
- [45] Y. de Castro and F. Gamboa. ‘Exact Reconstruction using Beurling Minimal Extrapolation’. In: *arXiv.org* (Mar. 2011). arXiv: 1103.4951v2. URL: <http://arxiv.org/abs/1103.4951v2> (cit. on p. 4).
- [46] A. Chatalic, V. Schellekens, F. Houssiau, Y.-A. De Montjoye, L. Jacques and R. Gribonval. ‘Compressive Learning with Privacy Guarantees’. In: *Information and Inference* (2021). URL: <https://hal.inria.fr/hal-02496896> (cit. on pp. 5, 6).
- [47] P. L. Combettes and J.-C. Pesquet. ‘Proximal splitting methods in signal processing’. In: *Fixed-point algorithms for inverse problems in science and engineering*. Springer, 2011, pp. 185–212 (cit. on p. 4).
- [48] P. Di Lorenzo, P. Banelli, S. Barbarossa and S. Sardellitti. ‘Distributed Adaptive Learning of Graph Signals’. In: *IEEE Transaction on Signal Processing* 65.16 (2017) (cit. on p. 5).
- [49] P. M. Djuric and R. C. *Cooperative and Graph Signal Processing: Principle and Applications*. Academic Press, 2018 (cit. on p. 5).
- [50] M. Elad. *Sparse and Redundant Representations*. English. From Theory to Applications in Signal and Image Processing. Springer, 2010. URL: http://books.google.fr/books?id=d5b61JI9BvAC%5C&printsec=frontcover%5C&dq=sparse+and+redundant+representations%5C&hl=%5C&cd=1%5C&source=gbs%5C_api (cit. on p. 4).
- [51] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. New York, NY: Springer, 2013. DOI: 10.1007/978-0-8176-4948-7. URL: <http://link.springer.com/10.1007/978-0-8176-4948-7> (cit. on p. 4).
- [52] J. Friedman, T. Hastie and R. Tibshirani. ‘Sparse inverse covariance estimation with the graphical lasso’. In: *Biostatistics* 9.3 (2008), pp. 432–441 (cit. on p. 5).
- [53] B. Girault, P. Gonçalves and E. Fleury. ‘Translation on Graphs: An Isometric Shift Operator’. In: *IEEE Signal Processing Letters* 22.12 (Dec. 2015), pp. 2416–2420. DOI: 10.1109/LSP.2015.2488279. URL: <https://hal.inria.fr/hal-01221562> (cit. on p. 5).
- [54] A. Gonon, N. Brisebarre, R. Gribonval and E. Riccietti. ‘Approximation speed of quantized vs. unquantized ReLU neural networks and beyond’. In: *IEEE Transactions on Information Theory* 69.6 (June 2023), pp. 3960–3977. DOI: 10.1109/TIT.2023.3240360. URL: <https://hal.science/hal-03672166> (cit. on p. 14).
- [55] S. Gratton, V. Mercier, E. Riccietti and P. L. Toint. ‘A block-coordinate approach of multi-level optimization with an application to physics-informed neural networks’. In: *Computational Optimization and Applications* 89.2 (2024), pp. 385–417 (cit. on p. 11).
- [56] R. Gribonval, G. Blanchard, N. Keriven and Y. Traonmilin. ‘Compressive Statistical Learning with Random Feature Moments’. en. In: *Mathematical Statistics and Learning* (2021). URL: <https://hal.inria.fr/hal-01544609> (visited on 15/04/2021) (cit. on pp. 5, 6).

- [57] R. Gribonval, G. Blanchard, N. Keriven and Y. Traonmilin. ‘Statistical Learning Guarantees for Compressive Clustering and Compressive Mixture Modeling’. In: *Mathematical Statistics and Learning* 3.2 (Aug. 2021). This preprint results from a split and profound restructuring and improvements of <https://hal.inria.fr/hal-01544609v2> It is a companion paper to <https://hal.inria.fr/hal-01544609v3>, pp. 165–257. DOI: 10.4171/msl/21. URL: <https://hal.inria.fr/hal-02536818> (cit. on p. 16).
- [58] R. Gribonval, T. Mary and E. Riccietti. ‘Optimal quantization of rank-one matrices in floating-point arithmetic—with applications to butterfly factorizations’. working paper or preprint. June 2023. URL: <https://inria.hal.science/hal-04125381> (cit. on p. 14).
- [59] R. Gribonval, T. Mary and E. Riccietti. ‘Scaling is all you need: quantization of butterfly matrix products via optimal rank-one quantization’. In: *Actes du GRETSI 2023*. Actes du GRETSI 2023 2023-1193. Grenoble, France: GRETSI - Groupe de Recherche en Traitement du Signal et des Images, Aug. 2023, pp. 497–500. URL: <https://inria.hal.science/hal-04146706> (cit. on p. 14).
- [60] R. Jenatton, J.-Y. Audibert and F. Bach. ‘Structured Variable Selection with Sparsity-Inducing Norms’. English. In: *Journal of Machine Learning Research* 12 (2011). Publisher: Massachusetts Institute of Technology Press, pp. 2777–2824. URL: <http://hal.inria.fr/inria-00377732> (cit. on p. 4).
- [61] S. Kumar, J. Ying, J. V. de M. Cardoso and D. Palomar. ‘A unified Framework for Structured Graph Learning via Spectral Constraints’. In: *Journal of Machine Learning Research* 21 (2020), pp. 1–60 (cit. on p. 5).
- [62] J. Larsson, Q. Klopfenstein, M. Massias and J. Wallin. ‘Coordinate Descent for SLOPE’. In: *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*. Valencia, Spain, Apr. 2023. URL: <https://hal.science/hal-04211867> (cit. on p. 4).
- [63] G. Lauga, A. Repetti, E. Riccietti, N. Pustelnik, P. Gonçalves and Y. Wiaux. ‘A multilevel framework for accelerating uSARA in radio-interferometric imaging’. In: *European Signal Processing Conference (EUSIPCO)*. Lyon, France, Aug. 2024. DOI: 10.23919/EUSIPCO63174.2024.10715263. URL: <https://inria.hal.science/hal-04509285> (cit. on p. 18).
- [64] G. Lauga, E. Riccietti, N. Pustelnik and P. Gonçalves. ‘Méthodes multi-niveaux pour la restauration d’images hyperspectrales’. In: *Colloque GRETSI, September 2023* (2023) (cit. on p. 18).
- [65] G. Lauga, E. Riccietti, N. Pustelnik and P. Gonçalves. ‘Méthodes proximales multi-niveaux pour la restauration d’images’. In: *GRETSI’22 - 28ème Colloque Francophone de Traitement du Signal et des Images*. Nancy, France, Sept. 2022. URL: <https://inria.hal.science/hal-03739496> (cit. on p. 18).
- [66] G. Lauga, E. Riccietti, N. Pustelnik and P. Gonçalves. ‘Multilevel FISTA for image restoration’. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*. IEEE, Rhodes, Greece, June 2023. DOI: 10.48550/arXiv.2210.15940. URL: <https://hal.science/hal-03831180> (cit. on p. 18).
- [67] H. T. V. Le, M. Foare, A. Repetti and N. Pustelnik. ‘Embedding Blake-Zisserman Regularization in Unfolded Proximal Neural Networks for Enhanced Edge Detection’. 2024. URL: <https://hal.archives-ouvertes.fr/hal-04771534v1> (cit. on p. 15).
- [68] Q.-T. Le. ‘Algorithmic and theoretical aspects of sparse deep neural networks’. Theses. Ecole normale supérieure de lyon - ENS LYON, Dec. 2023. URL: <https://inria.hal.science/tel-04329531> (cit. on p. 15).
- [69] Q.-T. Le, E. Riccietti and R. Gribonval. ‘Does a sparse ReLU network training problem always admit an optimum?’ In: *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*. Advances in Neural Information Processing Systems 36 (NeurIPS 2023). New Orleans (Louisiane), United States, Dec. 2023. URL: <https://inria.hal.science/hal-04108849> (cit. on p. 15).
- [70] Q.-T. Le, E. Riccietti and R. Gribonval. ‘Spurious Valleys, NP-hardness, and Tractability of Sparse Matrix Factorization With Fixed Support’. In: *SIAM Journal on Matrix Analysis and Applications* (2022). URL: <https://hal.science/hal-03364668> (cit. on p. 15).

- [71] Q.-T. Le, L. Zheng, E. Riccietti and R. Gribonval. ‘Fast learning of fast transforms, with guarantees’. In: *ICASSP 2022 - IEEE International Conference on Acoustics, Speech and Signal Processing*. This paper is associated to code for reproducible research available at <https://hal.inria.fr/hal-03552956>. Singapore, Singapore, May 2022. DOI: [10.1109/ICASSP43922.2022.9747791](https://doi.org/10.1109/ICASSP43922.2022.9747791). URL: <https://inria.hal.science/hal-03438881> (cit. on pp. 14, 15).
- [72] S. Marcotte, R. Gribonval and G. Peyré. ‘Abide by the Law and Follow the Flow: Conservation Laws for Gradient Flows’. In: *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*. Advances in Neural Information Processing Systems 36 (NeurIPS 2023). New Orleans (Louisiane), United States, Dec. 2023. URL: <https://inria.hal.science/hal-04150576> (cit. on p. 13).
- [73] F. Marini, M. Porcelli and E. Riccietti. ‘A multilevel stochastic regularized first-order method with application to training’. In: *arXiv preprint arXiv:2412.11630* (2024) (cit. on p. 18).
- [74] T. Moreau, M. Massias, A. Gramfort, P. Ablin, P.-A. Bannier, B. Charlier, M. Dagr eou, T. Dupre la Tour, G. Durif, C. F. Dantas et al. ‘Benchopt: Reproducible, efficient and collaborative optimization benchmarks’. In: *Advances in Neural Information Processing Systems 35* (2022), pp. 25404–25421 (cit. on p. 19).
- [75] C. Pouliquen, P. Gonalves, M. Massias and T. Vayer. ‘Implicit Differentiation for Hyperparameter Tuning the Weighted Graphical Lasso’. In: *GRETSI 2023 - XXIX eme Colloque Francophone de Traitement du Signal et des Images*. Grenoble (France), France, Aug. 2023, pp. 1–4. URL: <https://hal.science/hal-04151796> (cit. on p. 11).
- [76] A. Rahimi and B. Recht. ‘Random features for large-scale kernel machines’. In: 2007 (cit. on p. 6).
- [77] F. Roosta-Khorasani and M. Mahoney. ‘Sub-sampled Newton methods’. In: *Math. Program.* 174 (2019), pp. 293–326. DOI: [10.1007/s10107-018-1346-5](https://doi.org/10.1007/s10107-018-1346-5) (cit. on p. 6).
- [78] D. Shuman, S. Narang, P. Frossard, A. Ortega and P. Vandergheynst. ‘The Emerging Field of Signal Processing on Graphs’. In: *IEEE Signal Processing Magazine* (May 2013), pp. 83–98 (cit. on p. 5).
- [79] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Sch olkopf and G. R. G. Lanckriet. ‘Hilbert Space Embeddings and Metrics on Probability Measures.’ In: *JMLR* 11 (2010), pp. 1517–1561. URL: <http://dblp.org/rec/journals/jmlr/SriperumbudurGFSL10> (cit. on p. 6).
- [80] P. Stock and R. Gribonval. ‘An Embedding of ReLU Networks and an Analysis of their Identifiability’. In: *Constructive Approximation* 57 (2023), pages 853–899. DOI: [10.1007/s00365-022-09578-1](https://doi.org/10.1007/s00365-022-09578-1). URL: <https://hal.science/hal-03292203> (cit. on p. 13).
- [81] I. Todic and P. Frossard. ‘Dictionary Learning’. In: *IEEE Signal Processing Magazine* 28.2 (), pp. 27–38. DOI: [10.1109/MSP.2010.939537](https://doi.org/10.1109/MSP.2010.939537). URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5714407> (cit. on p. 4).
- [82] H. Van Assel, T. Vayer, R. Flamary and N. Courty. ‘SNEkhorn: Dimension Reduction with Symmetric Entropic Affinities’. In: *Thirty-seventh Annual Conference on Neural Information Processing Systems (NeurIPS)*. NeurIPS 2023 conference paper. New Orleans, United States, Dec. 2023. URL: <https://hal.science/hal-04103326> (cit. on p. 17).
- [83] H. Van Assel, C. Vincent-Cuaz, T. Vayer, R. Flamary and N. Courty. ‘Interpolating between Clustering and Dimensionality Reduction with Gromov-Wasserstein’. In: *NeurIPS OTML Workshop*. New Orleans, United States, Dec. 2023. URL: <https://hal.science/hal-04356797> (cit. on p. 17).
- [84] T. Vayer, E. Lasalle, R. Gribonval and P. Gonalves. ‘Compressive Recovery of Sparse Precision Matrices’. working paper or preprint. Dec. 2023. DOI: [10.48550/arXiv.2311.04673](https://doi.org/10.48550/arXiv.2311.04673). URL: <https://hal.science/hal-04275341> (cit. on p. 16).
- [85] Y. Wang, Z. Jiang, X. Chen, P. Xu, Y. Zhao, Y. Lin and Z. Wang. ‘E2-train: Training state-of-the-art cnns with over 80% energy savings’. In: *Advances in Neural Information Processing Systems*. 2019, pp. 5138–5150 (cit. on p. 6).
- [86] G. Yang, T. Zhang, P. Kirichenko, J. Bai, A. G. Wilson and C. De Sa. ‘SWALP: Stochastic weight averaging in low precision training’. In: *International Conference on Machine Learning*. 2019, pp. 7015–7024 (cit. on p. 6).

-
- [87] Z. Yao, A. Gholami, S. Shen, K. Keutzer and M. W. Mahoney. ‘ADAHESIAN: An adaptive second order optimizer for machine learning’. In: *arXiv preprint arXiv:2006.00719* (2020) (cit. on p. 6).
- [88] J. Ying, J. V. de M. Cardoso and D. Palomar. ‘Nonconvex Sparse Graph Learning under Laplacian Constrained Graphical Model’. In: *34th Conference on Neural Information Processing Systems*. 2020 (cit. on p. 5).
- [89] L. Zheng, G. Puy, E. Riccietti, P. Pérez and R. Gribonval. ‘Factorisation butterfly par identification algorithmique de blocs de rang un’. In: *XXIXème Colloque Francophone de Traitement du Signal et des Images*. Grenoble, France, Aug. 2023. URL: <https://inria.hal.science/hal-04145743> (cit. on p. 15).
- [90] L. Zheng, E. Riccietti and R. Gribonval. ‘Efficient Identification of Butterfly Sparse Matrix Factorizations’. In: *SIAM Journal on Mathematics of Data Science* (2022). URL: <https://inria.hal.science/hal-03362626> (cit. on p. 15).