

# 2025 Activity Report

RESEARCH CENTRE: Inria Paris Centre

IN PARTNERSHIP WITH: Ecole normale supérieure de Paris

  
Project-Team

# ARGO

Learning, graphs and distributed optimization





## **Project-Team ARGO**

*Creation of the Project-Team: 2023 October 01*

Each year, Inria research teams publish an Activity Report presenting their work and results over the reporting period. These reports follow a common structure, with some optional sections depending on the specific team. They typically begin by outlining the overall objectives and research programme, including the main research themes, goals, and methodological approaches. They also describe the application domains targeted by the team, highlighting the scientific or societal contexts in which their work is situated. The reports then present the highlights of the year, covering major scientific achievements, software developments, or teaching contributions. When relevant, they include sections on software, platforms, and open data, detailing the tools developed and how they are shared. A substantial part is dedicated to new results, where scientific contributions are described in detail, often with subsections specifying participants and associated keywords. Finally, the Activity Report addresses funding, contracts, partnerships, and collaborations at various levels, from industrial agreements to international cooperations. It also covers dissemination and teaching activities, such as participation in scientific events, outreach, and supervision. The document concludes with a presentation of scientific production, including major publications and those produced during the year.

## Keywords

### Computer sciences and digital sciences

- A3.4. – Machine learning and statistics
- A3.5.1. – Analysis of large graphs
- A6.1.2. – Stochastic Modeling
- A6.2.3. – Probabilistic methods
- A6.2.6. – Optimization
- A7.1. – Algorithms
- A7.1.3. – Graph algorithms
- A8.1. – Discrete mathematics, combinatorics
- A8.2. – Optimization
- A8.7. – Graph theory
- A8.8. – Network science
- A8.9. – Performance evaluation
- A9.2. – Machine learning
- A9.2.3. – Reinforcement learning
- A9.2.8. – Deep learning
- A9.9. – Distributed AI, Multi-agent

### Other research topics and application domains

- B4. – Energy
- B7. – Transport and logistics
- B8.1. – Smart building/home
- B8.5. – Smart society
- B9.5.6. – Data science

## Contents

<b>Project-Team ARGO</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>5</b>
<b>2 Overall objectives</b>	<b>7</b>
<b>3 Research program</b>	<b>7</b>
3.1 Learning and algorithms on graphs . . . . .	7
3.2 Deep learning on structured data and new architectures for learning . . . . .	8
3.3 Distributed optimisation and control . . . . .	9
<b>4 Application domains</b>	<b>10</b>
<b>5 Highlights of the year</b>	<b>10</b>
5.1 Awards . . . . .	10
<b>6 Latest software developments, platforms, open data</b>	<b>10</b>
6.1 Latest software developments . . . . .	10
6.1.1 Farm2Python . . . . .	10
6.1.2 WFCRL . . . . .	11
6.1.3 COGNAC . . . . .	11
<b>7 New results</b>	<b>11</b>
7.1 Learning and algorithms on graphs . . . . .	11
7.1.1 High-dimensional statistical inference . . . . .	11
7.1.2 Graph algorithms and temporal graphs . . . . .	14
7.1.3 Stochastic matching and queueing networks . . . . .	16
7.2 Deep Learning on structured data and new architectures for learning . . . . .	16
7.3 Optimization and control . . . . .	17
7.3.1 Federated learning and decentralized optimization . . . . .	17
7.3.2 Markov decision processes and reinforcement learning . . . . .	18
7.3.3 Learning and control for energy networks . . . . .	19
7.4 Other results . . . . .	21
<b>8 Bilateral contracts and grants with industry</b>	<b>21</b>
8.1 Bilateral contracts with industry . . . . .	21
<b>9 Partnerships and cooperations</b>	<b>21</b>
9.1 International research visitors . . . . .	21
9.1.1 Visits of international scientists . . . . .	21
9.2 National initiatives . . . . .	22
9.2.1 Project REDEEM, PEPR IA . . . . .	22
9.2.2 Project AI-NRGY, PEPR TASE . . . . .	22
9.2.3 Challenge Inria FedMalin . . . . .	22
9.2.4 Challenge INRIA-EDF . . . . .	22
9.2.5 Challenge LLM4CODE . . . . .	22
9.2.6 Joint IFPEN - Inria laboratory . . . . .	23
9.2.7 GdR ROD . . . . .	23
9.3 Regional initiatives . . . . .	23
9.3.1 PRAIRIE-PSAI . . . . .	23

---

<b>10 Dissemination</b>	<b>23</b>
10.1 Promoting scientific activities	23
10.1.1 Invited talks	23
10.1.2 Scientific expertise	24
10.1.3 Research administration	24
10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach	24
10.2.1 Teaching	24
10.2.2 Supervision	25
10.2.3 Juries	26
10.3 Popularization	26
10.3.1 Productions (articles, videos, podcasts, serious games, ...)	26
<b>11 Scientific production</b>	<b>27</b>
11.1 Major publications	27
11.2 Publications of the year	28

# 1 Team members, visitors, external collaborators

## Research Scientists

- Ana Busic [Team leader, INRIA, Researcher, until Sep 2025, HDR]
- Ana Busic [Team leader, INRIA, Senior Researcher, from Oct 2025, HDR]
- Elisabetta Cornacchia [INRIA, Starting Research Position]
- Marc Lelarge [INRIA, Senior Researcher, from May 2025, HDR]
- Antoine Maillard [INRIA, Researcher]
- Laurent Massoulié [INRIA, Senior Researcher, from May 2025, HDR]
- Sean Meyn [INRIA, Chair, from May 2025 until Jul 2025]
- Kevin Scaman [INRIA, ISFP, until Jun 2025]
- Laurent Viennot [INRIA, Senior Researcher, HDR]

## Faculty Members

- Louise Budzynski [ENS PARIS, Associate Professor]
- Jean-Michel Fourneau [UVSQ, Professor, until Jun 2025, délégation, HDR]

## Post-Doctoral Fellows

- Ashok Krishnan Komalan Sindhu [INRIA, Post-Doctoral Fellow]
- Constantin Philippenko [INRIA, Post-Doctoral Fellow, until Aug 2025]
- Sushil Mahavir Varma [INRIA, Post-Doctoral Fellow, until Aug 2025]

## PhD Students

- Pierre Aguié [INRIA, from Oct 2025]
- Killian Bakong Epoune [INRIA]
- Julien Cardinal [INRIA, from Mar 2025]
- Jackson Carter [SORBONNE UNIVERSITE, from Oct 2025]
- Baptiste Corban [IFPEN]
- Romain Cosson [INRIA, until Aug 2025]
- Louann Coste [ENS PARIS, from Sep 2025]
- Andrea Vincenzo Dell'Abate [INRIA, from Oct 2025]
- Alexandra Kortchemski [INRIA, from Nov 2025]
- Jean Adrien Lagesse [ENS PARIS]
- Thomas Le Corre [ENS DE LYON, from Sep 2025]
- Thomas Le Corre [INRIA, until Aug 2025]
- Shu Li [INRIA]

- Antoine Lunven [MERCE, CIFRE, from Dec 2025]
- Jakob Maier [INRIA]
- Julien Moreau [Renault, CIFRE]
- David Robin [INRIA, until Jun 2025]
- Jules Sintès [INRIA]
- Martin Van Waerebeke [INRIA]
- Louis Vassaux [INRIA, from Sep 2025]
- Lucas Weber [DGA, until Feb 2025]

### **Interns and Apprentices**

- Pierre Aguié [INRIA, Intern, from Apr 2025 until Aug 2025]
- Julien Cardinal [INRIA, from Feb 2025 until Feb 2025]
- Justine Cauvi [INRIA, Intern, from Feb 2025 until Jul 2025]
- Giovanni Colage [INRIA, Intern, from Oct 2025]
- Andrea Vincenzo Dell'Abate [ENS PARIS, Intern, until Jun 2025]
- Alexandra Dinu [INRIA, Intern, from May 2025]
- Alexandra Kortchemski [INRIA, Intern, from May 2025 until Oct 2025]
- Antoine Lunven [INRIA, Intern, from Apr 2025 until Oct 2025]
- Maxime Muhlethaler [ENPC, until Sep 2025]
- Louis Vassaux [ENS PARIS, Intern, until Mar 2025]

### **Administrative Assistant**

- Marina Kovacic [INRIA]

### **Visiting Scientists**

- Alessio Giorlandino [SISSA, from Oct 2025]
- Seyoung Yun [KAIST, from May 2025 until Jul 2025]

### **External Collaborators**

- Pierre-Gabriel Berlureau [ENS PARIS, until Feb 2025]
- Sebastien Brasier [EPFL - Lausanne, from Jul 2025 until Aug 2025]
- Pierluigi Crescenzi [GSSI, L'Aquila]
- Jean-Michel Fourneau [UVSQ, from Jul 2025, HDR]

## 2 Overall objectives

The research activity of ARGO focuses on learning, optimization and control methods for graphs and networks. The challenges we aim to address are:

**Determine efficient polynomial-time algorithms for fundamental graph processing tasks** such as clustering and graph alignment; advance understanding of the "hard phase" for such graph problems which consists of problem instances with no known polynomial time algorithms while non-polynomial time algorithms are known to exist.

**Develop new deep learning architectures.** We envision to use graph theory and algorithms to better understand and improve neural networks, either by reducing their size or by enhancing their structure. Message passing is the dominant paradigm in Graphical Neural Networks (GNN) but has fundamental limitations due to its equivalence to the Weisfeiler-Lehman isomorphism test. We will investigate architectures breaking away from the message-passing schemes to develop more expressive GNN architectures.

**Develop distributed algorithms for federated learning,** achieving optimal performance for: supervised learning of a common model, under constraints of privacy and energy consumption; personalized learning of individual models; unsupervised learning of clustering and mixture models.

**Advance the theory of reinforcement learning:** by investigating convergence properties and connections with control theory. We also envision to develop new reinforcement learning algorithms for distributed systems.

ARGO is a spin-off of the DYOGENE project-team. DYOGENE was created in 2013 jointly with Département d'Informatique de l'École Normale Supérieure (DIENS).

## 3 Research program

The research activity of ARGO is structured around three methodological axes:

### 3.1 Learning and algorithms on graphs

**Participants:** Louise Budzynski, Ana Bušić, Jean-Michel Fourneau, Marc Lelarge, Antoine Maillard, Laurent Massoulié, Laurent Viennot.

**Information-theoretic versus computational limits:** The two basic lines of inquiry in statistical inference have long been: (i) to determine fundamental statistical (i.e., information-theoretic) limits; and (ii) to find efficient algorithms achieving these limits. However, for many structured inference problems, it is not clear if statistical optimality is compatible with efficient computation. Statistically optimal estimators often entail an infeasible exhaustive search. Conversely, for many settings the computationally efficient algorithms we know are statistically suboptimal, requiring higher signal strength or more data than is information-theoretically necessary. This phenomenon suggests that the information-theoretic limit on the signal-to-noise ratio (or the amount of data) for these problems, as studied since the beginning of mathematical statistics, is not the practically relevant benchmark for modern high-dimensional settings. Instead, the practically relevant benchmark is the fundamental statistical limit for computationally efficient algorithms. By now dozens of fundamental high-dimensional statistical estimation problems are conjectured to have different computational and statistical limits. These problems (for example, sparse linear regression or sparse phase retrieval) are ubiquitous in practice and well-studied theoretically, yet the central mysteries remain: What are the fundamental data limits for computationally efficient algorithms? How do we find optimal efficient algorithms? At a more basic level, are these statistical-computational gaps in various problems appearing for a common reason? Is there hope of building a widely applicable theory describing and explaining statistical-computational trade-offs?

**Algorithmic approaches and their computational limits:** Of particular interest to us are specific problems of learning on graphs which arise in many situations and hence are strongly motivated by various application scenarios. Identification of new algorithms and characterization of their computational limits

can thus have important consequences on various application areas and at the same time advance our global understanding of the above-mentioned discrepancy between statistical and computational limits.

Two examples of graph inference problems important for the team are:

- **Graph clustering**, also known as **community detection**, with a plethora of applications from recommender systems to inference of protein function in protein-protein interaction networks.
- **Graph alignment** is an important generic task, relevant for social network data de-anonymization, registration of medical imaging data, automatic machine translation, . . .

Our long term objective is to jointly augment our algorithmic toolkit and at the same time deepen our understanding of the so-called hard phase that arises when computational and information-theoretic limits do not match. The algorithms that we shall develop and analyze will at least initially be versions or variations of: Message-passing methods akin to belief propagation; spectral methods; semi-definite programming; Monte-Carlo Markov chain simulation.

**Graph algorithms for neural networks:** Neural networks can be represented as weighted graphs. We propose to study their analysis and optimization through the lens of graph algorithms. An important aspect of neural network optimization concerns the reduction of their size to save resources in terms of memory and energy. In particular, a large body of literature concerns *pruning* techniques, that consist in updating a network through the removal of links. It appears that classical network architectures can often be greatly reduced in size while preserving accuracy with such techniques. However, starting from a rather large network seems necessary for training. Pruning is reminiscent of graph sparsification which consists in summarizing a graph by a smaller one while preserving similar properties. For example, a spanner is a subgraph where distances are preserved up to some stretch factor. Or similarly, a cut sparsifier is a sparser weighted graph such that the weights of cuts are approximately preserved. Such techniques could bring new pruning approaches in the context of neural networks with tunable approximation with respect to size reduction.

**Temporal graphs:** Time has an important role in a large number of practical graphs ranging from social to transportation networks. Most prominently, this concerns graphs that are the result of interactions of people or objects such as customer-product purchase graphs. However, their time aspect is often ignored for solving tasks such as community detection or recommendation. In its simplest form, a temporal graph is a multigraph where each edge is labeled with a date.

Over the past decades numerous works have revisited various graph problems and concepts in this temporal context. The most central one is temporal connectivity which is based on defining a temporal path as a path forming a sequence of edges with increasing date labels. Clustering appears as a natural problem to revisit in this context. There are also graphs where the date labels can be chosen while respecting some scheduling constraints. This is the case for public transport networks where edges along a bus trip have to be scheduled one after another but where the starting time of each bus can be tuned for global optimization purposes. This inverse problem of defining a temporal graph from a graph raises interesting structural questions. For example, characterizing the graphs that can be transformed into a fully temporally connected temporal graph while having a minimum number of edges was a challenging question raised by the problem of all-to-all gossiping. Other problems such as maximizing flows or finding temporal matchings could lead to new interesting structural characterizations.

Our current efforts in this area are outlined in Section 7.1.

### 3.2 Deep learning on structured data and new architectures for learning

**Participants:** Marc Lelarge, Kevin Scaman.

Typical machine learning algorithms operating on structured data require representations of the often symbolic data as numerical vectors. Vector representations of the data range from handcrafted feature

vectors via automatically constructed graph kernels to learned representations, either computed by dedicated embedding algorithms or implicitly computed by learning architectures like graph neural networks. The performance of machine learning methods crucially depends on the quality of the vector representations. Given the importance of the topic, there is surprisingly little theoretical work on vector embeddings, especially when it comes to representing structural information that goes beyond metric information (that is, distances in a graph). The objects we want to embed either are graphs, possibly labelled or weighted, or more generally relational structures, or they are nodes of a (presumably large) graph or more generally elements or tuples appearing in a relational structure. When we embed entire graphs or structures, we speak of graph embeddings or relational structure embeddings; when we embed only nodes or elements we speak of node embeddings. The key theoretical questions we will ask about vector embeddings of objects are the following:

**Expressivity:** Which properties of objects are represented by the embedding? What is the meaning of the induced distance measure? Are there geometric properties of the latent space that represent meaningful relations on the objects?

**Complexity:** What is the computational cost of computing the vector embedding? What are efficient embedding algorithms? How can we efficiently retrieve semantic information of the embedded data, for example, answer queries? or solve combinatorial problems?

**Geometric Deep Learning:** Recently, geometric deep learning emerged as an attempt for geometric unification of a broad class of machine learning problems from the perspectives of symmetry and invariance. These principles not only underlie the breakthrough performance of convolutional neural networks and the recent success of graph neural networks but also provide a principled way to construct new types of problem-specific inductive biases.

Deep learning has brought in the past decade a revolution in data science and made possible many tasks previously thought to be beyond reach — whether computer vision, speech recognition, natural language translation, or playing Go. But we now have a zoo of different neural network architectures for different kinds of data and it is difficult to understand the relations between different methods. Geometric Deep Learning serves two purposes: first, to provide a common mathematical framework and unifying principles to derive the most successful neural network architectures, and second, give a constructive procedure to build future architectures in a principled way.

We consider high-dimensional problems with an additional structure that comes from the geometry of the input signal and explore ways to incorporate this geometric structure into the learning algorithms.

**Optimization for Geometric Deep Learning:** Optimization plays a key role in the training of neural networks, and the success or failure of a given architecture is in a large part driven by its amenability to gradient-based optimization methods. In this respect, structured data such as graphs or spatio-temporal time series pose new challenges, as the complexity and structure of the dataset impacts the loss landscape of the training objective in non-trivial, and sometimes detrimental, ways. Graphs are indeed high-dimensional objects, and possess a large variety of characteristics, ranging from *local*—such as node degree, neighborhood sizes or number of triangles—to *global*—such as connectivity, the presence of clusters or communities—. Efficient neural network architectures should be able to correctly identify the impact of such characteristics on the desired output, and thus be sufficiently expressive to encode all these characteristics. As a result, the parameters of the deep learning architecture are supposed to encode these patterns, structures and invariances, and the optimization algorithm used during training should be able to detect them in the data. Understanding when gradient descent can or cannot find a given pattern or invariance in the dataset can thus help deep learning practitioners know which architectures to use in which circumstances, and alleviate some of the issues related to the lack of transparency of neural networks by better describing the patterns that a given architecture is able to learn.

Our objective will be to better understand the expressive capabilities of structured deep learning architectures such as graph neural networks by investigating the relationships between their structure and their optimization.

Our current efforts in this area are outlined in Section 7.2.

### 3.3 Distributed optimisation and control

**Participants:** Ana Bušić, Sean Meyn, Jean-Michel Fourneau, Kevin Scaman, Laurent Massoulié.

An important trend is to consider distributed optimization and control settings, fuelled among others, by the growth of cloud computing, the proliferation of mobile devices, and increasing integration of distributed energy resources in power grids.

**Federated Learning and Distributed Optimization:** Federated learning is relevant to situations where learning tasks are to be performed on a dataset that cannot be centralized at one location, either for reasons of storage/communication resources, or for privacy reasons.

Many distinct learning scenarios can be envisioned in this framework, featuring a variety of objectives and constraints. The team has obtained state-of-the-art results for the **supervised learning** scenario where all agents involved seek to train a common model on the union of their individual datasets (Scaman et al. 2017, Scaman et al. 2018). Besides supervised learning of a common model, we will also tackle so-called **personalized learning**, whereby individual agents seek a model tailored to their individual dataset, yet expect to improve their learning experience from collaboration between one another, and **unsupervised learning**.

**Reinforcement learning:** Along with the sharp increase in visibility of the field, the rate at which new reinforcement learning algorithms are being proposed is at a new peak. While the surge in activity is creating excitement and opportunities, there is a gap in understanding of basic principles that these algorithms need to satisfy for successful application. We will concentrate our efforts on (i) *algorithm design and convergence properties*, (ii) *exploiting partial knowledge of the system dynamics and/or optimal policy*, (iii) *connections between control and reinforcement learning*.

**Distributed control and multi-agent RL:** Motivated by the needs of the modern power networks, we investigate **distributed control** approaches for large populations of agents, using controlled Markov decision processes and mean-field control.

In **multi-agent reinforcement learning**, our focus is on algorithms that take into account specific network dependence structure between agents.

Our current efforts in this area are outlined in Section 7.3.

## 4 Application domains

Our main applications are social networks, energy networks, and large language model based code assistants.

## 5 Highlights of the year

### 5.1 Awards

Mathieu Even, *accessit prix de thèse Gilles Kahn*.

Best paper award at NetGCoop 2025 conference for the paper "Achieving a Collective Target Through Incentives" by Ashok Krishnan Komalan Sindhu, Helene Le Cadre and Ana Busic.

## 6 Latest software developments, platforms, open data

### 6.1 Latest software developments

#### 6.1.1 Farm2Python

**Name:** Interfacing tool for wind farm control research

**Keywords:** Wind farm, Python

**Functional Description:** Python interfacing tool to wind farm simulators. Makes evaluating and comparing wind farm control algorithms easier for engineers and researchers looking to optimize energy production.

**Contact:** Claire Bizon Monroc

### 6.1.2 WFCRL

**Name:** Wind Farm Control Reinforcement Learning

**Keywords:** Wind farm, Reinforcement learning, Benchmarking, Python

**Functional Description:** Python benchmark library to evaluate reinforcement learning algorithms on wind farm control. Makes evaluating and comparing wind farm control algorithms easier for engineers and researchers looking to optimize energy production.

**Contact:** Claire Bizon Monroc

### 6.1.3 COGNAC

**Name:** Cooperative Graph-based Networked Agent Challenges

**Keywords:** Reinforcement learning, Multi-agent, Multi-Agents System, Machine learning, Graph, Network simulator

**Functional Description:** COGNAC is a Python-based benchmark suite offering flexible, graph-structured, cooperative multi-agent environments for MARL research. The package offers standardized minimal implementations of several well-known theoretical graph-based MARL problems taken from the literature, adapted for empirical benchmarking with modern RL tooling.

**Release Contributions:** Initial released. Deployed on Pypi index.

**URL:** <https://github.com/yojul/cognac>

**Contact:** Jules Sintès

## 7 New results

**Participants:** All ARGO.

### 7.1 Learning and algorithms on graphs

#### 7.1.1 High-dimensional statistical inference

**Evidence of replica symmetry breaking under the Nishimori conditions in epidemic inference on graphs.**

In Bayesian inference, computing the posterior distribution from the data is typically a nontrivial problem, which usually requires approximations such as mean-field approaches or numerical methods, like the Monte Carlo Markov chain. Being a high-dimensional distribution over a set of correlated variables, the posterior distribution can undergo the notorious replica symmetry-breaking transition. When that happens, several mean-field methods and virtually every Monte Carlo scheme cannot provide a reasonable approximation to the posterior and its marginals. Replica symmetry is believed to be guaranteed whenever the data are generated with known prior and likelihood distributions, namely under the so-called Nishimori conditions. In [16], we break this belief by providing a counterexample showing that under the Nishimori conditions, replica symmetry breaking arises. Introducing a simple, geometrical model that can be thought of as a patient-zero retrieval problem in a highly infectious regime of the epidemic Susceptible-Infectious model,

we show that under the Nishimori conditions there is evidence of replica symmetry breaking. We achieve this result by computing the instability of the replica symmetric cavity method toward the one-step replica symmetry-broken phase. The origin of this phenomenon—replica symmetry breaking under the Nishimori conditions—is likely due to the correlated disorder appearing in the epidemic models.

**Interacting copies of random-constraint satisfaction problems.** In [14], we study a system of  $\gamma = 2$  coupled copies of a well-known constraint satisfaction problem (random hypergraph bicoloring) to examine how the ferromagnetic coupling between the copies affects the properties of the solution space. We solve the replicated model by applying the cavity method to the supervariables that take  $2^\gamma$  values. Our results show that a coupling of strength  $\gamma$  between the copies decreases the clustering threshold  $\alpha_d(\gamma)$ , at which typical solutions shatter into disconnected components, therefore preventing numerical methods such as Monte Carlo Markov chains from reaching equilibrium in polynomial time. This result needs to be reconciled with the observation that, in models with coupled copies, denser regions of the solution space should be more accessible. Additionally, we observe a change in the nature of the clustering phase transition, from discontinuous to continuous, in a wide  $\gamma$  range. We investigate how the coupling affects the behavior of the belief propagation (BP) algorithm on finite-size instances and find that BP convergence is significantly impacted by the continuous transition. These results highlight the importance of better understanding algorithmic performance at the clustering transition and call for further exploration into the optimal use of reweighting strategies designed to enhance algorithmic performance.

**Exact threshold for approximate ellipsoid fitting of random points.** In [15], we consider the problem (P) of exactly fitting an ellipsoid (centered at 0) to  $n$  standard Gaussian random vectors in  $\mathbf{R}^d$ , as  $n, d \rightarrow \infty$  with  $n/d^2 \rightarrow \alpha > 0$ . This problem is conjectured to undergo a sharp transition: with high probability, (P) has a solution if  $\alpha < 1/4$ , while (P) has no solutions if  $\alpha > 1/4$ . So far, only a trivial bound  $\alpha > 1/2$  is known to imply the absence of solutions, while the sharpest results on the positive side assume  $\alpha \leq \eta$  (for  $\eta > 0$  a small constant) to prove that (P) is solvable. In this work we show a universality property for the minimal fitting error achievable by ellipsoids: we show that, to leading order, it coincides with the minimal error in a so-called “Gaussian equivalent” problem, for which the satisfiability transition can be rigorously analyzed. Our main results follow from this finding, and they are twofold. On the positive side, we prove that if  $\alpha < 1/4$ , there exists an ellipsoid fitting all the points up to a small error, and that the lengths of its principal axes are bounded above and below. On the other hand, for  $\alpha > 1/4$ , we show that achieving small fitting error is not possible if the length of the ellipsoid’s shortest axis does not approach 0 as  $d \rightarrow \infty$  (and in particular there does not exist any ellipsoid fit whose shortest axis length is bounded away from 0 as  $d \rightarrow \infty$ ). To the best of our knowledge, our work is the first rigorous result characterizing the expected phase transition in ellipsoid fitting at  $\alpha = 1/4$ . In a companion non-rigorous work, the second author and D. Kunisky give a general analysis of ellipsoid fitting using the replica method of statistical physics, which inspired the present work.

**Bilinear Sequence Regression: A Model for Learning from Long Sequences of High-Dimensional Tokens.** Current progress in artificial intelligence is centered around so-called large language models that consist of neural networks processing long sequences of high-dimensional vectors called tokens. Statistical physics provides powerful tools to study the functioning of learning with neural networks and has played a recognized role in the development of modern machine learning. The statistical physics approach relies on simplified and analytically tractable models of data. However, simple tractable models for long sequences of high-dimensional tokens are largely underexplored. Inspired by the crucial role models such as the single-layer teacher-student perceptron (also known as generalized linear regression) played in the theory of fully connected neural networks, in [18], we introduce and study the (BSR) as one of the most basic models for sequences of tokens. We note that modern architectures naturally subsume the BSR model due to the skip connections. Building on recent methodological progress, we compute the Bayes-optimal generalization error for the model in the limit of long sequences of high-dimensional tokens and provide a message-passing algorithm that matches this performance. We quantify the improvement that optimal learning brings with respect to vectorizing the sequence of tokens and learning via simple linear regression. We also unveil surprising properties of the gradient descent algorithms in the BSR model.

**Average-Case Matrix Discrepancy: Satisfiability Bounds.** Given a sequence of symmetric matrices  $\{A_i\}_{i=1}^n$ , and a margin  $\gamma$ , we investigate whether it is possible to find signs such that the operator norm of the signed sum satisfies  $\| \sum_{i=1}^n \sigma_i A_i \| \leq \gamma$ . Kunisky and Zhang (2023) recently introduced a random version of this problem, where the matrices are drawn from the Gaussian orthogonal ensemble. This model can be seen as a random variant of the celebrated Matrix Spencer conjecture and as a matrix-valued analog of the symmetric binary perceptron (SBP) in statistical physics. In [20], we establish a satisfiability transition in this problem. Our main results are twofold. First, we prove that the expected number of solutions with margin  $\gamma$  has a sharp threshold at a critical  $\gamma_c$ : for the problem is typically unsatisfiable, while for the average number of solutions becomes exponentially large. Second, combining a second-moment method with recent results from Altschuler (2023) on margin concentration in perceptron-type problems, we identify a second threshold  $\gamma_2$ , such that for the problem admits solutions with high probability. In particular, we establish that a system of Gaussian random matrices can be balanced so that the spectrum of the resulting matrix macroscopically shrinks compared to the typical semicircle law. Finally, under a technical assumption, we show that there exist values of  $\gamma$  for which the number of solutions has large variance, implying the failure of the second-moment method and uncovering a richer picture than in the vector-analog SBP problem. Our proofs rely on concentration inequalities and large deviation properties for the law of correlated Gaussian matrices under spectral norm constraints.

**Injectivity of ReLU networks: Perspectives from statistical physics.** When can the input of a ReLU neural network be inferred from its output? In other words, when is the network injective? In [21], we consider a single layer,  $x \mapsto \text{ReLU}(Wx)$ , with a random Gaussian  $m \times n$  matrix  $W$ , in a high-dimensional setting where  $n, m \rightarrow \infty$ . Recent work connects this problem to spherical integral geometry giving rise to a conjectured sharp injectivity threshold for  $\alpha = m/n$  by studying the expected Euler characteristic of a certain random set. We adopt a different perspective and show that injectivity is equivalent to a property of the ground state of the spherical perceptron, an important spin glass model in statistical physics. By leveraging the (non-rigorous) replica symmetry-breaking theory, we derive analytical equations for the threshold whose solution is at odds with that from the Euler characteristic. Furthermore, we use Gordon’s min–max theorem to prove that a replica-symmetric upper bound refutes the Euler characteristic prediction. Along the way we aim to give a tutorial-style introduction to key ideas from statistical physics in an effort to make the exposition accessible to a broad audience. Our analysis establishes a connection between spin glasses and integral geometry but leaves open the problem of explaining the discrepancies.

**Fundamental Limits of Matrix Sensing: Exact Asymptotics, Universality, and Applications.** In the matrix sensing problem, one wishes to reconstruct a matrix from (possibly noisy) observations of its linear projections along given directions. In [38], we consider this model in the high-dimensional limit: while previous works on this model primarily focused on the recovery of low-rank matrices, we consider in this work more general classes of structured signal matrices with potentially large rank, e.g. a product of two matrices of sizes proportional to the dimension. We provide rigorous asymptotic equations characterizing the Bayes-optimal learning performance from a number of samples which is proportional to the number of entries in the matrix. Our proof is composed of three key ingredients: (i) we prove universality properties to handle structured sensing matrices, related to the “Gaussian equivalence” phenomenon in statistical learning, (ii) we provide a sharp characterization of Bayes-optimal learning in generalized linear models with Gaussian data and structured matrix priors, generalizing previously studied settings, and (iii) we leverage previous works on the problem of matrix denoising. The generality of our results allow for a variety of applications: notably, we mathematically establish predictions obtained via non-rigorous methods from statistical physics in Erba et al. (2024) regarding Bilinear Sequence Regression, a benchmark model for learning from sequences of tokens, and in Maillard et al. (2024) on Bayes-optimal learning in neural networks with quadratic activation function, and width proportional to the dimension.

**Bayes-optimal learning of an extensive-width neural network from quadratically many samples.** In [31], we consider the problem of learning a target function corresponding to a single hidden layer neural network, with a quadratic activation function after the first layer, and random weights. We consider the asymptotic limit where the input dimension and the network width are proportionally large. Recent work [Cui et al., 2023] established that linear regression provides Bayes-optimal test error to learn such a function

when the number of available samples is only linear in the dimension. That work stressed the open challenge of theoretically analyzing the optimal test error in the more interesting regime where the number of samples is quadratic in the dimension. In this paper, we solve this challenge for quadratic activations and derive a closed-form expression for the Bayes-optimal test error. We also provide an algorithm, that we call GAMP-RIE, which combines approximate message passing with rotationally invariant matrix denoising, and that asymptotically achieves the optimal performance. Technically, our result is enabled by establishing a link with recent works on optimal denoising of extensive-rank matrices and on the ellipsoid fitting problem. We further show empirically that, in the absence of noise, randomly-initialized gradient descent seems to sample the space of weights, leading to zero training loss, and averaging over initialization leads to a test error equal to the Bayes-optimal one

**Graph Alignment via Birkhoff Relaxation.** We consider the graph alignment problem, wherein the objective is to find a vertex correspondence between two graphs that maximizes the edge overlap. The graph alignment problem is an instance of the quadratic assignment problem (QAP), known to be NP-hard in the worst case even to approximately solve. In [35], we analyze Birkhoff relaxation, a tight convex relaxation of QAP, and present theoretical guarantees on its performance when the inputs follow the Gaussian Wigner Model. More specifically, the weighted adjacency matrices are correlated Gaussian Orthogonal Ensemble with correlation  $\frac{1}{1+\sigma^2}$ .

### 7.1.2 Graph algorithms and temporal graphs

**Unweighted Layered Graph Traversal: Passing a Crown via Entropy Maximization.** Introduced by Papadimitriou and Yannakakis in 1989, layered graph traversal is a central problem in online algorithms and mobile computing that has been studied for several decades, and which now is essentially resolved in its original formulation. In [22], we demonstrate that what appears to be an innocuous modification of the problem actually leads to a drastic (exponential) reduction of the competitive ratio. Specifically, we present an algorithm that is  $O(\log^2 w)$ -competitive for traversing unweighted layered graphs of width  $w$ . Our algorithm chooses the agent’s position simply according to the probability distribution over the current layer that maximizes the sum of entropies of the induced distributions in the preceding layers.

**Certificates in P and Subquadratic-Time Computation of Radius, Diameter, and all Eccentricities in Graphs.** In the context of fine-grained complexity, in [27] we investigate the notion of certificate enabling faster polynomial-time algorithms. We specifically target radius (minimum eccentricity), diameter (maximum eccentricity), and all-eccentricity computations for which quadratic-time lower bounds are known under plausible conjectures. In each case, we introduce a notion of certificate as a specific set of nodes from which appropriate bounds on all eccentricities can be derived in subquadratic time when this set has sublinear size. The existence of small certificates is a barrier against SETH-based lower bounds for these problems. We indeed prove that for graph classes with small certificates, there exist randomized subquadratic-time algorithms for computing the radius, the diameter, and all eccentricities respectively. Moreover, these notions of certificates are tightly related to algorithms probing the graph through one-to-all distance queries and allow to explain the efficiency of practical radius and diameter algorithms from the literature. Our formalization enables a novel primal-dual analysis of a classical approach for diameter computation that leads to algorithms for radius, diameter and all eccentricities with theoretical guarantees with respect to certain graph parameters. This is complemented by experimental results on various types of real-world graphs showing that these parameters appear to be low in practice. Finally, we obtain refined results for several graph classes.

**Forbidden Patterns in Temporal Graphs Resulting from Encounters in a Corridor.** In [17], we study temporal graphs arising from mobility models, where vertices correspond to agents moving in space and edges appear each time two agents meet. We propose a rather natural one-dimensional model. If each pair of agents meets exactly once, we get a simple temporal clique where the edges are ordered according to meeting times. In order to characterize which temporal cliques can be obtained as such ‘mobility graphs’, we introduce the notion of forbidden patterns in temporal graphs. Furthermore, using a classical result in combinatorics, we count the number of such mobility cliques for a given number of agents, and show that

not every temporal clique resulting from the 1D model can be realized with agents moving with different constant speeds. For the analogous circular problem, where agents are moving along a circle, we provide a characterization via circular forbidden patterns. Our characterization in terms of forbidden patterns can be extended to the case where each edge appears at most once. We also study the problem where pairs of agents are allowed to cross each other several times, using an approach from automata theory. We observe that in this case, there is no finite set of forbidden patterns that characterize such temporal graphs and nevertheless give a linear-time algorithm to recognize temporal graphs arising from this model.

**On the Complexity of Computing a Fastest Temporal Path in Interval Temporal Graphs.** Temporal graphs arise when modeling interactions that evolve over time. They usually come in several flavors, depending on the number of parameters used to describe the temporal aspects of the interactions: time of appearance, duration, delay of transmission. In the point model, edges appear at specific points in time, while in the more general interval model, edges can be present over multiple time intervals. In both models, the delay for traversing an edge can change with each edge appearance. When time is discrete, the two models are equivalent in the sense that the presence of an edge during an interval is equivalent to a sequence of point-in-time occurrences of the edge. However, this transformation can drastically change the size of the input and has complexity issues. Indeed, in [42], show a gap between the two models with respect to the complexity of the classical problem of computing a fastest temporal path from a source vertex to a target vertex, i.e. a path where edges can be traversed one after another in time and such that the total duration from source to target is minimized. It can be solved in near-linear time in the point model, while we show that the interval model requires quadratic time under classical assumptions of fine-grained complexity. With respect to linear time, our lower bound implies a factor of the number of vertices, while the best known algorithm has a factor of the number of underlying edges. Interestingly, we show that near-linear time is possible in the interval model when restricted to all delays being zero, i.e. traversing an edge is instantaneous.

**Foremost, Fastest, Shortest: Temporal Graph Realization under Various Path Metrics.** In [43], we follow the current trend on temporal graph realization, where one is given a property  $P$  and the goal is to determine whether there is a temporal graph, that is, a graph where the edge set changes over time, with property  $P$ . We consider the problems where as property  $P$ , we are given a prescribed matrix for the duration, length, or earliest arrival time of pairwise temporal paths. That is, we are given a matrix  $D$  and ask whether there is a temporal graph such that for any ordered pair of vertices  $(s, t)$ ,  $D_{s,t}$  equals the duration (length, or earliest arrival time, respectively) of any temporal path from  $s$  to  $t$  minimizing that specific temporal path metric. For shortest and earliest arrival temporal paths, we are the first to consider these problems as far as we know. We analyze these problems for many settings like: strict and non-strict paths, periodic and non-periodic temporal graphs, and limited number of labels per edge (that is, limited occurrence number per edge over time). In contrast to all other path metrics, we show that for the earliest arrival times, we can achieve polynomial-time algorithms in periodic and non-periodic temporal graphs and for strict and non-strict paths. However, the problem becomes NP-hard when the matrix does not contain a single integer but a set or range of possible allowed values. As we show, the problem can still be solved efficiently in this scenario, when the number of entries with more than one value is small, that is, we develop an FPT-algorithm for the number of such entries. For the setting of fastest paths, we achieve new hardness results that answers an open question by Klobas, Mertzios, Molter, and Spirakis [Theor. Comput. Sci. '25] about the parameterized complexity of the problem with respect to the vertex cover number and significantly improves over a previous hardness result for the feedback vertex set number. When considering shortest paths, we show that the periodic versions are polynomial-time solvable whereas the non-periodic versions become NP-hard.

**Parameterized Restless Temporal Path.** Recently, Bumpus and Meeks introduced a purely temporal parameter, called vertex-interval-membership-width, which is promising for the design of fixed-parameter tractable (FPT) algorithms for vertex reachability problems in temporal graphs. In [26], we study this newly introduced parameter for the problem of restless temporal paths, in which the waiting time at each node is restricted. In this article, we prove that, in the interval model, where arcs are present for entire time intervals, finding a restless temporal path is NP-hard even if the vertex-interval-membership-width is equal to

three. We exhibit FPT algorithms for the point model, where arcs are present at specific points in time, both with uniform delay one and arbitrary positive delays. In the latter case, this comes with a slight additional computational cost.

### 7.1.3 Stochastic matching and queuing networks

**Performance paradoxes in matching systems are not that rare.** A Matching model describes the waiting times suffered by items before they match with other items and disappear immediately without service. The matching relation is described by a compatibility graph. It is an easy representation of multiple types synchronizations between items. When we add an edge in the compatibility graph, one expects that the expected total number of items decreases. Unexpectedly this is not always true and there is a performance paradox. Extending previous results, in [24], we show new matching models with performance paradoxes which seems to be more frequent than one can expect.

## 7.2 Deep Learning on structured data and new architectures for learning

**Bootstrap learning for combinatorial graph alignment with sequential GNNs.** Graph neural networks (GNNs) have struggled to outperform traditional optimization methods on combinatorial problems, limiting their practical impact. In [46], we address this gap by introducing a novel chaining procedure for the graph alignment problem—a fundamental NP-hard task of finding optimal node correspondences between unlabeled graphs using only structural information. Our method trains a sequence of GNNs where each network learns to iteratively refine similarity matrices produced by previous networks. During inference, this creates a bootstrap effect: each GNN improves upon partial solutions by incorporating discrete ranking information about node alignment quality from prior iterations. We combine this with a powerful architecture that operates on node pairs rather than individual nodes, capturing global structural patterns essential for alignment that standard message-passing networks cannot represent. Extensive experiments on synthetic benchmarks demonstrate substantial improvements: our chained GNNs achieve over 3× better accuracy than existing methods on challenging instances, and uniquely solve regular graphs where all competing approaches fail. When combined with traditional optimization as post-processing, our method substantially outperforms state-of-the-art solvers on the graph alignment benchmark.

### **Optimal Adaptive Time-Frequency Representation via Differentiable Short-Time Fourier Transform.**

The short-time Fourier transform (STFT) is widely used for analyzing non-stationary signals. However, its performance is highly sensitive to its parameters, and manual or heuristic tuning often yields suboptimal results. To overcome this limitation, in [19], we propose a unified differentiable formulation of the STFT that enables gradient-based optimization of its parameters. This approach addresses the limitations of traditional STFT parameter tuning methods, which often rely on computationally intensive discrete searches. It enables fine-tuning of the time-frequency representation (TFR) based on any desired criterion. The efficacy of the proposed differentiable STFT in enhancing TFRs is demonstrated through experiments on both simulated and real-world data.

**Assimilation of Sparse Vehicle Trajectories with a Macroscopic Traffic Model.** Accurate macroscopic traffic prediction could lead to smarter vehicles and safer roads. It is however dependent on reliable and frequent measurements. Recent research in scientific machine learning has shown progress on reconstruction tasks using probe vehicles, hinting at potential gains for a sequential prediction model. In [40], we adopt the classical data assimilation framework and gather insight on how to design effective predictors. Our results highlight the effectiveness of macroscopic solvers given sparse data while confirming the flexibility of Kalman filtering. However, we report the difficulty of calibrating the solution given the low amount of measures and their high variance.

**Minif2f in Rocq: Automatic Translation Between Proof Assistants - A Case Study.** While the Minif2f dataset exists for Lean, Isabelle/HOL, and MetaMath, it has not been formalized in Rocq, limiting cross-system comparisons in automated theorem proving. In [41], we investigate whether state-of-the-art LLMs can automatically translate formal theorems between proof assistants. Using a three-stage methodology

from basic prompting to multi-turn conversations with error feedback, we successfully translated 478 out of 488 theorems (98%) from Minif2f to Rocq. Expert validation of 150 translations confirmed high accuracy, with only three errors. This work provides a complete Rocq formalization of Minif2f and demonstrates the viability of LLM-based cross-proof-assistant translation.

**Tacq -Context Aware Tactic Recommendation for Rocq.** Despite recent impressive achievements of Large Language Models (LLMs) in formal mathematics using proof assistants, the state-of-the-art has mostly focused on mathematics competitions where problems typically involve simple and well-understood concepts. This is unfortunately far from current practice in formal mathematics, where experts must navigate large libraries of lemmas and manipulate complex constructions.

In [36], we focus on the problem of next tactic recommendation for Rocq. A key issue is to prompt the model with enough context to understand the current goal when the proof relies on large libraries with numerous dependencies and where specialized notations are pervasive. We present a tool that can extract notations and dependencies from the current goal and annotate them with natural language docstrings. We show that such augmented contexts improve the ability of state-of-the-art models to generate valid tactics on the challenging MathComp library.

## 7.3 Optimization and control

### 7.3.1 Federated learning and decentralized optimization

**In-depth Analysis of Low-rank Matrix Factorisation in a Federated Setting.** In [32], we analyze a distributed algorithm to compute a low-rank matrix factorization on  $N$  clients, each holding a local dataset  $S_i \in \mathbb{R}^{n_i \times d}$ , mathematically, we seek to solve  $\min_{U_i \in \mathbb{R}^{n_i \times r}, V \in \mathbb{R}^{d \times r}} 1/2 \sum_{i=1}^N \|S_i - U_i V^T\|_F^2$ . Considering a power initialization of  $V$ , we rewrite the previous smooth non-convex problem into a smooth strongly-convex problem that we solve using a parallel Nesterov gradient descent potentially requiring a single step of communication at the initialization step. For any client  $i$  in  $\{1, \dots, N\}$ , we obtain a global  $V$  in  $\mathbb{R}^{d \times r}$  common to all clients and a local variable  $U_i$  in  $\mathbb{R}^{n_i \times r}$ . We provide a linear rate of convergence of the excess loss which depends on  $\sigma_{\max}/\sigma_r$ , where  $\sigma_r$  is the  $r^{\text{th}}$  singular value of the concatenation  $S$  of the matrices  $(S^i)_{i=1}^N$ . This result improves the rates of convergence given in the literature, which depend on  $\sigma_{\max}^2/\sigma_{\min}^2$ . We provide an upper bound on the Frobenius-norm error of reconstruction under the power initialization strategy. We complete our analysis with experiments on both synthetic and real data.

**Noiseless Privacy-Preserving Decentralized Learning.** Decentralized learning (DL) enables collaborative learning without a server and without training data leaving the users' devices. However, the models shared in DL can still be used to infer training data. Conventional defenses such as differential privacy and secure aggregation fall short in effectively safeguarding user privacy in DL, either sacrificing model utility or efficiency. In [23], we introduce Shatter, a novel DL approach in which nodes create virtual nodes (VNs) to disseminate chunks of their full model on their behalf. This enhances privacy by (i) preventing attackers from collecting full models from other nodes, and (ii) hiding the identity of the original node that produced a given model chunk. We theoretically prove the convergence of Shatter and provide a formal analysis demonstrating how Shatter reduces the efficacy of attacks compared to when exchanging full models between nodes. We evaluate the convergence and attack resilience of Shatter with existing DL algorithms, with heterogeneous datasets, and against three standard privacy attacks. Our evaluation shows that Shatter not only renders these privacy attacks infeasible when each node operates 16 VNs but also exhibits a positive impact on model utility compared to standard DL. In summary, Shatter enhances the privacy of DL while maintaining the utility and efficiency of the model.

**Adaptive collaboration for online personalized distributed learning with heterogeneous clients.** In [48], we study the problem of online personalized decentralized learning with  $N$  statistically heterogeneous clients collaborating to accelerate local training. An important challenge in this setting is to select relevant collaborators to reduce gradient variance while mitigating the introduced bias. To tackle this, we introduce a gradient-based collaboration criterion, allowing each client to dynamically select peers with similar gradients during the optimization process. Our criterion is motivated by a refined and more general theoretical analysis

of the All-for-one algorithm, proved to be optimal in Even et al. (2022) for an oracle collaboration scheme. We derive excess loss upper-bounds for smooth objective functions, being either strongly convex, non-convex, or satisfying the Polyak-Łojasiewicz condition; our analysis reveals that the algorithm acts as a variance reduction method where the speed-up depends on a *sufficient variance*. We put forward two collaboration methods instantiating the proposed general schema; and we show that one variant preserves the optimality of All-for-one. We validate our results with experiments on synthetic and real datasets.

**When to Forget? Complexity Trade-offs in Machine Unlearning.** Machine Unlearning (MU) aims at removing the influence of specific data points from a trained model, striving to achieve this at a fraction of the cost of full model retraining. In [37], we analyze the efficiency of unlearning methods and establish the first upper and lower bounds on minimax computation times for this problem, characterizing the performance of the most efficient algorithm against the most difficult objective function. Specifically, for strongly convex objective functions and under the assumption that the forget data is inaccessible to the unlearning method, we provide a phase diagram for the unlearning complexity ratio – a novel metric that compares the computational cost of the best unlearning method to full model retraining. The phase diagram reveals three distinct regimes: one where unlearning at a reduced cost is infeasible, another where unlearning is trivial because adding noise suffices, and a third where unlearning achieves significant computational advantages over retraining. These findings highlight the critical role of factors such as data dimensionality, the number of samples to forget, and privacy constraints in determining the practical feasibility of unlearning.

**Stab-SGD: Noise-Adaptivity in Smooth Optimization with Stability Ratios.** In the context of smooth stochastic optimization with first order methods, in [33], we introduce the stability ratio of gradient estimates, as a measure of local relative noise level, from zero for pure noise to one for negligible noise. We show that a schedule-free variant (Stab-SGD) of stochastic gradient descent obtained by just shrinking the learning rate by the stability ratio achieves real adaptivity to noise levels (i.e. without tuning hyperparameters to the gradient’s variance), with all key properties of a good schedule-free algorithm: neither plateau nor explosion at initialization, and no saturation of the loss. We believe this theoretical development reveals the importance of estimating the local stability ratio in the construction of well-behaved (last-iterate) schedule-free algorithms, particularly when hyperparameter-tuning budgets are a small fraction of the total budget since noise-adaptivity and cheaper horizon-free tuning are most crucial in this regime.

### 7.3.2 Markov decision processes and reinforcement learning

**COGNAC: Cooperative Graph-based Networked Agent Challenges for Multi-Agent Reinforcement Learning.** Many controlled complex systems have an inherent network structure, such as power grids, traffic light systems, or computer networks. Automatically controlling these systems is highly challenging due to their combinatorial complexity. Standard single-agent reinforcement learning (RL) approaches often struggle with the curse of dimensionality in such settings. In contrast, the multi-agent paradigm offers a promising solution by distributing decision-making, thereby addressing both algorithmic and combinatorial challenges. In [34], we introduce COGNAC (COoperative Graph-based Networked Agent Challenges), a collection of cooperative graph-structured environments designed to facilitate experiments across different graph sizes and topologies. COGNAC bridges the gap between theoretical research in network control and practical multi-agent RL (MARL) applications by offering a flexible, scalable platform with a suite of simple yet highly challenging problems rooted in networked environments. Our benchmarks also support the development and evaluation of decentralized and distributed learning algorithms, motivated by the growing interest in more sustainable and frugal AI systems. Experiments on COGNAC show that independent actor-critic learning (IPPO) yields the highest-quality joint policies while scaling robustly to large network sizes with minimal hyperparameter tuning. Value-based independent learning (IDQL) typically needs substantially more training and is less reliable on combinatorial tasks. In contrast, standard Centralized-Training Decentralized-Execution (CTDE) methods and fully centralized training are slower to converge, less stable, and struggle to generalize to larger, more interdependent networks. These results suggest that CTDE approaches likely need extra information or inter-agent communication to fully capture the underlying network structure of each problem.

### **Efficient Solving of Large Single Input Superstate Decomposable Markovian Decision Process.**

Solving Markov Decision Processes (MDPs) remains a central challenge in sequential decision-making, especially when dealing with large state spaces and long-term optimization criteria. A key step in Bellman dynamic programming algorithms is the policy evaluation, which becomes computationally demanding in infinite-horizon settings such as average-reward or discounted-reward formulations. In the context of Markov chains, aggregation and disaggregation techniques have for a long time been used to reduce complexity by exploiting structural decompositions. In [47], we extend these principles to a structured class of MDPs. We define the Single-Input Superstate Decomposable Markov Decision Process (SISDMDP), which combines Chiu’s single-input decomposition with Robertazzi’s single-cycle recurrence property. When a policy induces this structure, the resulting transition graph can be decomposed into interacting components with centralized recurrence. We develop an exact and efficient policy evaluation method based on this structure. This yields a scalable solution applicable to both average and discounted reward MDPs.

### **A slot-based energy storage decision-making approach for optimal Off-Grid telecommunication operator.**

In [13], a slot-based energy storage approach is proposed for decision-making in the context of an Off-Grid telecommunication operator. We consider network systems powered by solar panels, where harvest energy is stored in a battery that can also be sold when fully charged. To reflect real-world conditions, we account for non-stationary energy arrivals and service demands that depend on the time of day, as well as the failure states of PV panel. The network operator we model faces two conflicting objectives: maintaining the operation of its infrastructure and selling (or supplying to other networks) surplus energy from fully charged batteries. To address these challenges, we developed a slot-based Markov Decision Process (MDP) model that incorporates positive rewards for energy sales, as well as penalties for energy loss and battery depletion. This slot-based MDP follows a specific structure we have previously proven to be efficient in terms of computational performance and accuracy. From this model, we derive the optimal policy that balances these conflicting objectives and maximizes the average reward function. Additionally, we present results comparing different cities and months, which the operator can consider when deploying its infrastructure to maximize rewards based on location-specific energy availability and seasonal variations. Experimental results show that our proposed algorithm outperforms classical methods in large-scale scenarios. While Relative Value Iteration algorithm remains competitive on smaller instances, its convergence time increases significantly under strict precision requirements (e.g.,  $\epsilon < 10^{-10}$ ). In contrast, our method maintains both speed and robustness, solving MDPs with up to  $2 \times 10^5$  states and 100 actions in under one hour, whereas standard approaches exceed  $10^4$  seconds.

### **7.3.3 Learning and control for energy networks**

**Moment Constrained Optimal Transport for Thermostatically Controlled Loads.** Controlling large populations of thermostatically controlled loads (TCLs), such as water heaters, poses significant challenges due to the need to balance global constraints (e.g., grid stability) with individual requirements (e.g., physical limits and quality of service). In [45], we introduce a novel framework based on Moment Constrained Optimal Transport (MCOT) for distributed control of TCLs. By formulating the control problem as an optimal transport problem with moment constraints, our approach integrates global consumption constraints and physical feasibility conditions into the control design. This problem with high (or infinite) dimensionality can be reduced to a much lower finite-dimensional problem. The structure of this problem allows for computing the gradient with Monte Carlo methods by generating trajectories of TCLs. Contrary to all previous work, in our MCOT framework, it is possible to choose the sampling law, which considerably speeds up the calculations. This algorithm mitigates the need for extensive state-space discretization and significantly reduces computational complexity compared to existing methods. Numerical experiments in a water heater case study demonstrate that our MCOT-based method effectively coordinates TCLs under various constraints. We further extend our approach to an online setting, illustrating its practical applicability on simulated data.

### **Moment Constrained Optimal Transport for Energy Demand Management of Heterogeneous Loads.**

In [25], we address the problem of coordinating a large population of heterogeneous electrical loads, such as electric vehicles (EVs) and water heaters (WHs), under global operational constraints. We extend the Moment Constrained Optimal Transport for Control (MCOT-C) framework to accommodate multiple

classes of agents with distinct dynamics and cost structures. Our formulation relies on a meanfield limit that captures agent heterogeneity through class-specific distributions. We propose a scalable gradient descent algorithm and a Model Predictive Control (MPC) scheme that enables online adaptation of this algorithm to uncertain or progressively revealed agent information. The proposed approach is validated through numerical experiments on real datasets for EVs and WHs, demonstrating the effectiveness of this method in enforcing global constraints while preserving agent-level dynamics.

**Mean-Field Control of Heterogeneous Piecewise Deterministic Markov Processes under Grid Constraints.** In [30], we investigate mean-field control problems for large populations of agents modeled by piecewise deterministic Markov processes. Motivated by the growing need for demand-side management in power systems, we develop a decentralized control framework that accounts for key operational grid constraints such as maximum aggregate power and ramping limits. These constraints are critical to ensure the feasibility and reliability of control strategies in real-world power grids. Furthermore, we address the limitation of classical mean-field approaches that assume agent homogeneity by extending the formulation to heterogeneous populations, where agent-specific characteristics (such as charging power or battery capacity) are encoded through fixed parameters. The resulting framework enables the scalable and feasible coordination of various flexible loads. Theoretical developments are illustrated through applications to electric vehicle charging, demonstrating the impact of the proposed constraints and heterogeneity modeling on the system's aggregate behavior.

**Mean Field Control of Thermostatically Controlled Loads as Piecewise Deterministic Markov Processes.**

In [44], we present a mean-field control approach for Piecewise Deterministic Markov Processes (PDMPs), specifically designed for controlling a large number of agents. By modeling the interactions of a large number of agents through an aggregate cost function, the proposed method mitigates the high dimensionality of the problem by focusing on a representative agent. The contribution of this work is the application of a PDMP-based mean-field control framework to the coordination of a large population of Thermostatically Controlled Loads (TCLs). Adapting this framework to TCLs requires incorporating a quality-of-service constraint ensuring that each agent's temperature remains within a specified comfort range. To achieve this, an additional jump intensity is introduced so that agents are very likely to switch between heating and cooling modes when they reach the boundaries of their temperature range. This extension to TCLs is demonstrated through Water Heaters (WHs) control, with a decentralized algorithm based on a dual formulation and stochastic gradient descent. The numerical results obtained illustrate this approach on two examples (signal tracking and taking into account energy price).

**Achieving a Collective Target Through Incentives.** In many games, the payoff of individual users is a function of a collective outcome that is common to all agents. Often, the users may be interested in jointly steering this outcome to a desired value. In [28], we present such a scenario, where the collective outcome is strictly monotone in the joint strategy of the agents. Further, the agents are irrational in their perception of a random cost component in their payoff. This irrationality is modelled using prospect theory. A coordinator steers the game to a desired collective outcome, by designing incentives. These incentives modify the responses of the users. Owing to the potential structure of the game, the system converges to a Nash equilibrium at which the desired collective outcome is obtained.

**How Irrationality Shapes Nash Equilibria: A Prospect-Theoretic Perspective.** Noncooperative games with uncertain payoffs have been classically studied under the expected-utility theory framework, which relies on the strong assumption that agents behave rationally. However, simple experiments on human decision makers found them to be not fully rational, due to their subjective risk perception. Prospect theory was proposed as an empirically-grounded model to incorporate irrational behaviours into game-theoretic models. But, how prospect theory shapes the set of Nash equilibria when considering irrational agents, is still poorly understood. To this end, in [29], we study how prospect theoretic transformations may generate new equilibria while eliminating existing ones. Focusing on aggregative games, we show that capturing users' irrationality can preserve symmetric equilibria while causing the vanishing of asymmetric equilibria. Further, there exist value functions which map uncountable sets of equilibria in the expected-utility maximization framework to

finite sets. This last result may shape some equilibrium selection theories for human-in-the-loop systems where computing a single equilibrium is insufficient and comparison of equilibria is needed.

## 7.4 Other results

**Attacking and Fixing the Android Protected Confirmation Protocol.** Android Protected Confirmation (APC) is an authentication protocol designed by Google. It leverages the extra security of the Trusted Execution Environment (TEE) to secure transactions even in the presence of a compromised OS. The intended security guarantee for APC is that if a transaction has been signed under APC, then the user must have previously given its explicit consent, even if an attacker has gained root access to the victim’s Android OS. In [39], we present a security analysis of APC in the Universal Composability (UC) framework. We uncover two attacks on the design of the protocol which allow a root adversary to issue transactions without the user consenting to them. We provide an attack implementation on a Google Pixel phone, and propose light-weight fixes. Finally, we specify the ideal UC functionality capturing the intended security guarantees for APC, and prove that the fixed protocol UC-realizes it.

# 8 Bilateral contracts and grants with industry

## 8.1 Bilateral contracts with industry

**Participants:** Marc Lelarge, Julien Moreau.

CIFRE PhD thesis of Julien Moreau with RENAULT, advisor Marc Lelarge.

# 9 Partnerships and cooperations

## 9.1 International research visitors

### 9.1.1 Visits of international scientists

#### Inria International Chair

**Participants:** Sean Meyn, ARGO and SIERRA teams.

Prof Sean Meyn (University of Florida) is holding Inria International Chair from 2019 to 2025.

#### Other international visits to the team

**Alessio Giorlandino**

**Status** PhD student

**Institution of origin:** SISSA (Trieste)

**Country:** Italy

**Dates:** October 6th - April 6th 2026

**Context of the visit:** Research visit to Antoine Maillard

**Mobility program/type of mobility:** Research stay

**Seyoung Yun****Status** Professor**Institution of origin:** KAIST**Country:** Italy**Dates:** from May 2025 until Jul 2025**Context of the visit:** Research visit to ARGO team**Mobility program/type of mobility:** Research stay**9.2 National initiatives****9.2.1 Project REDEEM, PEPR IA**

Project **REDEEM** aims to explore new distributed learning approaches that are resilient, robust to noise and adversarial attacks, and respectful of privacy. These distributed approaches should make it possible to go beyond current federated learning. From a theoretical point of view, REDEEM aims to provide a solid foundation for the proposed approaches, particularly in the case of malicious protagonists participating in the learning phase, and with the overriding objective of ensuring data confidentiality as far as possible. In addition to new approaches to distributing learning, REDEEM also aims for efficient implementations, by offering the community open-source code and tools.

**9.2.2 Project AI-NRGY, PEPR TASE**

The objective of **AI-NRGY** is to address the major constraints of tomorrow's energy networks (highly distributed, dynamic, heterogeneous, critical and sometimes volatile) by contributing to the implementation of distributed intelligence solutions leveraging different computing methods (edge, fog and cloud computing), and by proposing a software architecture as well as the methods, models and algorithms necessary for the implementation of distributed intelligence solutions likely to accelerate the digitalization of energy networks.

**9.2.3 Challenge Inria FedMalin**

Members of ARGO participate in the **FedMalin** Inria défi on Federated Learning.

**9.2.4 Challenge INRIA-EDF**

Ashok Krishnan Komalan Sindhu is a post-doc within **INRIA-EDF** challenge on Managing tomorrow's power systems, collaborating with A. Bušić and H  l  ne Le Cadre (INOCS, INRIA Lille).

The challenge entitled "managing tomorrow's power systems" aims to imagine and develop new tools (methods, regulatory approach, algorithms, software) and new sources of information (meters, sensors at higher temporal and spatial resolutions, weather forecasts, electric vehicle charging stations, etc.) to support strategic and operational decisions for economically, ecologically and resilient management of new power systems as part of the ecological transition and and climate change.

More specifically, the operational implementation of the scenarios defined by RTE for the evolution of the French power system requires the development of mechanisms and tools to enable decisions to be taken from the long term to the short term.

**9.2.5 Challenge LLM4CODE**

M. Lelarge participates to the challenge **LLM4CODE**: Reliable and Productive Code Assistants Based on Large Language Models.

Generative AI, in particular the recent Large Language Models (LLMs), show great promise for software developments. Specialized models are now able to perform an impressive variety of programming tasks: solving programming exercises, assisting software developers, or even generating mechanized proofs. Yet,

many challenges still need to be addressed to build reliable and productive LLM-based coding assistants: improving the quality of the generated code, increasing the developers' confidence in the generated code, enabling interaction with other software development tools (verification, test), and providing new capabilities (automated migration and evolution of software).

The goal of the Challenge Inria LLM4Code is to leverage LLM capabilities to build code assistants that can enhance both reliability and productivity.

### 9.2.6 Joint IFPEN - Inria laboratory

The joint laboratory IFPEN - Inria Convergence HPC / AI / HPDA for the energetic transition is a strategic partnership between IFPEN and Inria under the form of a "contrat-cadre", started in June 2020 in continuation of previous collaborations. The goal of this laboratory without premises is to promote the collaborations between the two institutions on account of the energy transition. The two organizations propose here funds for thesis and post-doctoral projects on the basis of an annual call for proposals.

Co-supervised PHD thesis within IFPEN - Inria laboratory:

- Baptiste Corban, started in November 2024, co-supervised by A. Bušić, Donatien Dubuc (IFPEN), and Jiamin Zhu (IFPEN).

Patent granted in 2025: [49].

### 9.2.7 GdR ROD

A. Bušić is coordinating, with E. Hyon (LIP 6), the **working group COSMOS** (Stochastic optimization and control, modeling and simulation), of **GdR-ROD** (Recherche Opérationnelle et Décision).

## 9.3 Regional initiatives

### 9.3.1 PRAIRIE-PSAI

**PR[AI]RIE-PSAI** is one of the 9 French AI-Clusters financed by France 2030.

Created in 2019 by CNRS, Inria, Institut Pasteur, PSL University, Université de Paris Cité, and a club of industrial partners, PaRis Artificial Intelligence Research InstitutE (PR[AI]RIE) was one of the 4 Interdisciplinary Institutes for AI research set up as part of the national strategy for Artificial Intelligence announced by the President of the French Republic in March 2018.

Starting in 2024 it has evolve to become the PR[AI]RIE-PSAI and cover the triptych of research-training-innovation.

## 10 Dissemination

### 10.1 Promoting scientific activities

#### 10.1.1 Invited talks

Ana Basic:

- Plenary talk at **ROADEF 2025 conference**.
- Invited talk at **8th Workshop on Cognition & Control**.

Marc Lelarge:

- Long lecture at **Journées ALEA 2025**
- Tutorial lecture at **Ecole d'automne du PEPR TASE**

Antoine Maillard:

- January 2025: Spin glasses and related topics day, at the Inhomogeneous Random Systems annual workshop, Institut Henri Poincaré, Paris.
- February 2025: Workshop "Towards a theory of typical-case algorithmic hardness" in Les Houches, France.
- February 2025: Probability Seminar, École Normale Supérieure de Lyon, France.
- June 2025: Data Science Seminar, SISSA, Trieste, Italy.
- June 2025: Invited short lecture (3h) on statistical physics and learning at the ASCAI meeting, Orsay, France.
- August 2025: Workshop "Statistical Physics and Machine Learning: moving forward", Cargèse, France.
- September 2025: "Workshop Random Matrix Theory for Learning and Statistical Physics", ZiF, Bielefeld, Germany.
- December 2025: Maths-IA seminar, Institut de Mathématiques de Bordeaux, Bordeaux, France.

### 10.1.2 Scientific expertise

M. Lelarge, member of the steering Committee of Fondation Sciences Mathématiques de Paris (FSMP).

### 10.1.3 Research administration

- A. Bušić, member of the laboratory board of DIENS, until October 2025.
- A. Bušić, member of the selection committee for CRCN/ISFP positions, Inria Paris.
- L. Massoulié, "délégué scientifique" of the Inria Paris center.
- L. Viennot, member of the committee for health and working conditions (FSS) at Paris Inria research center.

## 10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

### 10.2.1 Teaching

Most of the permanent researchers teach and are responsible for courses proposed within DIENS at the L3/M1 level and in M2 master programs at PSL and/or Paris Saclay: MASH, MASEF, MPRI, MVA.

- L. Budzynski is MdC at DIENS, Ecole Normale Supérieure, PSL.
- A. Bušić is Professeur Attaché IA at PSL University.
- M. Lelarge is Professeur Attaché at ENS, PSL and responsible for the « cursus maths-info » at ENS.
- L. Massoulié is Professeur Chargé de Cours at Ecole Polytechnique, Centre de mathématiques appliquées.

Some courses in 2025:

- L. Viennot teaches "Theory of practical graph algorithms" (12h), M2 level, MPRI master program.
- L. Budzynski. teaches "Structures et algorithmes aléatoires" (24h), L3 level, DI ENS
- L. Budzynski. teaches "Algorithmics" (60h), L2 level, CPES
- A. Bušić and L. Massoulié teach "Foundations of network models" (24h), M2 level, MPRI
- A. Bušić teaches "Reinforcement learning" (24h), M2 level MASH/MASEF
- A. Bušić teaches "Modèles et algorithmes des réseaux" (24h), M1 level DI ENS
- A. Bušić teaches an introductory course to reinforcement learning (18h), within the PSL Data science & AI for academics training programme, aimed to a broad scientific audience (all disciplines within PSL).

- M. Lelarge is teaching a machine learning course at ENS (M1) and in the master ICFP (M2)
- M. Lelarge is teaching a course on LLM for code and proof, M2 level, at master IASD and MVA
- L. Massoulié is teaching "Inference in large random graphs", M2 level, Mathématiques de l'Aléatoire, Université Paris Saclay.

### 10.2.2 Supervision

PhD defended:

- Lucas Weber, Exploiting Partial System Knowledge in Reinforcement Learning for Admission Control and Electricity Storage Optimization, defended 17 January 2025, supervised by A. Bušić and J. Zhu (IFPEN)
- David Robin, Construction and convergence of provably-correct neural networks, defended 27 June 2025, supervised by M. Lelarge
- Romain Cosson, Collective algorithms for exploration and decision-making, defended 11 July 2025, supervised by L. Massoulié
- Thomas Le Corre, Distributed control of flexible loads in power grids, defended 22 October 2025, supervised by A. Bušić

PhD in progress:

- Jakob Maier, since October 2022, supervised by L. Massoulié
- Killian Bakong Epoune, since September 2023, supervised by L. Massoulié and K. Scaman
- Martin Van Waerebeke, since September 2023, supervised by K. Scaman, Marco Lorenzi (EPIONE team) et de Giovanni Neglia (NEO team)
- Jean Adrien Lagesse, since September 2023, supervised by M. Lelarge
- Shu Li, since September 2023, supervised by A. Bušić and J.-M. Fourneau
- Baptiste Corban IFPEN, since November 2024, supervised by A. Bušić, D. Dubuc (IFPEN) and J. Zhu (IFPEN)
- Jules Sintès, since November 2024, supervised by A. Bušić
- Jules Viennot, since September 2024, supervised by M. Lelarge
- Julien Moreau, since September 2024, supervised by M. Lelarge
- Julien Cardinal, since March 2025, supervised by A. Bušić
- Louis Vassaux, since September 2025, supervised by L. Massoulié
- Pierre Aguié, since September 2025, supervised by L. Massoulié
- Andrea Vincenzo Dell'abate, since September 2025, supervised by L. Budzynski and L. Massoulié
- Louann Coste, since September 2025, supervised by L. Viennot
- Jackson Carter, since October 2025, supervised by R. Cornette (ISYEB), M. Lelarge, and Ronan ALLAIN (CR2P)
- Alexandra Kortchemski, since November 2025, supervised by L. Massoulié and A. Bušić
- Antoine Lunven, since December 2025, supervised by A. Bušić and A. Bechihi (MERCÉ)

### 10.2.3 Juries

Reviewer in PHD juries:

- Ana Bušić: Rémy Hosseinkhan-Boucher, "On Learning-Based Control of Dynamical Systems" (Université Paris-Saclay), advisors: Anne Vilnat, Lionel Mathelin, Onofrio Semeraro
- Ana Bušić: Elie Kadoche, "Deep reinforcement learning for wind farm flow control" (Institut Polytechnique de Paris), advisor: Pascal Bianchi
- Ana Bušić: Daniel Mastropietro, "Fleming-Viot particle systems for faster exploration in reinforcement learning and combinatorial optimisation" (Université de Toulouse), advisors: Urtzi Ayesta, Matthieu Jonckheere
- Ana Bušić: Hamed El Ferchichi, "Learning Structured Rankings and Matchings in Stochastic Bandit Models" (ENSAE, Institut Polytechnique de Paris), advisors: Matthieu Lerasle and Vianney Perchet
- Marc Lelarge: Yassine Abbahaddou, "Key Principles of Graph Machine Learning : Representation, Robustness, and Generalization" (Institut Polytechnique de Paris), advisor: Michalis Vazirgiannis

Examinator in HDR juries:

- Laurent Massoulié: Malisa Vucinic (PSL)

Examinator in PHD juries:

- Ana Bušić: president of the jury for Chiara Mignacco, "A mathematical study of policy orchestration for reinforcement learning" (Université Paris-Saclay), advisors: Gilles Stoltz, Matthieu Jonckheere
- Ana Bušić: president of the jury for Francesco Morri, "Game-Theoretic Learning in Intelligent Marketplaces", advisors: Luce Brotcorne, Hélène Le Cadre
- Ana Bušić: jury member for Mohammed Abdullah, "Optimizing resource allocation in future wireless networks under uncertainty" (Institut Polytechnique de Paris), advisors: Tijani Chahed, Salah Eddine El Ayoubi
- Ana Bušić: jury member for Konstantinos Parginos, "Advanced Methods for Enhancing the Interpretability of AI tools in the Energy Sector" (Mines ParisTech, PSL), advisors: Georges Kariniotakis, Ricardo Bessa
- Marc Lelarge: president of the jury for Félix Marcoccia, "Topology optimization in mobile wireless networks using machine learning" (Sorbonne Université), advisors: Paul Mühlethaler, Cédric Adjih
- Laurent Massoulié: jury member for Mathieu Molina, "The impact of fairness in resource allocation" (ENSAE, Institut Polytechnique de Paris), advisors: Vianney Perchet, Patrick Loiseau, Nicolas Gast
- Laurent Viennot: jury member for Maxime Dupuy, "Adding some continuity in discrete ship weather routing" (Institut Polytechnique de Paris), advisors: Leo Liberti et Claudia D'Ambrosio.

## 10.3 Popularization

### 10.3.1 Productions (articles, videos, podcasts, serious games, ...)

A. Basic, C. Bizon Monroc, J. Sintès. Article: "[L'apprentissage par renforcement au service de l'optimisation](#)".

## 11 Scientific production

### 11.1 Major publications

- [1] A. Bandeira and A. Maillard. ‘Exact threshold for approximate ellipsoid fitting of random points’. In: *Electronic Journal of Probability* 30.none (1st Jan. 2025). DOI: [10.1214/25-EJP1378](https://doi.org/10.1214/25-EJP1378). URL: <https://inria.hal.science/hal-05495104>.
- [2] A. Braunstein, L. Budzynski, M. Mariani and F. Ricci-Tersenghi. ‘Evidence of replica symmetry breaking under the Nishimori conditions in epidemic inference on graphs’. In: *Physical Review E* 111.6 (2025), p. 064308. DOI: [10.1103/tbd5-25vf](https://doi.org/10.1103/tbd5-25vf). URL: <https://hal.science/hal-05442589>.
- [3] F. Brunelli, P. Crescenzi and L. Viennot. ‘Making Temporal Betweenness Computation Faster and Restless’. In: *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2024, Barcelona, Spain, August 25-29, 2024*. KDD ’24: The 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. Barcelona, Spain: ACM, 24th Aug. 2024, pp. 163–174. DOI: [10.1145/3637528.3671825](https://doi.org/10.1145/3637528.3671825). URL: <https://hal.science/hal-04901946>.
- [4] N. Cammardella, A. Bušić and S. Meyn. ‘Kullback–Leibler-Quadratic Optimal Control’. In: *SIAM Journal on Control and Optimization* 61.5 (25th Oct. 2023), pp. 3234–3258. DOI: [10.1137/21M1433654](https://doi.org/10.1137/21M1433654). URL: <https://inria.hal.science/hal-04476982>.
- [5] F. F. Dragan, G. Ducoffe, M. Habib and L. Viennot. ‘Certificates in P and Subquadratic-Time Computation of Radius, Diameter, and all Eccentricities in Graphs’. In: *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. New Orleans (LA), United States, 12th Jan. 2025, pp. 2157–2193. DOI: [10.1137/1.9781611978322.70](https://doi.org/10.1137/1.9781611978322.70). URL: <https://inria.hal.science/hal-01729748>.
- [6] L. Ganassali, M. Lelarge and L. Massoulié. ‘Correlation detection in trees for planted graph alignment’. In: *The Annals of Applied Probability* 34.3 (1st June 2024). DOI: [10.1214/23-AAP2020](https://doi.org/10.1214/23-AAP2020). URL: <https://hal.science/hal-04893755>.
- [7] L. Ganassali, L. Massoulié and G. Semerjian. ‘Statistical limits of correlation detection in trees’. In: *The Annals of Applied Probability* 34.4 (1st Aug. 2024), pp. 3701–3734. DOI: [10.1214/23-AAP2048](https://doi.org/10.1214/23-AAP2048). URL: <https://hal.science/hal-04893751>.
- [8] C. B. Monroc, A. Bušić, D. Dubuc and J. Zhu. ‘WFCRL: A Multi-Agent Reinforcement Learning Benchmark for Wind Farm Control’. In: *Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track. Vancouver, Canada, 10th Dec. 2024. URL: <https://hal.science/hal-04864926>.
- [9] D. Robin, K. Bakong and K. Scaman. ‘Stab-SGD: Noise-Adaptivity in Smooth Optimization with Stability Ratios’. In: *NeurIPS 2025 - The Thirty-Ninth Annual Conference on Neural Information Processing Systems*. San Diego (CA), United States, 6th Dec. 2025. URL: <https://hal.science/hal-05497993>.
- [10] K. Scaman, M. Even, B. L. Bars and L. Massoulié. ‘Minimax Excess Risk of First-Order Methods for Statistical Learning with Data-Dependent Oracles’. In: *AISTATS 2024 - International Conference on Artificial Intelligence and Statistics*. Valencia, Spain, 2024. URL: <https://hal.science/hal-04895743>.
- [11] M. van Waerebeke, M. Lorenzi, G. Neglia and K. Scaman. ‘When to Forget? Complexity Trade-offs in Machine Unlearning’. In: *ICML 2025 - International Conference in Machine Learning*. Vancouver (BC), Canada, 13th July 2025. URL: <https://hal.science/hal-05464347>.
- [12] Y. Xu, A. Maillard, F. Krzakala and L. Zdeborová. ‘Fundamental Limits of Matrix Sensing: Exact Asymptotics, Universality, and Applications’. In: *38th Annual Conference on Learning Theory*. Lyon, France, 12th Sept. 2025. URL: <https://inria.hal.science/hal-05495125>.

## 11.2 Publications of the year

### International journals

- [13] Y. Ait El Mahjoub and J.-M. Fourneau. ‘A slot-based energy storage decision-making approach for optimal Off-Grid telecommunication operator’. In: *Computer Communications* 242 (23rd July 2025), p. 108273. DOI: [10.1016/j.comcom.2025.108273](https://doi.org/10.1016/j.comcom.2025.108273). URL: <https://hal.science/hal-04838301> (cit. on p. 19).
- [14] M. C. Angelini, L. Budzynski and F. Ricci-Tersenghi. ‘Interacting copies of random-constraint satisfaction problems’. In: *Physical Review E* 112.6 (12th Dec. 2025), p. 064117. DOI: [10.1103/86dj-6w3r](https://doi.org/10.1103/86dj-6w3r). URL: <https://hal.science/hal-05442592> (cit. on p. 12).
- [15] A. S. Bandeira and A. Maillard. ‘Exact threshold for approximate ellipsoid fitting of random points’. In: *Electronic Journal of Probability* 30.none (1st Jan. 2025). DOI: [10.1214/25-EJP1378](https://doi.org/10.1214/25-EJP1378). URL: <https://inria.hal.science/hal-05495104> (cit. on p. 12).
- [16] A. Braunstein, L. Budzynski, M. Mariani and F. Ricci-Tersenghi. ‘Evidence of replica symmetry breaking under the Nishimori conditions in epidemic inference on graphs’. In: *Physical Review E* 111.6 (2025), p. 064308. DOI: [10.1103/tbd5-25vf](https://doi.org/10.1103/tbd5-25vf). URL: <https://hal.science/hal-05442589> (cit. on p. 11).
- [17] M. Csikós, M. Habib, M.-H. Nguyen, M. Rabie and L. Viennot. ‘Forbidden Patterns in Temporal Graphs Resulting from Encounters in a Corridor’. In: *Journal of Computer and System Sciences* (Jan. 2025), p. 103620. DOI: [10.1016/j.jcss.2025.103620](https://doi.org/10.1016/j.jcss.2025.103620). URL: <https://hal.science/hal-04901976> (cit. on p. 14).
- [18] V. Erba, E. Troiani, L. Biggio, A. Maillard and L. Zdeborová. ‘Bilinear Sequence Regression: A Model for Learning from Long Sequences of High-Dimensional Tokens’. In: *Physical Review X* 15.2 (16th June 2025), p. 021092. DOI: [10.1103/14p2-vrxt](https://doi.org/10.1103/14p2-vrxt). URL: <https://inria.hal.science/hal-05495130> (cit. on p. 12).
- [19] M. Leiber, Y. Marnissi, A. Barrau, S. Meignen and L. Massoulié. ‘Optimal Adaptive Time-Frequency Representation via Differentiable Short-Time Fourier Transform’. In: *IEEE Transactions on Signal Processing* 73 (29th Oct. 2025), pp. 5047–5059. DOI: [10.1109/TSP.2025.3624477](https://doi.org/10.1109/TSP.2025.3624477). URL: <https://hal.science/hal-05486435> (cit. on p. 16).
- [20] A. Maillard. ‘Average-Case Matrix Discrepancy: Satisfiability Bounds’. In: *Random Structures and Algorithms* 67.3 (29th Oct. 2025). DOI: [10.1002/rsa.70033](https://doi.org/10.1002/rsa.70033). URL: <https://inria.hal.science/hal-05495133> (cit. on p. 13).
- [21] A. Maillard, A. S. Bandeira, D. Belius, I. Dokmanić and S. Nakajima. ‘Injectivity of ReLU networks: Perspectives from statistical physics’. In: *Applied and Computational Harmonic Analysis* 76 (Apr. 2025), p. 101736. DOI: [10.1016/j.acha.2024.101736](https://doi.org/10.1016/j.acha.2024.101736). URL: <https://inria.hal.science/hal-05495127> (cit. on p. 13).

### International peer-reviewed conferences

- [22] X. Bai, C. Coester and R. Cosson. ‘Unweighted Layered Graph Traversal: Passing a Crown via Entropy Maximization’. In: SODA 2025 - Symposium on Discrete Algorithms. New Orleans (Louisiane), United States: Society for Industrial and Applied Mathematics, 7th Jan. 2025, pp. 3884–3900. DOI: [10.1137/1.9781611978322.131](https://doi.org/10.1137/1.9781611978322.131). URL: <https://hal.science/hal-04894327> (cit. on p. 14).
- [23] S. Biswas, M. Even, A.-M. Kermarrec, L. Massoulié, R. Pires, R. Sharma and M. de Vos. ‘Noiseless Privacy-Preserving Decentralized Learning’. In: PETS 2025 - 25th Privacy Enhancing Technologies Symposium. Vol. 2025. Bristol, United Kingdom, Jan. 2025, pp. 824–844. DOI: [10.56553/popets-2025-0043](https://doi.org/10.56553/popets-2025-0043). URL: <https://hal.science/hal-04920384> (cit. on p. 17).
- [24] A. Bušić, J.-M. Fourneau, A. Lunven and S. Li. ‘Performance Paradoxes in Matching Systems are not that Rare’. In: NETGCOOP 2025 - 12th international conference on Network games, artificial intelligence, control and optimization. Vol. LNCS-16173. Lecture Notes in Computer Science. Bilbao, Spain: Springer Nature Switzerland, 8th Nov. 2025, pp. 25–34. DOI: [10.1007/978-3-032-09315-8\\_3](https://doi.org/10.1007/978-3-032-09315-8_3). URL: <https://hal.science/hal-05441137> (cit. on p. 16).

- [25] J. Cardinal, T. Le Corre and A. Bušić. ‘Moment Constrained Optimal Transport for Energy Demand Management of Heterogeneous Loads’. In: NETGCOOP 2025 - 12th International Conference of Networks, Games, Control and Optimization. Bilbao, Spain, Nov. 2025. URL: <https://hal.science/hal-05375478> (cit. on p. 19).
- [26] J. Cauvi and L. Viennot. ‘Parameterized Restless Temporal Path’. In: FCT 2025 - 25th International Symposium on Fundamentals of Computation Theory. Wroclaw, Poland, 2025. URL: <https://inria.hal.science/hal-05144137> (cit. on p. 15).
- [27] F. F. Dragan, G. Ducoffe, M. Habib and L. Viennot. ‘Certificates in P and Subquadratic-Time Computation of Radius, Diameter, and all Eccentricities in Graphs’. In: Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). New Orleans (LA), United States, 12th Jan. 2025, pp. 2157–2193. DOI: [10.1137/1.9781611978322.70](https://doi.org/10.1137/1.9781611978322.70). URL: <https://inria.hal.science/hal-01729748> (cit. on p. 14).
- [28] A. Krishnan, H. Le Cadre and A. Bušić. ‘Achieving a Collective Target Through Incentives’. In: *Network Games, Artificial Intelligence, Control and Optimization: 12th International Conference, NETGCOOP 2025, Bilbao, Spain, October 8–10, 2025, Proceedings*. NETGCOOP 2025 - 12th International Conference of Networks, Games, Control and Optimization. Bilbao, Spain, 9th Nov. 2025. DOI: [10.1007/978-3-032-09315-8\\_6](https://doi.org/10.1007/978-3-032-09315-8_6). URL: <https://hal.science/hal-05424025> (cit. on p. 20).
- [29] A. Krishnan K. S., H. Le Cadre and A. Bušić. ‘How Irrationality Shapes Nash Equilibria: A Prospect-Theoretic Perspective’. In: *Proceedings of 64th IEEE Conference on Decision and Control (CDC) 2025*. 64th IEEE Conference on Decision and Control (CDC) 2025. Rio de Janeiro, Brazil, 2025, pp. 4428–4433. URL: <https://hal.science/hal-05036781> (cit. on p. 20).
- [30] T. Le Corre and A. Bušić. ‘Mean-Field Control of Heterogeneous Piecewise Deterministic Markov Processes under Grid Constraints’. In: SmartGridComm 2025 - IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids. Toronto, Canada, Nov. 2025. URL: <https://hal.science/hal-05375490> (cit. on p. 20).
- [31] A. Maillard, E. Troiani, S. Martin, F. Krzakala and L. Zdeborová. ‘Bayes-optimal learning of an extensive-width neural network from quadratically many samples’. In: *Advances in Neural Information Processing Systems 37*. NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver, Canada: Neural Information Processing Systems Foundation, Inc. (NeurIPS), 31st Jan. 2025, pp. 82085–82132. DOI: [10.52202/079017-2609](https://doi.org/10.52202/079017-2609). URL: <https://inria.hal.science/hal-05495117> (cit. on p. 13).
- [32] C. Philippenko, K. Scaman and L. Massoulié. ‘In-depth Analysis of Low-rank Matrix Factorisation in a Federated Setting’. In: AAAI 2025 - 39th Annual AAI Conference on Artificial Intelligence. Philadelphia, United States, 25th Feb. 2025. URL: <https://hal.science/hal-04893622> (cit. on p. 17).
- [33] D. Robin, K. Bakong and K. Scaman. ‘Stab-SGD: Noise-Adaptivity in Smooth Optimization with Stability Ratios’. In: NeurIPS 2025 - The Thirty-Ninth Annual Conference on Neural Information Processing Systems. San Diego (CA), United States, 6th Dec. 2025. URL: <https://hal.science/hal-05497993> (cit. on p. 18).
- [34] J. Sintès and A. Bušić. ‘COGNAC: Cooperative Graph-based Networked Agent Challenges for Multi-Agent Reinforcement Learning’. In: NeurIPS 2025 - 39th Annual Conference on Neural Information Processing Systems. San Diego, United States, 2nd Dec. 2025. URL: <https://hal.science/hal-05488783> (cit. on p. 18).
- [35] S. M. Varma, I. Waldspurger and L. Massoulié. ‘Graph Alignment via Birkhoff Relaxation’. In: *NeurIPS 2025*. NeurIPS 2025 - Thirty-Ninth Annual Conference on Neural Information Processing Systems. San Diego (CA), United States, 2nd Dec. 2025. URL: <https://inria.hal.science/hal-05387695> (cit. on p. 14).
- [36] J. Viennot, G. Baudart, E. J. Gallego Arias, M. Lelarge and T. Stoskopf. ‘Tacq -Context Aware Tactic Recommendation for Rocq’. In: JFLA 2026 – 37es Journées Francophones des Langues Applicatifs. Vol. JFLA 2026 – 37es Journées Francophones des Langues Applicatifs. Oberbronn, France, 27th Jan. 2026. URL: <https://hal.science/hal-05428141> (cit. on p. 17).

- [37] M. van Waerebeke, M. Lorenzi, G. Neglia and K. Scaman. ‘When to Forget? Complexity Trade-offs in Machine Unlearning’. In: ICML 2025 - International Conference in Machine Learning, Vancouver (BC), Canada, 13th July 2025. URL: <https://hal.science/hal-05464347> (cit. on p. 18).
- [38] Y. Xu, A. Maillard, F. Krzakala and L. Zdeborová. ‘Fundamental Limits of Matrix Sensing: Exact Asymptotics, Universality, and Applications’. In: COLT 2025 - 38th Annual Conference on Learning Theory, Lyon, France, 12th Sept. 2025. URL: <https://inria.hal.science/hal-05495125> (cit. on p. 13).

### Conferences without proceedings

- [39] M. Arapinis, V. Danos, M. Racouchot, D. Robin and T. Zacharias. ‘Attacking and Fixing the Android Protected Confirmation Protocol’. In: The 10th IEEE European Symposium on Security and Privacy, Venice, Italy: IEEE, 30th June 2025, pp. 458–479. DOI: [10.1109/EuroSP63326.2025.00034](https://doi.org/10.1109/EuroSP63326.2025.00034). URL: <https://hal.science/hal-05038009> (cit. on p. 21).
- [40] J. Moreau and M. Lelarge. ‘Assimilation of Sparse Vehicle Trajectories with a Macroscopic Traffic Model’. In: Machine Learning for the Physical Sciences @ NeurIPS 2025, San Diego, CA, United States, 6th Dec. 2025. URL: <https://inria.hal.science/hal-05501423> (cit. on p. 16).
- [41] J. Viennot, G. Baudart, E. J. Gallego Arias and M. Lelarge. ‘Minif2f in Rocq: Automatic Translation Between Proof Assistants -A Case Study’. In: 5th Workshop on Mathematical Reasoning and AI (MathAI@NeurIPS 2025), San Diego (CA), United States, 6th Dec. 2025. URL: <https://hal.science/hal-05467134> (cit. on p. 16).

### Reports & preprints

- [42] G. Aubian, F. Brunelli, F. F. Dragan, G. Ducoffe, M. Habib, A. Ibiapina and L. Viennot. *Complexity Gaps between Point and Interval Temporal Graphs for some Reachability Problems*. 17th Jan. 2025. URL: <https://inria.hal.science/hal-04863343> (cit. on p. 15).
- [43] J. Cauvi, N. Morawietz and L. Viennot. *Foremost, Fastest, Shortest: Temporal Graph Realization under Various Path Metrics*. Oct. 2025. URL: <https://inria.hal.science/hal-05293265> (cit. on p. 15).
- [44] T. L. Corre, A. Séguret and A. Bušić. *Mean Field Control of Thermostatically Controlled Loads as Piecewise Deterministic Markov Processes*. 3rd Nov. 2025. URL: <https://hal.science/hal-05499491> (cit. on p. 20).
- [45] T. Le Corre, J. Cardinal and A. Bušić. *Moment Constrained Optimal Transport for Thermostatically Controlled Loads*. 27th Aug. 2025. URL: <https://hal.science/hal-05226794> (cit. on p. 19).
- [46] M. Lelarge. *Bootstrap learning for combinatorial graph alignment with sequential GNNs*. 3rd Oct. 2025. URL: <https://hal.science/hal-05405438> (cit. on p. 16).
- [47] Y. A. E. Mahjoub, J.-M. Fourneau and S. Alouah. *Efficient Solving of Large Single Input Superstate Decomposable Markovian Decision Process*. 1st Aug. 2025. DOI: [10.48550/arXiv.2508.00816](https://doi.org/10.48550/arXiv.2508.00816). URL: <https://hal.science/hal-05198828> (cit. on p. 19).
- [48] C. Philippenko, B. Le Bars, K. Scaman and L. Massoulié. *Adaptive collaboration for online personalized distributed learning with heterogeneous clients*. 9th July 2025. URL: <https://hal.science/hal-05175097> (cit. on p. 17).

### Patents

- [49] E. Bouba, D. Dubuc, J. Zhu, C. B. Monroc and A. Bušić. ‘Method for controlling a wind farm using a reinforcement learning method’. US12241455B2 (United States). 4th Mar. 2025. URL: <https://hal.science/hal-04864919> (cit. on p. 23).