RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

Inria

2025 *Activity Report*

RESEARCH CENTRE: **Inria Centre at Université Côte d'Azur**

Project-Team

NEO

Network Engineering and Operations

Project-Team NEO

Creation of the Project-Team: 2017 December 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

A1.1.11. – Quantum architectures
A1.2.4. – QoS, performance evaluation
A1.2.6. – Sensor networks
A1.2.9. – Social Networks
A1.2.11. – Quantum communications
A1.3. – Distributed Systems
A1.3.1. – Web
A1.3.4. – Peer to peer
A1.3.5. – Cloud
A1.3.6. – Fog, Edge
A1.5. – Complex systems
A1.5.1. – Systems of systems
A1.5.2. – Communicating systems
A3.1.2. – Data management, quering and storage
A3.3.3. – Big data analysis
A3.4. – Machine learning and statistics
A3.5. – Social networks
A3.5.2. – Recommendation systems
A4.1. – Threat analysis
A4.8. – Privacy-enhancing technologies
A5.9. – Signal processing
A5.9.2. – Estimation, modeling
A5.9.4. – Signal processing over graphs
A5.9.6. – Optimization tools
A6.1.1. – Continuous Modeling (PDE, ODE)
A6.1.2. – Stochastic Modeling
A6.2.2. – Numerical probability
A6.2.3. – Probabilistic methods
A6.2.4. – Statistical methods
A6.2.6. – Optimization
A6.2.7. – HPC for machine learning
A6.4.1. – Deterministic control
A6.4.2. – Stochastic control
A6.4.6. – Optimal control
A7.1. – Algorithms
A7.1.1. – Distributed algorithms
A7.1.2. – Parallel algorithms

A7.1.4. – Quantum algorithms

A8.1. – Discrete mathematics, combinatorics

A8.2.1. – Operations research

A8.6. – Information theory

A8.8. – Network science

A8.9. – Performance evaluation

A8.11. – Game Theory

A9.2. – Machine learning

A9.2.1. – Supervised learning

A9.2.2. – Unsupervised learning

A9.2.3. – Reinforcement learning

A9.2.4. – Optimization and learning

A9.2.5. – Bayesian methods

A9.6. – Decision support

A9.7. – AI algorithmics

A9.9. – Distributed AI, Multi-agent

A9.11. – Generative AI

Other research topics and application domains

B2.3. – Epidemiology

B2.5.1. – Sensorimotor disabilities

B3.1. – Sustainable development

B3.1.1. – Resource management

B4. – Energy

B4.3.4. – Solar Energy

B4.4. – Energy delivery

B4.4.1. – Smart grids

B4.5.1. – Green computing

B6. – IT and telecom

B6.2. – Network technologies

B6.2.1. – Wired networks

B6.2.2. – wireless networks

B6.3.3. – Network Management

B6.3.4. – Social Networks

B6.4. – Internet of things

B6.6. – Embedded systems

B8.1. – Smart building/home

B9.2.1. – Music, sound

B9.5.1. – Computer science

B9.5.2. – Mathematics

Contents

1 Team members, visitors, external collaborators

Research Scientists

- Alain Jean-Marie [Team leader, INRIA, Senior Researcher, until Mar 2025]
- Giovanni Neglia [Team leader, INRIA, Senior Researcher, from Apr 2025, HDR]
- Sara Alouf [INRIA, Researcher, HDR]
- Eitan Altman [INRIA, Senior Researcher, HDR]
- Konstantin Avrachenkov [INRIA, Senior Researcher, HDR]
- Alain Jean-Marie [INRIA, Senior Researcher, from Apr 2025]
- Samir Medina Perlaza [INRIA, Researcher]
- Philippe Nain [INRIA, Emeritus, HDR]
- Giovanni Neglia [INRIA, Senior Researcher, until Mar 2025, HDR]

Faculty Member

- Michela Chessa [UNIV COTE D'AZUR, Associate Professor Delegation, from Sep 2025]

Post-Doctoral Fellows

- Khushboo Agarwal [INRIA, Post-Doctoral Fellow, until Oct 2025]
- Guodong Sun [INRIA, Post-Doctoral Fellow, from Apr 2025]
- Ying Zheng [INRIA, Post-Doctoral Fellow, from Dec 2025]

PhD Students

- Yaiza Bermudez [INRIA]
- Gaspard Gerard Philippe Berthelier [EDF R D, from Feb 2025]
- José Francisco Daunas Torres [Univ Sheffield, until Sep 2025]
- Ibtihal El Mimouni [NSP-SmartProfile, CIFRE]
- Louis Hauseux [UNIV COTE D'AZUR]
- Ahmad Nasser [NOKIA]
- Isidoor Pinillo Esquivel [INRIA, from Sep 2025]
- Julian Alfonso Santos Bustos [ORANGE, CIFRE]
- Adrien Sardi [NOKIA]
- Kyrylo Tymchenko [UNIV COTE D'AZUR, from Oct 2025]
- Jingye Wang [INRIA, from Sep 2025]
- Xufeng Zhang [INRIA]
- Xinying Zou [INRIA]

Interns and Apprentices

- Antonio Honsell [INRIA, Intern, from Jun 2025 until Aug 2025]
- Maxime Nicaise [INRIA, Intern, from Sep 2025]
- Pietro Tellarini [UNIV BOLOGNE, Intern, from Mar 2025 until Aug 2025]
- Zeev Weizmann [UNIV COTE D'AZUR, Intern, from May 2025 until Aug 2025]

Administrative Assistant

- Jane Desplanques [INRIA]

Visiting Scientist

- Daniel Richards Ravi Arputharaj [Freelance, until Feb 2025]

External Collaborator

- Patrick Brown [Freelance]

2 Overall objectives

Neo is an Inria project-team whose members are located in Sophia Antipolis (S. Alouf, K. Avrachenkov, G. Neglia, and S. M. Perlaza), in Avignon (E. Altman) and in Montpellier (A. Jean-Marie). E. Altman is also with the LINCS (Lab. for Information, Networking and Communication Sciences) in Paris. S. M. Perlaza is also with the ECE department at Princeton Univ., N.J. USA; and the Mathematics Department of the Univ. de la Polynésie française (Laboratoire GAATI), Faaa, Tahiti.

The team is positioned at the intersection of Operations Research and Network Science. By using the tools of Stochastic Operations Research, we model situations arising in several application domains, involving networking in one way or the other. The aim is to understand the rules and the effects in order to influence and control them so as to engineer the creation and the evolution of complex networks.

3 Research program

The problems studied in Neo involve generally optimization, dynamic systems or randomness, and often all at the same time. The techniques we use to tackle these problems are those of Stochastic Operations Research, Applied Probabilities and Information Theory.

Stochastic Operations Research is a collection of modeling, optimization and numerical computation techniques, aimed at assessing the behavior of man-made systems driven by random phenomena, and at helping to make decisions in such a context.

The discipline is based on applied probability and focuses on effective computations and algorithms. Its core theory is that of Markov chains over discrete state spaces. This family of stochastic processes has, at the same time, a very large modeling capability and the potential of efficient solutions. By "solution" is meant the calculation of some *performance metric*, usually the distribution of some random variable of interest, or its average, variance, etc. This solution is obtained either through exact "analytic" formulas, or numerically through linear algebra methods. Even when not analytically or numerically tractable, Markovian models are always amenable to "Monte-Carlo" simulations with which the metrics can be statistically measured.

An example of this is the success of classical Queueing Theory, with its numerous analytical formulas. Another important derived theory is that of the Markov Decision Processes, which allows to formalize *optimal* decision problems in a random environment. This theory allows to characterize the optimal decisions, and provides algorithms for calculating them.

Strong trends of Operations Research are: a) an increasing importance of multi-criteria multi-agent optimization, and the correlated introduction of Game Theory in the standard methodology; b) an increasing

concern of (deterministic) Operations Research with randomness and risk, and the consequent introduction of topics like Chance Constrained Programming and Stochastic Optimization. Data analysis is also more and more present in Operations Research: techniques from statistics, like filtering and estimation, or Artificial Intelligence like clustering, are coupled with modeling in Machine Learning techniques like Q-Learning.

4 Application domains

4.1 Network Science

Network Science is a multidisciplinary body of knowledge, principally concerned with the emergence of global properties in a network of individual agents. These global properties emerge from "local" properties of the network, namely, the way agents interact with each other. The central model of "networks" is the graph (of Graph Theory/Operations Research). Nodes represent the different entities managing information and taking decisions, whereas, links represent the fact that entities interact, or not. Links are usually equipped with a "weight" that measures the intensity of such interaction. Adding evolution rules to this quite elementary representation leads to dynamic network models, the properties of which Network Science tries to analyze.

A classical example of properties sought in networks is the famous "six degrees of separation" (or "small world") property: how and why does it happen so frequently? Another ubiquitous property of real-life networks is the Zipf or "scale-free" degree distribution. Some of these properties, when properly exploited, lead to successful business opportunities: just consider the PageRank algorithm of Google, which miraculously connects the relevance of some Web information with the relevance of the other information that points to it.

4.2 Network Engineering

In its primary acceptation, Network Science involves little or no engineering: phenomena are assumed to be "natural" and emerge without external interventions. However, the idea comes fast to intervene in order to modify the outcome of the phenomena. This is where NEO is positioned. Beyond the mostly descriptive approach of Network Science, we aim at using the techniques of Operations Research so as to engineer complex networks.

To quote two examples: controlling the spread of diseases through a "network" of people is of primarily interest for mankind. Similarly, controlling the spread of information or reputation through a social network is of great interest in the Internet. Precisely, given the impact of web visibility on business income, it is tempting (and quite common) to manipulate the graph of the web by adding links so as to drive the PageRank algorithm to a desired outcome.

Another interesting example is the engineering of community structures. Recently, thousands of papers have been written on the topic of community *detection* problem. In most of the works, the researchers propose methods, most of the time, heuristics, for detecting communities or dense subgraphs inside a large network. Much less effort has been put in the understanding of community formation process and even much less effort has been dedicated to the question of how one can influence the process of community formation, e.g. in order to increase overlap among communities and reverse the fragmentation of the society.

Our ambition is to reach an understanding of the behavior of complex networks that will make us capable of influencing or producing a certain property in a given network. For this purpose, we will develop families of models to capture the essential structure, dynamics, and uncertainty of complex networks. The "solution" of these models will provide the correspondence between metrics of interest and model parameters, thus opening the way to the synthesis of effective control techniques.

In the process of tackling real, very large size networks, we increasingly deal with large graph data analysis and the development of decision techniques with low algorithmic complexity, apt at providing answers from large datasets in reasonable time.

5 Social and environmental responsibility

5.1 Impact of research results

Some of NEO's research is devoted to environmental issues, either related to water management, or to the exploitation of renewable resources. The involvement in the Chile-funded project MICCHI (Section 10.1.3) aims at connecting these theoretical results with actual drought problems.

Several research actions direclty aim at reducing carbon/energy footprint in the IT sector. Some of these actions are carried out within the FedMalin Inria challenge (Section 10.4). In particular, in [42], NEO researchers propose algorithms to reduce the carbon footprint of cross-silo federated learning.

6 Highlights of the year

6.1 Awards

- Ibtihal El Mimouni was awarded Best Poster Award at SophIA Summit 2025 for the poster "A bandit approach for responsible email recommender systems."
- Louis Hauseux was finalist of the student competition Prix Pierre Laffitte 2025.
- Giovanni Neglia was awarded a Chair on "Distributed Machine Learning over the Internet" by the Interdisciplinary Institute for Artificial Intelligence 3IA Côte d'Azur, in the theme "Core Elements of AI."
- Giovanni Neglia was recognized as top reviewer for the Conference on Uncertainty in Artificial Intelligence (UAI), July 21-25, 2025, Brazil

6.2 Keynotes

Konstantin Avrachenkov was invited to give a plenary talk at the 20th Workshop on Modelling and Mining Networks (WAW 2025).

6.3 Courtesy Appointments

Samir Medina Perlaza was re-appointed "Visiting Research Collaborator" in the Department of Electrical and Computer Engineering at Princeton University for the academic year 2025–2026. He was also re-appointed "Associate Researcher" in the Laboratory of Algebraic Geometry and Applications to Information Theory (GAATI) at the Université de la Polynésie Française for the academic year 2025–2026.

7 Latest software developments, platforms, open data

7.1 Latest software developments

7.1.1 Marmote

Name: MARkovian MOdeling: The Environment

Keyword: Markov model

Functional Description: Marmote is a library for modeling with Markov chains. It is written in C++ with a Python interface. It consists in a reduced set of high-level abstractions for constructing state spaces, transition structures and Markov chains (discrete-time and continuous-time). It provides the ability of constructing hierarchies of Markov models, from the most general to the particular, and equip each level with specifically optimized solution methods. The current release features the library MarmoteMDP for modeling Markov Decision Processes and solving them.

This software was started within the ANR MARMOTE project: ANR-12-MONU-00019 under the name marmoteCore. Within the Marmote project, the code conforms the latest C++ standards and the library is available on multiple platforms via a conda distribution.

Release Contributions: Version 1.3.1 was released at the end of 2025. Together with previous version 1.3.0, it brought many improvements to the 'user experience'. To quote a few: generalized use of the logging facility, generalized use of the 'policy' feature, including in numerical algorithms, improved formatting of large objects (sets, matrices, ...). Technical improvements concern primarily the computation of hitting times in Markov chains, and the related phase-type distributions. The RLGL algorithm for computing stationary distributions, which is an exclusivity of Marmote, was enriched with an "adaptive relaxation" strategy.

A complete set of examples is available on the project's website, both for the C++ and the Python languages, and both for Marmote and MarmoteMDP, the library for manipulating Markov Decision Processes. For Python, a notebook-based tutorial demonstrates the principal functionalities of the library.

News of the Year: The software was awarded an Inria ADT (Action of Technological Development). Thanks to the help of software engineers of Inria's SED (Service Experimentation Development), versions 1.3.0 and 1.3.1 were developed during the year. See above for the improvements brought by these versions.

URL: https://marmote.gitlabpages.inria.fr/marmote/

Publications: hal-03770430v1, hal-03781620v1, hal-02395100v1, hal-04176076v1

Contact: Alain Jean-Marie

Participants: Alain Jean-Marie, Patrick Brown, Emmanuel Hyon

Partner: Université Paris Nanterre

8 New results

8.1 Stochastic Modeling

Participants: Sara Alouf, Eitan Altman, Konstantin Avrachenkov, Alain Jean-Marie, Kyrylo Tymchenko.

Stochastic modeling is a core methodological pillar of Neo: it provides principled abstractions to capture uncertainty, temporal variability, and randomness in networked systems, and to derive both qualitative insights and quantitative performance guarantees. Over the last year, our work combined foundational advances and application-driven modeling, spanning Markov processes and queueing systems, as well as stochastic analyses of data-stream algorithms, time-series models, and distributed storage mechanisms.

8.1.1 Markov Processes

Markov processes are a cornerstone of Neo's stochastic modeling work, providing a unifying framework for analyzing complex random dynamics in networked systems. Over the last year, we pursued complementary contributions spanning (i) fundamental properties of continuous-time Markov processes, (ii) decision-making under partial observability in Markovian environments, and (iii) software support for Markov modeling and numerical analysis through continued development of the Marmote library.

In [14], Konstantin Avrachenkov together with Flora Spieksma (Leiden University, Netherlands) studies the deviation matrix of denumerable state space, multi-chain Markov processes in continuous time. The deviation matrix is a measure of the cumulative deviation from the limiting probabilities and it plays an important role in many application domains. First, the authors provide several equivalent necessary and

sufficient conditions for the existence of the deviation matrix. Next, they study a relation between the deviation matrix and rank-one perturbations of Markov processes. Based on a rank-one perturbation, they derive a versatile formula for the deviation matrix and apply this formula to Markov processes with restart. Along the way, they establish several new properties of rank-one perturbed Markov processes. They feel that those properties can be useful in their own right.

In [28], Konstantin Avrachenkov, together with Madhu Dhiman and Kavitha Veerarun (IIT Bombay, India), studies Markov Decision Processes (MDPs) with intermittent state information, with periods of missing observations. Linear programming (LP) methods can play a crucial role in solving MDPs, in particular, with constraints. However, the resultant belief MDPs lead to infinite dimensional LPs, even when the original MDP is with finite state and action spaces. The verification of strong duality becomes non-trivial. This work investigates the conditions for no duality gap in average-reward finite Markov decision process with intermittent state observations. The authors first establish that in such MDPs, the belief MDP is unichain if the original Markov chain is recurrent. Furthermore, they establish strong duality of the problem, under the same assumption. Finally, the authors provide a wireless channel example, where the belief state depends on the last channel state received and the age of the channel state. The numerical results indicate interesting properties of the solution.

The development of the `Marmote` library has continued, led by Alain Jean-Marie and Emmanuel Hyon of Université Paris-Nanterre. The software was awarded an Inria ADT (Action of Technological Development), during which several software engineers of Inria's SED (Service Experimentation Development) helped improve the robustness and portability of the library. This year's improvements have been primarily targeted at a better 'user experience'. Through an adequate configuration, `Marmote` programmers have access to a fine control on the behavior of the program: its algorithms and its output. On the side of Markov modeling and numerical analysis, the principal improvement is about the computation of hitting times and the related phase-type distributions. Also, the RLGL algorithm, co-developed by Konstantin Avrachenkov and Patrick Brown, was enriched with a "adaptive relaxation" strategy for vector updates. Version 1.3.1 of `Marmote` has been released at the end of 2025.

8.1.2 Queueing Systems

Queueing systems with multiple queues can be controlled in different ways depending on who makes the decision. In one setting, each arriving customer chooses which queue to join, leading to load-balancing policies that aim to reduce congestion and delay. In another setting, a single server chooses which queue to serve (and when to switch), yielding polling-type control problems where stability and delay trade-offs depend on the switching rule and its overhead. Over the last year, we contributed to both perspectives.

Join-the-shortest-queue and its variants have often been used in solving load balancing problems. The aim of such policies is to minimize the average system occupation, e.g., the customer's system time. In [36], Eitan Altman, together with Andrea Fox and Francesco De Pellegrini (Avignon Université, France), Arnob Ghosh (New Jersey Institute of Technology, USA), and Ness Shroff (The Ohio State University, USA), extends the load balancing setting to include constraints that may be imposed, e.g., due to the communication network. First, the authors cast the problem in the framework of constrained MDPs: this permits them to address both action-dependent constraints, such as bandwidth limitation, and state-dependent constraints, such as minimum queue utilization. Hence, unlike the state-of-the-art approaches in load balancing, they derive new policies that satisfy the constraints while minimizing system occupancy. Extensive numerical simulations evaluate the policies' performance under various system settings.

In [27], Konstantin Avrachenkov, together with Kousik Das, Kavitha Veeraruna (IIT Bombay, India) and Vartika Singh (University of Colorado, US), considers a polling system with two queues, where a single server is attending the queues in a cyclic order and requires non-zero switching times to switch between the queues. The aim is to identify a fairly general and comprehensive class of Markovian switching policies that renders the system stable, potentially a class of policies that can cover the Pareto frontier related to individual-queue-centric performance measures like the stationary expected number of waiting customers in each queue. For instance, such a class of policies has been identified recently for a polling system near the fluid regime (with large arrival and departure rates), and the authors aim to include that class. They also aim to include a second class that facilitates switching between the queues at the instance the occupancy in the opposite queue crosses a threshold and when that in the visiting queue is below a threshold (this inclusion facilitates design of 'robust' polling systems). Towards this, the authors consider a class of two-phase

switching policies, which includes the above mentioned classes. In the maximum generality, the policies can be represented by eight parameters, while two parameters are sufficient to represent the aforementioned classes. The authors provide simple conditions to identify the sub-class of switching policies that ensure system stability. By numerically tuning the parameters of the proposed class, they illustrate that the proposed class can cover the Pareto frontier for the stationary expected number of customers in the two queues.

8.1.3 Approximate Counting

Count-Min Sketch with Conservative Updates (CMS-CU) is a memory-efficient hash-based algorithm used to estimate the occurrences of items within a data stream. CMS-CU employs d hash functions and a total of m counters arranged in d rows, with each hash function mapping an item to one counter per row. In [29], Younes Ben Mazziane (Avignon Université, France) and Othmane Marfoq (Meta, USA)[1] study a similar algorithm, which they refer to as CU-S, where d hash functions map each item to d distinct counters in a single array of size m. Specifically, they present an analytical method to quantify the trade-off between memory usage and the Counting Error (CE) of an item in CU-S, defined as the discrepancy between the estimated and actual number of its occurrences. The first result of this paper shows that items absent from the stream experience the highest CE. They refer to their error as the Unseen items Error (UE). The second result is an upper bound on UE of any stream, expressed in terms of the Reference Error (RE), which is the UE of a reference stream where each item appears at most once. They also identify streams for which this bound is provably tight. The direct computation of RE involves dealing with an infinite state space Markov process similar to the Balls and Bins model with the Power of Choice, but more difficult to handle analytically. Instead, they construct two finite state space Markov processes, parametrized by a positive integer g, that bound the original chain and yield efficiently computable lower and upper bounds on RE. Increasing g narrows the gap between these bounds and enlarges the state space to $\binom{m+g-d}{g}$, thus increasing the computation time. For $d = m - 1$, $g = 1$, and as $m \to \infty$, they prove that the lower and upper bounds coincide. Their bounds are accurate for small values of g and for more general values of m and d, as shown by numerical computation. Finally, simulations on various streams confirm that bounding RE in this manner enables efficient and accurate computation of UE.

8.1.4 Characterization of Integrated Autoregressive Time Series

Integrated autoregressive time series are widely used to model non-stationary dynamics in economics, finance, and other applications, where observed variables exhibit persistent trends and become approximately stationary only after differencing. In [21], Konstantin Avrachenkov, together with with Phil Howlett (University of South Australia, Australia), Brendan Beare (University of Sydney, Australia), Massimo Franchi (Sapienza Universita di Roma, Italy), and John Boland (University of South Australia, Australia), proves an extended Granger-Johansen Representation Theorem (GJRT) for finite—or infinite—order integrated autoregressive time series on Banach space. The authors assume only that the resolvent of the autoregressive polynomial for the series is analytic on and inside the unit circle except for an isolated singularity at unity. If the singularity is a pole of finite order, the time series is integrated of the same order. If the singularity is an essential singularity, the time series is integrated of order infinity. When there is no deterministic forcing, the value of the series at each time is the sum of an almost surely convergent stochastic trend, a deterministic term depending on the initial conditions and a finite sum of embedded white noise terms in the prior observations. This is the extended GJRT. In each case the original series is the sum of two separate autoregressive time series on complementary subspaces—a singular component which is integrated of the same order as the original series and a regular component which is not integrated. The extended GJRT applies to all integrated autoregressive processes irrespective of the spatial dimension, the number of stochastic trends and cointegrating relations in the system and the order of integration.

8.1.5 Coding Schemes for Distributed Storage

Distributed storage over heterogeneous networks offers a cost-effective alternative to cloud storage. Most existing systems rely on erasure coding to ensure data durability despite significant peer unreliability. However,

[1] Both authors were members of the Neo team. Since most of the research was carried out while Younes Ben Mazziane was still a team member, he also listed his Neo affiliation.

these solutions face the fundamental trade-off between data safety and storage overhead, both of which are of great importance to users. In addition, resource consumption, particularly CPU and bandwidth usage, remains a major concern. In an ongoing research, Sara Alouf and Kyrylo Tymchenko, together with Frederic Giroire and Stephane Perennes (Coati Inria team), investigate novel coding schemes which combine reliability and fast coding operations. They use Markovian modeling to evaluate the performance of the coding schemes designed and compare them with classical erasure coding and locally repairable codes. This work has been submitted for publication and is under review.

8.2 Theory of Learning

Participants: Konstantin Avrachenkov, Yaiza Bermudez, Louis Hauseux, Alain Jean-Marie, José Francisco Daunas Torres, Ibtihal El Mimouni, Samir Medina Perlaza, Giovanni Neglia, Xinying Zou.

8.2.1 Information-theoretic Foundations for Statistical Learning

Over the last year, we developed measure-theoretic and information-theoretic tools that clarify how learning objectives and regularization relate to probability measures.

In [52, 47], Yaiza Bermudez and Samir Medina Perlaza, together with Gaetan Bisson (Université de la Polynésie française, France) and Iñaki Esnaola (University of Sheffield, United Kingdom), present rigorous statements and formal proofs for both foundational and advanced folklore theorems on the Radon-Nikodym derivative. The cases of conditional and marginal probability measures are carefully considered, which leads to an identity involving the sum of mutual and lautum information suggesting a new interpretation for such a sum.

In [23, 56], Samir Medina Perlaza, together with Gaetan Bisson (Université de la Polynésie française, France), presents closed-form expressions for the variation of the expectation of a given function due to changes in the probability measure (probability distribution drifts). These expressions unveil interesting connections with Gibbs probability measures, information projections, Pythagorean identities for relative entropy, mutual information, and lautum information.

In [54, 31], José Francisco Daunas Torres and Samir Medina Perlaza, together with Iñaki Esnaola (University of Sheffield, United Kingdom), introduce the dual formulation of empirical risk minimization (ERM) with f-divergence regularization (ERM-fDR). The authors connect the solution of the dual optimization problem to the ERM-f DR to the notion of normalization function introduced as an implicit function. This dual approach leverages the Legendre-Fenchel transform and the implicit function theorem to provide a nonlinear ODE expression to the normalization function. Furthermore, the nonlinear ODE expression and its properties provide a computationally efficient method to calculate the normalization function of the ERM-fDR solution under a mild condition.

In [18], José Francisco Daunas Torres and Samir Medina Perlaza, together with Iñaki Esnaola (University of Sheffield, United Kingdom) and H. Vincent Poor (Princeton University, USA), analyze the effect of relative entropy asymmetry in the context of empirical risk minimization with relative entropy regularization (ERM-RER). The authors consider two regularizations: (a) the relative entropy of the measure to be optimized with respect to a reference measure (Type-I ERM-RER); and (b) the relative entropy of the reference measure with respect to the measure to be optimized (Type-II ERM-RER). The main result is the characterization of the solution to the Type-II ERM-RER problem and its key properties. By comparing the well-understood Type-I ERM-RER with Type-II ERM-RER, the effects of entropy asymmetry are highlighted. The analysis shows that in both cases, regularization by relative entropy forces the support of the solution to collapse into the support of the reference measure, introducing a strong inductive bias that negates the evidence provided by the training data. The authors show finally that Type-II regularization is equivalent to Type-I regularization with an appropriate transformation of the empirical risk function.

In [58], Xinying Zou and Samir Medina Perlaza, together with Iñaki Esnaola (University of Sheffield, United Kingdom), address the challenge of learning models that stay reliable under distribution shifts when only finite training data is available. The authors propose a novel training-dependent minimax problem to design learning algorithms that are robust to the worst-case datagenerating distribution. For the ambiguity

sets, they construct Kullback-Leibler (KL) divergence neighborhoods on both model and data distributions. They obtain an analytical solution to this minimax problem, referred to as the robust learner. They show that the robust learner follows a Gibbs distribution, in which the prior can be chosen as any baseline learning algorithm to be robustified. Such a robust learner minimizes the worst-case generalization gap within a KL-divergence neighborhood of unseen new data and it provides a smaller generalization error compared with the baseline learning algorithm Q. Under certain conditions, the robust learner also guarantees a smaller expected testing error than Q. They also provide a training-dependent PAC-Bayes bound on the robust learner's performance on unseen data. Closed-form expressions for generalization error and expected loss of the robust learner are given in terms of mutual information, lautum information, and KL-divergence. As a by-product, they show that the proposed minimax problem admits a two-player zero-sum game formulation, for which a unique Nash equilibrium exists. This enhances the understanding of learning algorithm robustification. Numerical experiments validate the applicability of the results and the benefits of robustification.

8.2.2 Unlearning

Machine unlearning 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 [45], Giovanni Neglia, together with Martin van Waerebeke, Kevin Scaman (Argo Inria team), and Marco Lorenzi (Epione Inria team), analyzes the efficiency of unlearning methods. The authors 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, they 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.

8.2.3 Learning Strategies for Email Marketing

Email marketing is increasingly criticized due to ethical concerns, as bulk email campaigns often result in spam, reduced engagement, and negative user experiences. In addition, there is increasing awareness of the environmental impact, as these large-scale campaigns contribute to carbon emissions. To address these issues, in [35], Ibtihal El Mimouni and Konstantin Avrachenkov introduce QWIC-Fair (Q-learning Whittle Index with Context and Fairness), an algorithm that operates within a Contextual Restless Multi-Armed Bandit framework. QWIC-Fair leverages implicit feedback to learn the dynamics of user interactions and thus target users with relevant content. In this model, each user represents an arm of the bandit, evolving as a Markov Decision Process that captures state transitions reflecting their interactions with email contents, while accounting for contextual information. The algorithm also incorporates a fairness constraint to ensure balanced selection and to avoid repetitive targeting of the same users. The experiments conducted, using synthetic and real-world data, show that QWIC-Fair outperforms existing email marketing approaches.

In [34], Ibtihal El Mimouni and Konstantin Avrachenkov introduce DQWIC, a novel algorithm that combines Deep Reinforcement Learning and Whittle index theory within the Contextual Restless Multi-Armed Bandit framework for the discounted criterion. The authors design DQWIC to learn in evolving environments typical of real-world applications, such as recommender systems, where user preferences and environmental dynamics evolve over time. In particular, they apply DQWIC to the problem of optimizing email recommendations, where it tackles the dual challenges of enhancing content relevance and reducing spam messages, thereby addressing ethical concerns related to intrusive emailing. The algorithm leverages two neural networks: a Q-network for approximating action-value functions and a Whittle-network for estimating Whittle indices, both of which integrate contextual features to inform decision-making. In addition, the inclusion of context allows them to handle many heterogeneous users in a scalable way. The learning process occurs through a two time scale stochastic approximation, with the Q-network updated frequently to minimize the loss between predicted and target Q-values, and the Whittle-network updated on a slower time scale. To evaluate its effectiveness, they conducted experiments in partnership with Smartprofile, a company

specializing in digital marketing. Their results, derived from both synthetic and real-world data, show that DQWIC outperforms existing email marketing baselines.

8.2.4 Graph Clustering

Graph clustering (aka community detection) is one of the fundamental problems in data science which consists of partitioning graph nodes into disjoint communities. In [19, 40], Konstantin Avrachenkov, together with Lucas Lopes Felipe and Daniel Sadoc Menasché (UFRJ, Brazil), presents a game-theoretic perspective on the Constant Potts Model (CPM) for partitioning graphs into disjoint communities, emphasizing its efficiency, robustness, and accuracy. Efficiency: CPM is reinterpreted as a potential hedonic game by decomposing its global Hamiltonian into local utility functions, where the local utility gain of each agent matches the corresponding increase in global utility. Leveraging this equivalence, the authors prove that local optimization of the CPM objective via better-response dynamics converges in pseudo-polynomial time to an equilibrium partition. Robustness: The authors introduce and relate two stability criteria, namely a strict criterion based on a novel notion of robustness—requiring nodes to simultaneously maximize neighbors and minimize non-neighbors within communities—and a relaxed utility function based on a weighted sum of these objectives, controlled by a resolution parameter. Accuracy: In community tracking scenarios, where initial partitions are used to bootstrap the Leiden algorithm with partial ground-truth information, the experiments reveal that robust partitions yield higher accuracy in recovering ground-truth communities.

In [38], Konstantin Avrachenkov and Louis Hauseux, together with Nahuel Soprano-Loto (Mathnet Inria team), study Gibbs distributions with competing interactions and propose a higher-order extension of the Swendsen-Wang dynamics that incorporates triangular bonds. The new dynamics preserves the same stationary distribution, alleviates frustration, and yields markedly better sampling. When applied to a synthetic Euclidean-graph community-detection benchmark, the proposed algorithm outperforms existing methods.

8.2.5 Density-based Clustering

Many clustering algorithms are based on density estimates in $\mathbb{R}^d$. Building geometric graphs on a dataset $X \subset \mathbb{R}^d$ is an elegant way to achieve this. In fact, the connected components of a geometric graph match exactly with the high-density clusters of the 1-Nearest Neighbor density estimator or Single-Linkage algorithm. In [20], Louis Hauseux and Konstantin Avrachenkov, together with Josiane Zerubia (Ayana Inria team), analyze and generalize the classical Single-Linkage clustering algorithm, which performs hierarchical clustering by iteratively merging the two closest clusters. Single-Linkage and its robust version are still widely used in modern clustering techniques like the state-of-the-art HDBSCAN. Single-Linkage can be understood from three perspectives: (i) it conducts persistent analysis on geometric graphs; (ii) it identifies high-density clusters using the 1-Nearest Neighbor density estimator; and (iii) it is implemented via the minimum spanning tree of the data. The authors extend Single-Linkage to higher-order interactions by replacing geometric graphs with hypergraphs and introducing a stricter notion of connected components, named K-polyhedra. Specifically, for $K = 2$, their method employs "triangle connectivity". They prove that K-polyhedra correspond to high-density clusters of the K-Nearest Neighbors density estimator. In practice, this approach is implemented by identifying a minimum K-tree. The authors also introduce original geometric optimizations for efficiently computing the 2-generalization of Single-Linkage in low-dimensional Euclidean spaces. Experimental results demonstrate that even when $K = 2$ is used, the proposed method already surpasses the state-of-the-art clustering methods on synthetic and real-world datasets.

8.2.6 Conjectural Learning

In the context of dynamic games in which players have limited information, Alain Jean-Marie, together with Mabel Tidball (INRAE, Montpellier, France) and Tania Jiménez (Avignon Université, France), introduces in [60] the family of Conjectural Learning procedures. With Conjectural Learning, agents form conjectures about what the opponent will play, as a function of their action or some state variable and may revise these conjectures at each interaction. The authors prove general properties of Conjectural Learning procedures, comparing their steady-states to the cooperative solution of a corresponding static game with complete information. They then specify simple functional forms of conjectures and analyze the resulting dynamic

systems, in terms of steady states and their relation with Pareto solutions of the complete information framework. They illustrate the transient and stationary behavior of these Conjectural Learning procedures in the Fish War model of Levhari and Mirman.

In a related forthcoming publication, the same authors endow the otherwise selfish and myopic agents with a certain degree of forward-looking behavior, in the form of a subjective valuation of future stocks, incorporated in their utility. They analyze the interaction of three specific conjectural learning procedures and this short-range forward-looking behavior, in the case of a management problem of groundwater resource with two symmetric players. The performance is measured with both the asymptotic stock and the total discounted gain. They conduct numerical experiments using data from the La Mancha aquifer in Spain. They conclude that learning processes can be Pareto improving at each period. However, this result depends on the valuation that agents place on future stocks, the inertia of the learning process and the initial value of conjectures, which must be appropriately chosen. In particular, if the valuation of the future stock is too small, the performance of these schemes is worse than that of the farsighted Nash solution, in terms of both profits and resource levels.

8.3 Distributed Learning

Participants: Yaiza Bermudez, Samir Medina Perlaza, Giovanni Neglia, Daniel Richards Ravi Arputharaj.

Federated Learning (FL) is a distributed machine learning paradigm in which multiple clients collaboratively train a shared model while keeping their raw data local. FL typically takes two forms: cross-device FL, involving a very large number of intermittently available edge devices (e.g., smartphones, IoT), and cross-silo FL, involving a smaller number of reliable organizations (e.g., hospitals, banks) with larger local datasets. Over the last year, our research tackled key bottlenecks in federated learning by (i) advancing theoretical understanding, with new guarantees for FL under Markovian data streams and an information-theoretic framework that yields closed-form characterizations of generalization error in FL; (ii) developing scalable collaboration mechanisms, including Bayesian one-shot FL that exploits multimodality in local objectives and communication-efficient peer-graph cooperation to help clients identify similar peers; (iii) improving learning under realistic constraints, through carbon-aware client selection and scheduling that leverages slack time and fairness to reduce emissions while preserving accuracy under heterogeneity; and (iv) strengthening the privacy-evaluation toolbox with new, highly effective reconstruction and attribute-inference attacks, including the first model-based attribute inference attack tailored to federated regression and a geometric attack enabling perfect reconstruction of very large batches without prior distributional knowledge.

8.3.1 Theoretical Insights on Federated Learning Algorithms

A sound theoretical understanding is essential to characterize when FL works, how fast it converges, and how well it generalizes under the statistical dependencies and heterogeneity that arise in practice. Over the last year, we contributed to this agenda with complementary results on optimization under dependent client data streams and on generalization error in FL through an information-theoretic lens.

Most theoretical and empirical FL studies rely on the assumption that clients have access to pre-collected data sets, with limited investigation into scenarios where clients continuously collect data. In many real-world applications, particularly when data is generated by physical or biological processes, client data streams are often modeled by non-stationary Markov processes. Unlike standard i.i.d. sampling, the performance of FL with Markovian data streams remains poorly understood due to the statistical dependencies between client samples over time. In [39], Giovanni Neglia, together with Tan-Khiem Huynh, Malcolm Egan, and Jean-Marie Gorce (Maracas Inria team), investigates whether FL can still support collaborative learning with Markovian data streams. Specifically, the authors analyze the performance of Minibatch SGD (Stochastic Gradient Descent), Local SGD, and a variant of Local SGD with momentum. They answer affirmatively under standard assumptions and smooth non-convex client objectives: the sample complexity is proportional to the inverse of the number of clients, with a communication complexity comparable to the i.i.d. scenario. However, the sample complexity for Markovian data streams remains higher than for i.i.d. sampling. Their analysis is validated via experiments with real pollution monitoring time series data.

In [53], Yaiza Bermudez and Samir Medina Perlaza, together with Iñaki Esnaola (University of Sheffield, United Kingdom) and H. Vincent Poor (Princeton University, USA), characterize the generalization error of FL systems through a novel statistical framework. Central to this framework is the concept of a meta-federated learning algorithm, defined as a probability measure over a client's local models conditioned on the datasets of all participating clients. By means of this abstraction, the authors state several fundamental properties of FL systems and derive closed-form expressions for the generalization error. More specifically, they extend to FL the method of gaps, originally introduced for non-federated settings, and obtain closed-form expressions for the generalization error in terms of classical information measures, including relative entropy, mutual information, and lautum information. A central role in these new expressions is played by specific Gibbs probability measures (Gibbs algorithms). More importantly, they reveal that the challenge of evaluating the generalization error in FL is reduced to two distinct tasks: (a) measuring the dependence of client model choices on the datasets of all clients; and (b) distinguishing the meta-federated learning algorithm from a Gibbs algorithm trained solely on local data. Through these findings, they establish new links between generalization in FL, mismatched hypothesis testing, Shannon's information measures, and Pythagorean identities for the generalization error.

8.3.2 Scalable Algorithms

Scalability is a central requirement for federated and distributed learning systems: algorithms must remain effective when the number of clients grows, communication is constrained, and the amount of data per client varies widely. Over the last year, we investigated scalability both by reducing the number of communication rounds in FL and by enabling large numbers of clients to collaborate with limited per-client resources.

One-Shot FL enables multiple clients to cooperatively learn a global model in a single round of communication with a central server. In [44], Giovanni Neglia, together with Jacopo Talpini and Marco Savi (University of Milano-Bicocca, Italy), analyzes the One-Shot FL problem through the lens of Bayesian inference. The authors propose FedBEns, an algorithm that leverages the inherent multimodality of local loss functions to find better global models. Their algorithm leverages a mixture of Laplace approximations for the clients' local posteriors, which the server then aggregates to infer the global model. They conduct extensive experiments on various datasets, demonstrating that the proposed method outperforms competing baselines that typically rely on unimodal approximations of the local losses.

While FL clients may want to collaborate because they do not have enough data to learn an accurate model on their own, collaboration also introduces a bias–variance trade-off when local data distributions differ. A key challenge is for each client to identify clients with similar distributions while learning the model, a problem that remains largely unresolved. In [37], Giovanni Neglia, together with Franco Galante and Emilio Leonardi (Politechnic University of Turin, Italy), focuses on a particular instance of this challenge, where each client collects samples from a real-valued distribution over time to estimate its mean. Existing algorithms face impractical per-client space and time complexities (linear in the number of clients $|A|$). To address scalability challenges, the authors propose a framework where clients self-organize into a graph, allowing each client to communicate with only a selected number of peers r. They propose two collaborative mean estimation algorithms: one employs a consensus-based approach, while the other uses a message-passing scheme, with complexity $O(r)$ and $O(r \log |A|)$, respectively. They establish conditions for both algorithms to yield asymptotically optimal estimates and they provide a theoretical characterization of their performance.

8.3.3 Carbon-aware Distributed Learning

Training large-scale machine learning models incurs substantial carbon emissions. Since FL distributes computation across geographically dispersed clients, it offers a natural framework to exploit regional and temporal variations in Carbon Intensity (CI) through carbon-aware participation and scheduling.

In [42], Daniel Richards Ravi Arputharaj and Giovanni Neglia, together with Charlotte Rodriguez (Accenture Labs, France) and Angelo Rodio (Linköping University, Sweden), investigate how to reduce emissions in FL through carbon-aware client selection and training scheduling. They first quantify the emission savings of a carbon-aware scheduling policy that leverages slack time—permitting a modest extension of the training duration so that clients can defer local training rounds to lower-carbon periods. They then examine the performance trade-offs of such scheduling, which stem from statistical heterogeneity among clients, selection bias in participation, and temporal correlation in model updates. To leverage these trade-offs,

they construct a carbon-aware scheduler that integrates slack time, α-fair carbon allocation, and a global fine-tuning phase. Experiments on real-world CI data show that their scheduler outperforms slack-agnostic baselines, achieving higher model accuracy across a wide range of carbon budgets, with especially strong gains under tight carbon constraints.

8.3.4 Privacy Attacks to Federated Learning

Although FL keeps raw data local, the training phase can leak sensitive information through exchanged updates, enabling reconstruction attacks that recover either *attributes* of the training data or the data themselves. Over the last year, we strengthened the empirical toolbox for privacy evaluation in FL with two attacks that substantially broaden the scope and effectiveness of reconstruction in realistic settings.

First, while attribute inference attacks (AIA) have been widely studied for classification, their impact on *federated regression* had remained largely unexplored. In [32], Giovanni Neglia, together with Francesco Diana, Chuan Xu, and Frederic Giroire (Coati Inria team), Othmane Marfoq (Meta, USA) and Eoin Thomas (Amadeus, France), addresses this gap by proposing novel model-based AIAs specifically designed for regression tasks in FL environments. The authors' approach considers scenarios where adversaries can either eavesdrop on exchanged messages or directly interfere with the training process. They benchmark their proposed attacks against state-of-the-art methods using real-world datasets. The results demonstrate a significant increase in reconstruction accuracy, particularly in heterogeneous client datasets, a common scenario in FL. The efficacy of their model-based AIAs makes them better candidates for empirically quantifying privacy leakage for federated regression tasks.

Second, existing data reconstruction attacks often rely on assumptions about clients' data distributions and their effectiveness degrades sharply when batch sizes exceed a few tens of samples. In [33], Giovanni Neglia, together with Francesco Diana, André Nusser, and Chuan Xu (Coati Inria team), introduces a novel data reconstruction attack that overcomes these limitations. The method leverages a new geometric perspective on fully connected layers to craft malicious model parameters, enabling the perfect recovery of arbitrarily large data batches in classification tasks without any prior knowledge of clients' data. Through extensive experiments on both image and tabular datasets, the authors demonstrate that their attack outperforms existing methods and achieves perfect reconstruction of data batches two orders of magnitude larger than the state of the art.

8.4 Game Theory and Applications

Participants: Khushboo Agarwal, Eitan Altman, Konstantin Avrachenkov.

8.4.1 Games with Irrational Players

The classical game theory considers rational players and proposes Nash equilibrium (NE) as the solution. However, real-world scenarios rarely feature rational players; instead, players make inconsistent and irrational decisions. Often, irrational players exhibit herding behavior by simply following the majority. In [26], Khushboo Agarwal and Konstantin Avrachenkov, together with Veeraruna Kavitha and Raghupati Vyas (IIT Bombay, India), consider a mean-field game with α-fraction of rational players and the rest being herding-irrational players. For such a game, the authors introduce a novel concept of equilibrium named α-Rational NE (in short, α-RNE). They extensively analyze the α-RNEs and their implications in games with two actions. Due to herding-irrational players, new equilibria may arise, and some classical NEs may be deleted. They establish that the rational players are not harmed but benefit from the presence of irrational players. More interestingly, in some examples, the rational players attain higher utility (under α-RNE) than even the social optimal utility (in the classical setting), by leveraging upon the herding behavior of irrational players. Surprisingly, the irrational players may also benefit by not being rational. They observe that irrational players do not lose compared to some classical NEs for participation and bandwidth-sharing games. Importantly, in bandwidth-sharing game, the irrational players also receive utility near social optimal utility. Such examples indicate that it may sometimes be 'rational' to be irrational.

In [13], Khushboo Agarwal and Konstantin Avrachenkov, together with Veeraruna Kavitha and Raghupati Vyas (IIT Bombay, India), consider one more realistic behavioral game dynamics where the players choose actions in a turn-by-turn manner and exhibit two prominent behavioral traits—α-fraction of them are myopic players who strategically choose optimal actions against the empirical distribution of the previous plays, while the rest exhibit herding behavior by choosing the most popular action till then. The utilities are realized for all, at the end of the game, and each player gets to play only once. The analysis focuses on scenarios when players encounter two possible choices, common in applications like participation games (e.g., crowd-sourcing) or minority games. To begin with, the authors derive the almost sure mean-field limits of such dynamics. The proof is constructive and progressively narrows down the potential limit set and finally establishes the existence of a unique limit for almost all sample paths. The authors argue that the dynamics at the limit is captured by a differential inclusion (and not the usual ordinary differential equation) due to the discontinuities arising from the switching behavioral choices. It is noteworthy that the presented methodology can be easily modified to analyze the avoid-the-crowd behavior, in place of herding behavior. The work is concluded with two interesting examples, named participation game and routing game, which encapsulate several real-life scenarios.

8.4.2 Strategic Queueing with Information Cost

Consider an M/M/1-type queue where joining attains a known reward, but a known waiting cost is paid per time unit spent queueing. In the 1960s, Naor showed that any arrival optimally joins the queue if its length is less than a known threshold. Yet acquiring knowledge of the queue length often brings an additional cost, e.g., website loading time or data roaming charge. Therefore, their model presents any arrival with three options: join blindly, balk blindly, or pay a known inspection cost to make the optimal joining decision by comparing the queue length to Naor's threshold. In a recent paper, Hassin and Roet-Green prove that a unique Nash equilibrium always exists and classify regions where the equilibrium probabilities are non-zero. In [17], Konstantin Avrachenkov and Eitan Altman, together with Jake Clarkson (National Highways, United Kingdom), complement these findings with new closed-form expressions for the equilibrium probabilities in the majority of cases. Further, Hassin and Roet-Green show that minimizing inspection cost maximizes social welfare. Envisaging a queue operator choosing where to invest, the authors of [17] compare the effects of lowering inspection cost and increasing the queue-joining reward on social welfare. They prove that the former dominates and that the latter can even have a detrimental effect on social welfare.

8.4.3 Kelly Mechanism

The Kelly mechanism is a proportional allocation auction widely adopted in decentralized resource allocation systems to share an infinitely divisible resource among competing agents.

In [41], Eitan Altman, together with Cleque Marlain Mboulou-Moutoubi, Younes Ben Mazziane, and Francesco De Pellegrini (Avignon Université, France), analyze the sequential game the Kelly allocation induces when agents have α-fair utilities and behave strategically. The authors' main result proves that synchronous best-response updates drive bids to the unique Nash equilibrium at a linear rate for $\alpha \in \{0, 1, 2\}$. Extensive simulations reveal that best-response dynamics reach equilibrium significantly faster than previously proposed no-regret learning algorithms.

When agents are aware of the allocation mechanism, their interactions form a game. The properties of its Nash equilibria are well understood under the simplifying assumption of unbounded budgets. In [30], Eitan Altman, together with Younes Ben Mazziane, Cleque-Marlain Mboulou-Moutoubi, and Francesco De Pellegrini (Avignon Université, France), analyzes the game in a more realistic budget-constrained setting, motivated by its optimality in terms of the liquid price of anarchy. Specifically, the authors establish a sufficient condition for the uniqueness of the Nash equilibrium and design a distributed sequential learning procedure that provably converges to the equilibrium. In particular, their sufficient condition holds when the payoff functions of the agents are of the proportional fair type in the allocated fraction. Finally, extensive numerical experiments shed light on the interplay between the heterogeneity of the payoff functions and the agents' budgets.

8.4.4 Applications to Energy Markets

In [22, 59], Eitan Altman, together with Hélène Le Cadre, Mathis Guckert (INOCS Inria team) and Mandar Datar (CEA-Leti, France), consider a peer-to-peer electricity market modeled as a private network game, where end users minimize their cost by computing their demand and controllable generation. Their nominal demand constitutes sensitive information that they might want to keep private. The authors prove that the private network game admits a unique variational equilibrium, which depends on the private information of all end users. Thus, to update their strategy, end users rely on randomized readings. They introduce a data aggregator, which aims to learn the end users' private information, while remunerating them depending on the quality of their readings. Using performative prediction, they define a decision-dependent game explicitly taking into account the distribution shift caused by the endusers' hidden ability. The decision-dependent game coincides with a Stackelberg game when the end users' hidden abilities are best responses. Further, the market robustness can be quantified by evaluating the efficiency loss as the difference between the social cost in the performatively stable equilibrium and the optimum. The authors show that under mild assumptions, the performatively stable equilibrium can be found by distributed and sequential variants of the repeated stochastic gradient method while they propose a two-timescale stochastic approximation method to learn Stackelberg equilibrium. Finally, they formulate the data aggregator's optimal contract design as a bilevel optimization problem that they cast as a more tractable non-linear non-convex optimization problem which can be solved using simulated annealing. Simulations on small and large scale problem instances illustrate the results.

8.5 Applications in Telecommunications

Participants: Sara Alouf, Samir Medina Perlaza, Philippe Nain, Giovanni Neglia, Xufeng Zhang.

Over the last year, we pursued several research contributions with direct applications to telecommunications systems. They span scalable online learning methods for caching, fundamental trade-offs in simultaneous information and energy transmission, and the analysis of covert communications under sequential detection constraints.

8.5.1 Scalable Online Learning for Caching

Online learning algorithms address sequential decision-making problems where, at each round, a learner selects an action, observes feedback (e.g., a loss or reward), and updates its decision rule on the fly. A key appeal of these methods is their worst-case, distribution-free guarantee: they ensure *sublinear regret* with respect to the best fixed decision in hindsight, without requiring statistical assumptions on how requests are generated. This makes online learning particularly attractive for caching, where traffic can be non-stationary and hard to predict. Over the last year, we studied how to make such no-regret caching policies practical at scale, by addressing constraints that arise in real deployments—limited observability of requests, and reduced computational and memory budgets of online algorithms operating over very large catalogs.

Most existing algorithms involve computationally expensive operations and require knowledge of all past requests, which may not be feasible in practical scenarios such as femtocaching, where a base station (BS) jointly decides the content of many edge caches and visibility of all requests at the BS requires constant communication between these caches and the BS. To capture this constraint, in [15], Sara Alouf and Giovanni Neglia, together with Younes Ben Mazziane (Université Avignon, France) and Francescomaria Faticanti (ENS, Lyon, France), study a single cache problem under a more restrictive setting, that they refer to as the Bernoulli Partial Observability (BPO) model, in which the caching policy only observes a request with probability p, reflecting the fraction of requests forwarded from the edge caches to the BS in the femtocaching example. They propose a policy, based on the classic online learning algorithm Follow-the-Perturbed-Leader (FPL), that achieves an asymptotically optimal regret bound of $O(\sqrt{CT/p})$ under BPO in $O(1)$ amortized time complexity as T goes to infinity, where C is the cache size and T is the number of requests. Moreover, they show that their policy extends to bipartite caching albeit with a sublinear α-regret for $\alpha = 1 - 1/e$ and a

higher computational cost. The experimental evaluation compares the proposed solution with classic caching policies and validates the proposed approach using both synthetic and real-world request traces.

Caching policies based on online learning algorithms often suffers high computation complexity, which hinders their practical adoption. In [16], Giovanni Neglia and Xufeng Zhang, together with Damiano Carra (University of Verona, Italy), introduce a new variant of the gradient-based online caching policy that achieves groundbreaking logarithmic computational complexity relative to catalog size, while also providing regret guarantees. This advancement allows the authors to test the policy on large-scale, real-world traces featuring millions of requests and items—a significant achievement, as such scales have been beyond the reach of existing policies with regret guarantees. The regret guarantees and the low complexity are also maintained in cases where items have non-uniform sizes. To the best of their knowledge, the proposed solution is the only low-complexity no-regret policy for such a case, and their experimental results demonstrate for the first time that the regret guarantees of gradient-based caching policies offer substantial benefits in practical scenarios.

Online learning algorithms provide robust performance in caching problems but require substantial memory to store per-file historical data, limiting their scalability to large-catalog systems. To overcome this challenge, in [46], Xufeng Zhang, Sara Alouf, and Giovanni Neglia propose a dimensionality reduction algorithm based on the Follow-the-Perturbed-Leader framework and the Johnson–Lindenstrauss lemma. Their method significantly reduces memory consumption while preserving sublinear regret, making it well-suited for caching under resource constraints. Experiments on both synthetic and real-world traces demonstrate its advantages over other memory-efficient approaches.

8.5.2 Simultaenous Information and Energy Transmission in Wireless Networks

Simultaneous wireless information and power transfer aims to use the same radio-frequency signal to convey information and deliver energy that can be harvested by the receiver. A central challenge is the inherent trade-off between reliable communication and sufficient energy transfer, especially in practical regimes with finite block-length codes and finite input constellations.

In [25], Samir Medina Perlaza, together with Sadaf Ul Zuhra and H. Vincent Poor (Princeton University, USA) and Mikael Skoglund (KTH Royal Institute of Technology, Sweden), characterize the trade-offs between information and energy transmission over an additive white Gaussian noise channel in the finite block-length regime with finite channel input symbols. The authors characterize these trade-offs in the form of inequalities involving the information transmission rate, energy transmission rate, decoding error probability (DEP) and energy outage probability (EOP) for a given finite block-length code. The first set of results identify a set of necessary conditions that a given code must satisfy for simultaneous information and energy transmission. They propose a novel method for constructing a family of codes that can satisfy a target information rate, energy rate, DEP and EOP. Finally, achievability results identify the set of tuples of information rate, energy rate, DEP and EOP that can be simultaneously achieved by the constructed family of codes.

8.5.3 Covert Communications

Covert communications study how a transmitter (or adversary) can operate while keeping its very presence hard to detect, typically by ensuring that any statistical test at a warden cannot reliably distinguish "signal-present" from "signal-absent." This viewpoint naturally connects to sequential detection: if an observer monitors a process for a change, a covert adversary may adapt its behavior so that the change remains barely detectable and detection is delayed as much as possible under a prescribed false-alarm constraint.

In [57, 24], Philippe Nain, together with Amir Reza Ramtin and Don Towsley (University of Massachusetts Amherst, USA), investigate the problem of covert quickest change detection in a continuous-time setting, where a Brownian motion experiences a drift change at an unknown time. Unlike classical formulations, the authors consider a covert adversary who adjusts the post-change drift $\mu = \mu(\gamma)$ as a function of the false alarm constraint parameter γ, with the goal of remaining undetected for as long as possible. Leveraging the exact expressions for the average detection delay (ADD) and average time to false alarm (AT2FA) known for the continuous-time CuSum procedure, the authors rigorously analyze how the asymptotic behavior of ADD evolves as $\mu(\gamma) \to 0$ with increasing γ. Their results reveal that classical detection delay characterizations no longer hold in this regime. They derive sharp asymptotic expressions for the ADD under various convergence rates of $\mu(\gamma)$, identify precise conditions for maintaining covertness, and characterize the total damage inflicted by the adversary. They show that the adversary achieves maximal damage when the drift scales as

$\mu(\gamma) = \Theta(1/\sqrt{\gamma})$, marking a fundamental trade-off between stealth and impact in continuous time detection systems.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

NEO has contracts with EDF (see §9.1.1), Hivenet (see §9.1.2), Nokia (see §9.1.3), NSP SmartProfile (see §9.1.4), and Orange Labs (see §9.1.5).

9.1.1 Cifre contract with EDF "Automated and responsible recommendation systems for digital marketing" (February 2025 – January 2028)

Participants: Gaspard Gerard Philippe Berthelier, Samir Medina Perlaza, Giovanni Neglia.

- Contractor: EDF
- Collaborators: Etienne Le Naour, Tahar Nabil, Richard Niamke

In collaboration with EDF, in the framework of Gaspard Gerard Philippe Berthelier's PhD thesis, we develop federated learning methods tailored to time series forecasting in the energy sector. Our project explores personalized FL to allow each client to learn specialized models while benefiting from collaborative training. We also investigate the use of public or non-sensitive data to improve representation learning and model initialization. This work aims to deliver privacy-preserving, adaptive FL methods suitable for real-world industrial time series applications.

9.1.2 Inria challenge with Hivenet (September 2025–August 2029)

Participant: Giovanni Neglia.

Project Acronym: CUPSELI

Project Title: Collaborative Unified Platform for a Scalable and Efficient Learning Infrastructure

Duration: September 2025–August 2029

Abstract: The CUPSELI challenge ambitions to push the boundaries of distributed computing and artificial intelligence to offer a sovereign, secure, and sustainable alternative to centralized cloud solutions. It mobilizes a large scientific community, bringing together 11 Inria research teams from six research centers: Rennes, Bordeaux, Lorraine, Côte d'Azur, Lyon, and Paris. The CUPSELI challenge addresses three major technological challenges: sustainable and distributed AI across diverse computing hardware, distributed and secure computing, and large-scale distributed computing. Two PhD students will be recruited during 2026 in the framework of this project.

9.1.3 Contracts with Nokia
Inria Challenge LearnNet

Participants: Ahmad Nasser, Giovanni Neglia, Ying Zheng.

Project Acronym: LearnNet

Project Title: Learning Networks

Duration: January 2024 - December 2027

Abstract: While machine learning is revolutionizing entire sectors of the digital economy and scientific research, its robust deployment in digital infrastructures raises many questions. The challenge Learning Networks (LearnNet) explores new avenues of research at the intersection of the fields of networks and learning. This challenge has two complementary objectives: rethinking the design of network protocols to serve machine learning applications, and exploring how learning can improve network management. Thus the LearnNet challenge studies the growing entanglement between the challenges of large-scale learning and network design.

LearnNet is a research project that spans 8 Inria research teams. The teams from Nokia are AIRL and NSSR.

Collaborators: Fabio Pianese, Chung Shue Chen (Nokia)

Publications in 2025: [39]

Inria Challenge SmartNet

Participants: Sara Alouf, Adrien Sardi.

Project Acronym: SmartNet

Project Title: AI Methods for Smart Network Management

Duration: January 2024 - December 2027

Abstract: The advent of virtualization, combined with the power of AI, has brought new opportunities in network management. To effectively address the challenges that come with this paradigm shift, the SmartNet project is dedicated to exploring the transformative potential of AI methods in enabling smart network management. The project strategically focuses on two key areas: slice provisioning and causal analysis of network malfunctions. The project is dedicated to the development of cutting-edge methods to respond effectively to the growing complexity of networks, particularly in multi-domain scenarios.

SmartNet is a research project that spans 5 Inria research teams. The teams from Nokia are MLS and NSSR.

Cifre contract "Energy efficient management/provisioning of Generative AI services for 6G networks " (January 2025 – December 2027) related to the Cifre thesis contract of A. Sardi.

Collaborator : Marie-Line Alberi-Morel (Nokia)

9.1.4 Cifre contract with NSP-SmartProfile "Automated and responsible recommendation systems for digital marketing" (August 2022 – May 2026)

Participants: Konstantin Avrachenkov, Ibtihal El Mimouni.

- Contractor: NSP-SmartProfile
- Collaborators: Hervé Baile, Julien Musso

SmartProfile is a marketing platform that allows to collect, to enhance and to analyze marketing data. Digital marketing campaigns continue to expand across all digital channels and media. The 'mass marketing' strategies implemented by most companies show limits in terms of performance and acceptance by clients, as well as in terms of their impact on the environment. In opposite to these practices, we believe that current technologies, particularly in terms of Artificial Intelligence (AI), should make marketing interactions more efficient and virtuous. Through this research project, we want to create an alternative solution to mass marketing by switching to an intelligent, automated and eco-responsible system, which will support the heterogeneity of data and the diversity of sectors, and whose purpose is to recommend the best content by determining the most relevant target and taking into account the communication constraints. This contract complements the Cifre thesis of Ibtihal El Mimouni. Relevant publications: [34, 35].

9.1.5 Cifre contract with Orange Labs "Analytical modeling of large-scale wireless networks integrating RIS" (September 2023 – September 2026)

Participants: Eitan Altman, Konstantin Avrachenkov, Julian Alfonso Santos Bustos.

- Contractor: Orange Labs
- Collaborators: Jean-Marc Kelif

A Reconfigurable Intelligent Surface (RIS) is a programmable surface structure that allows one to control the reflection of electromagnetic (EM) waves by changing the electric and magnetic properties of the surface. In the absence of RIS, short wavelentghs signals as in 5G, are subject to a huge attenuation when there is no direct line of sight channel. Within our collaboration, we shall evaluate and optimize the position of RIS.

This contract complements the Cifre thesis of J. Santos.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Inria associate team not involved in an IIL or an international program

LION

Participants: Khushboo Agarwal, Eitan Altman, Konstantin Avrachenkov, Samir Medina Perlaza.

Title: Learning In Operations and Networks

Duration: 2022 – 2025

Coordinator: Kavitha Veeraruna

Partners: Indian Institute of Technology Bombay (India)

Inria contact: Eitan Altman

Summary: Artificial Intelligence (AI) has affected all walks of life. We study its application in various domains like

1. Learning and Control in Healthcare: Our aim is to use novel AI methodologies, to predict the results of possible actions of involved decision-makers, using the available data.
2. Dual Learning Algorithms in wireless networks: We aim to develop learning algorithms for beam alignment in 5G Wireless networks to maintain high rates. We propose to use Age of Information as a metric.

3. Distributed and reinforcement learning: We will develop and analyze Deep Q-Network based learning algorithms and analyze their performance.

Publications in 2025: [13, 26, 27, 28].

10.1.2 STIC/MATH/CLIMAT AmSud projects

GSA

Participants: Konstantin Avrachenkov, Alain Jean-Marie.

Title: Graph Spectra and Applications

Program: MATH-AmSud

Duration: January 1, 2023 – December 31, 2025

Local supervisor: Konstantin Avrachenkov

Partners:

- V.Trevisan, L.E. Allem, A.M. França, C. Hoppen, UFRGS, UFF (Brésil)
- A. Pastine, Universidad Nacional de San Luis (Argentine)
- L. Medina, Universidad de Antofagasta (Chile)

Inria contact: Konstantin Avrachenkov

Latin America contact: Vilmar Trevisan

Summary: The present project proposes the establishment of a network of collaboration among Argentina, Chile, Brazil, and France, using the strength of 5 mathematics groups of 5 different institutions. The research topic of the proposal is Algebraic Graph Theory, an important and modern area of discrete mathematics. The proposal is structured in such a way that the training of highly qualified human resources and research activities are intertwined, this will ensure the generation of new knowledge in a relevant scientific area and leave permanent ties of collaboration between the different research groups beyond the completion of the project.

10.1.3 Participation in other International Programs

MICCHI

Participant: Alain Jean-Marie.

Title: Mecanismos e Incentivos Contra la Crisis HIdrica

Funding: Chile's Agencia Nacional de Investigación y Desarollo (ANID)

Partners:

- Universidad de O'Higgins, Chile (coordinator)
- Universidad de Chile, Chile
- Universidad de Barcelona, Spain
- INRAe, France

Duration: January 2024 - December 2025

Summary: The water crisis caused by global warming is one of the most important problems affecting agricultural regions such as the sixth chilean Region of O'Higgins. The main objective of this project is to investigate different mechanisms for the allocation of water resources in times of scarcity.

10.2 International research visitors

10.2.1 Visits of international scientists

Other international visits to the team: research stay

Kalle Alaluusua

Status PhD student

Institution of origin: Aalto University

Country: Finland

Dates: 14 to 18 April 2025

Context of the visit: collaboration on the topic of geometric network clustering

Mobility program/type of mobility: research stay

Damiano Carra

Status Full Professor

Institution of origin: University of Verona

Country: Italy

Dates: 6 to 10 October 2025

Context of the visit: collaboration on low-regret online learning

Mobility program/type of mobility: research stay

Kousic Das

Status Post-doc

Institution of origin: IIT Bombay

Country: India

Dates: 17 to 21 November 2025

Context of the visit: collaboration on stochastic processes, queueing theory and game theory (associated team LION)

Mobility program/type of mobility: research stay

Diego Goldsztajn

Status Post-doc

Institution of origin: University ORT Uruguay

Country: Uruguay

Dates: 27 October to 2 November 2025

Context of the visit: collaboration on Markov decision processes

Mobility program/type of mobility: research stay

Lorenzo Gregoris

Status PhD student

Institution of origin: Eindhoven University of Technology

Country: The Netherlands

Dates: 29 April to 28 May 2025

Context of the visit: collaboration on a new approach on Red Light Green Light Method for Solving Large Markov Chains

Mobility program/type of mobility: research stay

Vinay Kumar

Status Post-doc

Institution of origin: Eindhoven University of Technology

Country: The Netherlands

Dates: 14 to 18 April 2025

Context of the visit: collaboration on the topic of geometric network clustering

Mobility program/type of mobility: research stay

Lasse Leskela

Status Professor

Institution of origin: Aalto University

Country: Finland

Dates: 14 to 18 April 2025

Context of the visit: collaboration on the topic of geometric network clustering

Mobility program/type of mobility: research stay

Taisiia Morozova

Status PhD student

Institution of origin: Uppsala University

Country: Sweden

Dates: 2 to 28 June 2025

Context of the visit: collaboration on mean-field and clustering methods for 5G cellular systems

Mobility program/type of mobility: research stay

Angelo Rodio

Status Post-doc

Institution of origin: Linkoping University

Country: Sweden

Dates: 20 to 26 March 2025

Context of the visit: collaboration on semi-decentralized federated learning

Mobility program/type of mobility: research stay

Rajesh Sundaresan

Status Professor

Institution of origin: Indian Institute of Science, Bangalore

Country: India

Dates: 21 to 29 October 2025

Context of the visit: collaboration on stochastic perturbation of a dynamic system

Mobility program/type of mobility: research stay

Jacopo Talpini

Status Post-doc

Institution of origin: University of Milano-Bicocca

Country: Italy

Dates: 26 to 28 May 2025

Context of the visit: collaboration on federated learning in a single communication

Mobility program/type of mobility: research stay

Alexander Van Werde

Status Post-doc

Institution of origin: Munster University

Country: Germany

Dates: 11 to 18 October 2025

Context of the visit: collaboration on clustering of sparse geometric graphs

Mobility program/type of mobility: research stay

Kavitha Veeraruna

Status Professor

Institution of origin: IIT Bombay

Country: India

Dates: 17 to 21 November 2025

Context of the visit: collaboration on stochastic processes, queueing theory and game theory (associated team LION)

Mobility program/type of mobility: research stay

Uri Yechiali

Status Emeritus Professor

Institution of origin: Tel Aviv University

Country: Israel

Dates: 21 to 24 April 2025

Context of the visit: collaboration on strategic queues in a random environment

Mobility program/type of mobility: research stay

Other international visits to the team: internship

Antonio Honsell

Status intern (BSc)

Institution of origin: Bocconi University

Country: Italy

Dates: 16 June to 29 August 2025

Context of the visit: working on federated learning and privacy preservation

Mobility program/type of mobility: internship

Pietro Tellarini

Status intern (master/eng)

Institution of origin: University of Bologna

Country: Italy

Dates: 10 March to 31 August 2026

Context of the visit: working on similarity caching for text-to-image streaming models

Mobility program/type of mobility: internship

10.2.2 Visits to international teams

Research stays abroad

Konstantin Avrachenkov

Visited institution: Indian Institute of Technology (Bombay)

Country: India

Dates: 20 January - 2 February 2025

Context of the visit: Visit to Associate Team LION

Mobility program/type of mobility: research stay

Visited institution: University of Liverpool

Country: UK

Dates: 25-30 March 2025

Context of the visit: Visit to Dr. Alexey Piunovskiy

Mobility program/type of mobility: research stay

Visited institutions: UFRJ and UFRGS / University of Buenos Aires / ORT University Montevideo

Country: Brazil / Argentina / Uruguay

Dates: 24 November - 14 December 2025.

Context of the visit: Research visit, seminars

Mobility program/type of mobility: MATH-AmSud GSA

Yaiza Bermudez

Visited institutions: University of Cambridge / Sheffield University

Country: UK

Dates: 13-23 May 2025

Context of the visit: collaborations with Albert Guillen i Fabregas, Iñaki Esnaola

Mobility program/type of mobility: research stay

Visited institution: Universidad Carlos III de Madrid

Country: Spain

Dates: 15-19 December 2025

Context of the visit: collaboration with Tobias Koch

Mobility program/type of mobility: research stay

Alain Jean-Marie

Visited institution: Universidad O'Higgins

Country: Chile

Dates: 12 to 21 April 2025

Context of the visit: project MICCHI (§10.1.3)

Mobility program/type of mobility: research stay

Samir Medina Perlaza

Visited institution: GAATI Mathematic Laboratory, Université de la Polynésie française.

Country: Polynésie française

Dates: 31 January - 15 February 2025

Context of the visit: collaboration with Gaetan Bisson

Mobility program/type of mobility: research stay

Visited institution: Centre for Mathematical Sciences, University of Cambridge

Country: UK

Dates: 25-27 February 2025

Context of the visit: collaboration with Albert Gillen i Fabregas

Mobility program/type of mobility: seminar

Visited institution: Universidad Carlos III de Madrid

Country: Spain

Dates: 2-5 April 2025

Context of the visit: collaboration with Tobias Koch

Mobility program/type of mobility: seminar

Giovanni Neglia

Visited institution: Univ. Palermo

Country: Italy

Dates: 12 to 16 February 2025

Context of the visit: collaboration on federated learning

Mobility program/type of mobility: research stay, seminar

Visited institution: Univ. Federal Rio de Janeiro (UFRJ)

Country: Brazil

Dates: 17 to 27 October 2025

Context of the visit: collaboration with Daniel Figueiredo, Daniel Sadoc Menasche and Giulio Iacobelli on distributed AI

Mobility program/type of mobility: research stay, seminar

10.3 European initiatives

10.3.1 Horizon Europe

dAIEDGE

Participants: Sara Alouf, Alain Jean-Marie, Giovanni Neglia.

dAIEDGE project on cordis.europa.eu

Title: A network of excellence for distributed, trustworthy, efficient and scalable AI at the Edge

Duration: From September 1, 2023 to August 31, 2026

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- NVISO SA (NVISO), Switzerland
- UBOTICA TECHNOLOGIES LIMITED, Ireland
- UNIVERSITE COTE D'AZUR, France
- CSEM CENTRE SUISSE D'ELECTRONIQUE ET DE MICROTECHNIQUE SA - RECHERCHE ET DEVELOPPEMENT (CSEM), Switzerland
- VARJO TECHNOLOGIES OY, Finland
- FRAUNHOFER GESELLSCHAFT ZUR FORDERUNG DER ANGEWANDTEN FORSCHUNG EV (Fraunhofer), Germany
- THALES SIX GTS FRANCE SAS (THALES SIX GTS France), France
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (CEA), France
- INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM (IMEC), Belgium
- SOFIA UNIVERSITY ST KLIMENT OHRIDSKI (UNISOFIA), Bulgaria
- IDRYMA TECHNOLOGIAS KAI EREVNAS (FOUNDATION FOR RESEARCH AND TECHNOLOGYHELLAS), Greece
- FUNDACION INSTITUTO INTERNACIONAL DE INVESTIGACION EN INTELIGENCIA ARTIFICIAL Y CIENCIAS DE LA COMPUTACION, Spain
- BONSEYES COMMUNITY ASSOCIATION, Switzerland
- SINTEF AS (SINTEF), Norway
- DEUTSCHES FORSCHUNGSZENTRUM FUR KUNSTLICHE INTELLIGENZ GMBH (DFKI), Germany
- DEUTSCHES ZENTRUM FUR LUFT - UND RAUMFAHRT EV (DLR), Germany
- FUNDACION CENTRO DE TECNOLOGIAS DE INTERACCION VISUAL Y COMUNICACIONES VICOMTECH (VICOM), Spain
- FUNDINGBOX ACCELERATOR SP ZOO (FBA), Poland

- BLEKINGE TEKNISKA HOGSKOLA (BTH), Sweden
- EIDGENOESSISCHE TECHNISCHE HOCHSCHULE ZUERICH (ETH Zürich), Switzerland
- SYNOPSYS INTERNATIONAL LIMITED (SYNOPSYS), Ireland
- UNIVERSIDAD DE CASTILLA - LA MANCHA (UCLM), Spain
- SAFRAN ELECTRONICS & DEFENSE, France
- VERSES GLOBAL BV, Netherlands
- HAUTE ECOLE SPECIALISEE DE SUISSE OCCIDENTALE (HES-SO), Switzerland
- UNIVERSITY OF GLASGOW, United Kingdom
- STMICROELECTRONICS SRL, Italy
- Aegis Rider AG (Aegis Rider), Switzerland
- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS (CNRS), France
- KATHOLIEKE UNIVERSITEIT LEUVEN (KU Leuven), Belgium
- UNIVERSITA DEGLI STUDI DI MODENA E REGGIO EMILIA (UNIMORE), Italy
- THE UNIVERSITY OF EDINBURGH (UEDIN), United Kingdom
- HIPERT SRL, Italy
- UNIVERSIDAD DE SALAMANCA (USAL), Spain
- SORBONNE UNIVERSITE, France
- CENTRE D'EXCELLENCE EN TECHNOLOGIES DE L'INFORMATION ET DE LA COMMUNICATION (CETIC), Belgium

Inria contact: Giovanni Neglia

Coordinator: Alain Pagani (DFKI)

Summary: The proposal focuses on the Next Generation AI topic of the call HORIZON-CL4-2022-HUMAN-02-02. The vision of dAIEDGE Network of Excellence (NoE) is to strengthen and support the development of the dynamic European edge and distributed Artificial Intelligence (AI) ecosystem as an essential ingredient in the growth and competitiveness of European industrial sectors. The dAIEDGE Network aims to reinforce the research and innovation value chains to accelerate the digital and green transitions through advanced edge AI technologies, applications, and innovations, building on Europe's existing assets and industrial strengths. In parallel, it will fortify the edge AI research and industrial communities through technological developments beyond state of the art and become a dependable and strategic pillar for the European AI Lighthouse. This will be achieved by mobilizing and connecting the European AI and edge AI constituency, the relevant stakeholders, European partnerships, and projects, to provide roadmaps, guidelines and trends supporting the next-generation edge AI technologies. The key aim is to support and ensure rapid development, market uptake and open strategic sovereignty for Europe in the critical technologies for distributed edge AI (hardware, software, frameworks, tools). The dAIEDGE NoE will play a catalyst role in building a solid edge AI virtual network of research facilities and laboratories to benefit the European research and industrial community. The NoE multidisciplinary concept provide an arena for matchmaking, exchanging ideas, tools, and services, by bringing together the leading research centers, AI-on-demand platforms, digital innovation hubs, AI projects and initiatives. The ultimate goal for the dAIEDGE NoE is to support Europe to become a global center of excellence with unique human-centered edge AI competence addressing the social and economic challenges and the needs of the citizens and society.

Publications in 2025: [16, 32, 33, 37, 39, 42, 44, 46].

FINALITY

Participants: Sara Alouf, Giovanni Neglia, Isidoor Pinillo Esquivel, Jingye Wang.

FINALITY project on cordis.europa.eu

Title: saFe learNIng for lArge scaLe InTerconnected sYstems

Duration: From March 1, 2025 to February 28, 2029

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- SISTEMAS AVANZADOS DE TECNOLOGIA SA, Spain
- UNIVERSIDAD PUBLICA DE NAVARRA, Spain
- UNIVERSITE COTE D'AZUR, France
- ORANGE SA (Orange), France
- KUNGLIGA TEKNISKA HOEGSKOLAN (KTH), Sweden
- FUNDACION IMDEA NETWORKS (IMDEA NETWORKS), Spain
- AVIGNON UNIVERSITE, France
- ERICSSON AB (EAB), Sweden
- TELEFONICA INNOVACION DIGITAL SL, Spain
- NOKIA SPAIN SA, Spain
- NOKIA NETWORKS FRANCE, France
- SAFRAN PASSENGER INNOVATIONS GERMANY GMBH (SPI), Germany
- THE CYPRUS INSTITUTE (THE CYPRUS INSTITUTE), Cyprus
- UNIVERSIDAD CARLOS III DE MADRID (UC3M), Spain
- TECHNISCHE UNIVERSITEIT DELFT (TU Delft), Netherlands

Inria contact: Giovanni Neglia

Coordinator: Francesco De Pellegrini (Avignon Université, France)

Summary: FINALITY evolves the theoretical computer science curriculum focusing on the mastery of prompt and safe learning techniques for interconnected systems. The trainee team will develop and integrate innovative methodological tools specialized for AI-intensive resource allocation, particularly in the context of large-scale critical infrastructures for communication and computing. They will combine AI methods that are safe by respecting system boundaries and are prompt in adapting to the environmental changes. Throughout their research training, the FINALITY candidates will prioritize the principles of fairness and computational parsimony of AI methods. The FINALITY doctoral team will be supported by a world-class team of academic and industrial advisors, who work routinely on all the tools used in AI-based RA, advancing their theoretical foundations and their application in the industrial domain. They possess extensive experience in training doctoral students, and an excellent track record of joint research activities across the consortium. International exposure and dissemination are ensured by an extra-EU supervisory board

Publications in 2025: [16, 33, 39, 42, 44, 46].

10.4 National initiatives

NF-FOUNDS PC9 PEPR 5G

Participants: Khushboo Agarwal, Eitan Altman, Samir Medina Perlaza, Guodong Sun.

Project Acronym: NF-FOUNDS

Project Title: Networks of the Future - Foundations of Future Communications Networks

Program: ANR-22-PEFT-0010

Coordinator: CEA (Dmitri Kténas), CNRS (Serge Verdeyme), IMT (Daniel Koffman)

Duration: 2023 - 2030

Other Partners: EURECOM

Summary: The 5G network and the networks of the future represent a key issue for French and European industry, society and digital sovereignty. This is why the French government has decided to launch a dedicated national strategy. One of this strategy's priority ambitions is to produce significant public research efforts so the national scientific community contributes fully to making progress that clearly responds to the challenges of 5G and the networks of the future. In this context, the CNRS, the CEA and the Institut Mines-Télécom (IMT) are co-leading the '5G' acceleration PEPR to support upstream research into the development of advanced technologies for 5G and the networks of the future. Neo is involved in the theme "Networks and Telecommunications" and more specifically in the targeted projet 9 (PC9) Foundations of Future Communications Networks (FOUNDS).

Publications in 2025: [13, 18, 23, 25, 26, 31, 47, 52, 53, 54, 56].

ANR PARFAIT

Participants: Eitan Altman, Samir Medina Perlaza, Xinying Zou.

Project Acronym: PARFAIT

Project Title: Planning And leaRning For AI-Edge compuTing

Coordinator: Avignon Univ.

Duration: October 2021 - September 2025

Other Partners: Conservatoire National des Arts et Métiers (CNAM), Univ. Savoie Mont Blanc (USMB)

Summary: The PARFAIT project develops theoretical foundations for distributed and scalable resource allocation schemes on edge computing infrastructures tailored for AI-based processing tasks. Algorithmic solutions will be developed based on the theory of constrained, delayed, and distributed Markov decision processes to account for edge service orchestration actions and quantify the effect of orchestration policies. Furthermore, using both game and team formulations, the project will pave the way for a theory of decentralized orchestration, a missing building block necessary to match the application quest for data proximity and the synchronization problems that arise when multiple edge orchestrators cooperate under local or partial system view. Finally, to achieve efficient online edge service orchestration, such solutions will be empowered with reinforcement learning techniques to define a suit of orchestration algorithms able to at once adapt over time to the applications' load and cope with the uncertain information available from AI-based applications' footprints.

Publications in 2025: [18, 23, 31, 36, 47, 52, 53, 54, 56, 58].

Inria Challenge FedMalin

Participant: Giovanni Neglia.

Project Acronym: FedMalin

Project Title: FEDerated MAchine Learning over the INternet

Coordinator: Giovanni Neglia and Aurélien Bellet (PREMEDICAL Inria team)

Duration: November 2022 - November 2026

Summary: In many use-cases of Machine Learning (ML), data is naturally decentralized: medical data is collected and stored by different hospitals, crowdsensed data is generated by personal devices, etc. Federated Learning (FL) has recently emerged as a novel paradigm where a set of entities with local datasets collaboratively train ML models while keeping their data decentralized.

FedMalin is a research project that spans 10 Inria research teams and aims to push FL research and concrete use-cases through a multidisciplinary consortium involving expertise in ML, distributed systems, privacy and security, networks, and medicine. We propose to address a number of challenges that arise when FL is deployed over the Internet, including privacy and fairness, energy consumption, personalization, and location/time dependencies. FedMalin will also contribute to the development of open-source tools for FL experimentation and real-world deployments, and use them for concrete applications in medicine and crowdsensing. The FedMalin Inria Challenge is supported by Groupe La Poste, sponsor of the Inria Foundation.

Publications in 2025: [33, 37, 39, 42, 44].

DIAMOND

Participants: Yaiza Bermudez, Samir Medina Perlaza.

Project Acronym: DIAMOND

Project Title: Data-Injection Attacks in Supervised Machine Learning Systems

Funding Agency: Agence de l'Innovation de Défense (AID)

Principal Investigators: Samir Medina Perlaza and Iñaki Esnaola (University of Sheffield)

Duration: November 2024 - October 2027

Summary: This project aims at studying the impact on the generalization capabilities of federated learning systems of data-injection attacks (DIA) in the context of military applications. A DIA refers to any modification on the local training datasets aiming to tamper with the global performance. The focus is on the special class of Stealth DIAs (S-DIA), which exhibit the lowest-probability of detection. Before and after such attacks, training datasets exhibit empirical probabilities that are sufficiently close in relative entropy, which makes the probability of attack detection arbitrarily close to zero. The project funds the PhD scholarship of Yaiza Bermudez.

Publications in 2025: [47, 52, 53].

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organization

Steering committee chair, steering committee member

- Eitan Altman is
  - Member of the steering committee and founder of the "Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)";
  - Member of the steering committee and co-founder of the "Workshop on Networking Games Control and Optimization (NetGcoop)";
  - Member of the steering committee of the "International Conference on Performance Evaluation Methodologies and Tools (ValueTools)."

General chair, scientific chair

- Sara Alouf was General Co-Chair of the 36th Intl. Teletraffic Congress (ITC 36), held in Trondheim, Norway, June 2-5, 2025.
- Konstantin Avrachenkov was TPC Co-Chair and a member of the organizing committee of the IFIP WG 7.3 Performance 2025 – the 43rd International Symposium on Computer Performance, Modeling, Measurements and Evaluation, held in Amsterdam, The Netherlands, November 11-13, 2025.

Member of the conference program committees

- AAAI Conference on Artificial Intelligence (AAAI 2025), February 25-March 4, 2025, Philadelphia, Pennsylvania, USA (Samir Medina Perlaza);
- ACM SIGMETRICS 2025, *Winter* TPCs, June 9-13, 2025, Stony Brook, New York, USA (Konstantin Avrachenkov);
- ACM SIGMETRICS 2026, *Summer* and *Fall* TPC, June 8-12, 2026, Ann Arbor, Michigan, United States (Sara Alouf, Konstantin Avrachenkov);
- Annual Conference on Artificial Intelligence and Statistics (AISTATS), May 2-5, 2026, Tangier, Morocco (Samir Medina Perlaza, Giovanni Neglia);
- Conference on Game Theory and AI for Security (GameSec 2025), October 13-15, 2025, Athens, Greece (Konstantin Avrachenkov);
- Conference on Uncertainty in Artificial Intelligence (UAI 2025), July 21-25, 2025, Brazil (Giovanni Neglia—top reviewer, Konstantin Avrachenkov);
- European Conference on Networks and Communications & 6G Summit (EuCNC & 6G Summit 2025), June 3-6, 2025, Poznań, Poland (Samir Medina Perlaza);
- European Wireless 2025, October 27-29, 2025, Sophia-Antipolis, France (Samir Medina Perlaza);
- France's International Conference on Complex Systems (FRCCS 2025), May 21-23, 2025, Bordeaux, France, (Konstantin Avrachenkov, Alain Jean-Marie);
- IEEE Global Communications Conference (GLOBECOM 2025), December 8-12, 2025, Taipei, Taiwan (Samir Medina Perlaza);
- IEEE International Conference on Communications (ICC 2025), June 8-12, 2025, Montreal, Canada (Samir Medina Perlaza);

- IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm 2025), September 29-October 2, 2025, Toronto, Canada (Samir Medina Perlaza);
- IEEE International Conference on Communications in China (ICCC 2025), August 10-13, 2025, Shanghai, China (Samir Medina Perlaza);
- IEEE International Conference on Computer Communications (INFOCOM 2026), May 18-21, 2026, Tokyo, Japan (Sara Alouf);
- IEEE International Conference on Computing, Networking and Communications (ICNC 2025), February 17-20, 2025, Honolulu, Hawaii, USA (Samir Medina Perlaza);
- IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN 2025), May 26-29, 2025, Barcelona, Spain (Samir Medina Perlaza);
- IEEE International Symposium on Information Theory (ISIT 2025), June 22–27, 2025, Ann Arbor, Michigan, USA (Samir Medina Perlaza);
- IEEE Virtual Conference on Communications (VCC 2025), November 4-6, 2025, Virtual Conference (Samir Medina Perlaza);
- IEEE Wireless Communications and Networking Conference (WCNC 2025), March 24-27, 2025, Milan, Italy (Samir Medina Perlaza);
- International Conference on Complex Networks and their Applications, December 9-11, 2025, Binghamton, US (Konstantin Avrachenkov);
- International Conference on Distributed Computing and Intelligent Technology (ICDCIT 2025), January 8-11, 2025, Bhubaneswar, Odisha, India (Samir Medina Perlaza);
- International Conference on Machine Learning (ICML), July 13-19, 2025, Vancouver, Canada (Giovanni Neglia);
- International Conference of Networks, Games, Control and Optimization (NETGCOOP 2025), October 8-10, 2025, Bilbao, Spain (Khushboo Agarwal);
- International Symposium on Computer Performance, Modeling, Measurements and Evaluation (Performance 2025), November 11-13, 2025, Amsterdam, The Netherlands (Sara Alouf, Khushboo Agarwal);
- International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt 2025), May 26-29, 2025, Linköping, Sweden (Konstantin Avrachenkov);
- International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2025), October 21-23, 2025, Paris, France (Alain Jean-Marie);
- International Teletraffic Congress (ITC 36), 2-6 June 2025, Trondheim, Norway (Alain Jean-Marie);
- Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD 2025), June 10-13, 2025, Sydney, Australia (Konstantin Avrachenkov);
- SIAM Conference on Data Mining (SDM 2025), May 1-3, 2025, Alexandria Virginia, US (Konstantin Avrachenkov);
- Workshop on MAthematical performance Modeling and Analysis (MAMA 2025), June 13, 2025, Stony Brook, New York United States (Alain Jean-Marie, Philippe Nain);
- Workshop on Modelling and Mining Networks (WAW 2025), June 30 - July 3, 2025, Vilnius, Lithuania (Konstantin Avrachenkov);
- Workshop on Smart Antennas (WSA 2025), September 16-18, 2025, Erlangen, Germany (Samir Medina Perlaza).

Reviewer

- IEEE International Conference on Computer Communications (INFOCOM 2026), May 18-21, 2026, Tokyo, Japan (Khushboo Agarwal);
- IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN 2025) Guodong Sun;
- IEEE Wireless Communications and Networking Conference (WCNC 2025) Guodong Sun;
- IEEE International Conference on Communications (ICC 2025) Guodong Sun.
- IEEE International Symposium on Information Theory (ISIT 2026), June 28-July 3, 2026, Guangzhou, China (Khushboo Agarwal);

11.1.2 Journal

Member of the editorial boards

- ACM Transactions on Modeling and Performance Evaluation of Computing Systems (ACM ToMPECS) (Konstantin Avrachenkov, since 2016);
- AIMS (American Institute of Mathematical Sciences) Journal of Dynamics and Games (JDG) (Eitan Altman, since 2015);
- Birkhauser Journal on Dynamic Games and Applications (DGAA) (Eitan Altman, since 2012);
- CUP Probability in the Engineering and Informational Sciences (Konstantin Avrachenkov, since 2018);
- Elsevier Computer Communications (Sara Alouf, since 2021; Giovanni Neglia, since 2014);
- Elsevier Performance Evaluation (Konstantin Avrachenkov, Philippe Nain, members of Advisory Board since 2018);
- IEEE Transactions on Automatic Control (Konstantin Avrachenkov, since 2024);
- IEEE Transactions on Networking (Sara Alouf, since 2024, Eitan Altman, editor-at-large since 2013);
- Polynesian Journal of Mathematics (Samir Medina Perlaza, since 2024);
- Proceedings of the ACM on Measurement and Analysis of Computing Systems (ACM POMACS), (Konstantin Avrachenkov, member of Advisory Board since 2025);
- Springer Iran Journal of Computer Science (Eitan Altman, advisory board member);
- Taylor & Francis Stochastic Models (Konstantin Avrachenkov, since 2019).

Sara Alouf was the lead guest editor of a special issue of Performance Evaluation on extended papers from the 35th International Teletraffic Congress 2023 (see 2024's activity report). The preface has appeared in 2025 [48].

Reviewer - reviewing activities NEO members regularly perform reviews for journals such as Dynamic Games and Applications, IEEE Transactions on Networking, IEEE Transactions on Automatic Control, IEEE Transactions on Information Theory, IEEE Transactions on Wireless Communications, IEEE Transactions on Communications, IEEE Transactions on Network and Service Management, IEEE Transactions on Network Science and Engineering, Performance Evaluation, Elsevier Computer Communications, Elsevier Computer Networks.

11.1.3 Invited talks

- Konstantin Avrachenkov delivered
  - a plenary talk “Accuracy and Efficiency of Semi-Supervised Graph Clustering Methods” at the 20th Workshop on Modelling and Mining Networks (WAW 2025), June 30 - July 3, 2025, Vilnius, Lithuania;
  - an invited talk “Introduction to Graph Clustering” at Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil, December 8, 2025.
- Louis Hauseux delivered
  - a seminar "(Hyper-)graphs, percolation and clustering performances" at the MATHNET/DYOGENE Inria team at Inria-Paris center, March 2025;
  - a seminar "Méthodes de clustering avec graphes. Ou « Comment hacker (H)DBSCAN ? » Théorie, algorithmes & applications" at Valeo AI, Paris, November 2025.
- Giovanni Neglia delivered
  - an invited talk "Breaking Privacy in Federated Learning: Advances in Attribute Inference and Data Reconstruction Attacks" at University of Palermo, Italy, February 13, 2025;
  - an invited talk "Federated Learning" at the Colloquium on Networks and Learning, Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil, October 23, 2025.
- Samir Medina Perlaza delivered
  - a seminar "Variations of the Expectation due to Changes in the Measure: Applications to Generalization and Game Theory" in the series Cambridge Information Theory Seminar, Centre for Mathematical Sciences, University of Cambridge, February 26, 2025;
  - an invited talk "Generalization Error of Machine Learning Algorithms" at the Information Theory and Tapas Workshop, Universidad Carlos III de Madrid, April 2-5, 2025.

11.1.4 Leadership within the scientific community

- Eitan Altman
  - is Fellow Member of IEEE;
  - is Member of WG 7.3 of IFIP on Computer System Modeling;
  - is the elected Vice Chairman of WG 6.3 of IFIP on Performance of Communications Systems.
- Konstantin Avrachenkov is a member of Conseil Scientifique & Pédagogique EUR DS4H Univ. Côte d’Azur.
- Samir Medina Perlaza
  - is a member of the Digital Presence Committee of the IEEE Information Theory Society;
  - is a workpackage leader of the PEPR - Réseaux du Futur – A project funded by the French National Agency for Research (ANR) via the project n°ANR-22-PEFT-0010 of the France 2030 program;
  - is the organizer of the PC9 Seminar on Wireless Communications, a national online seminar part of the PEPR – Réseaux du Futur.
- Giovanni Neglia
  - holds a Chair by the Interdisciplinary Institute for Artificial Intelligence 3IA Côte d’Azur, in the theme "Core Elements of AI."

11.1.5 Research administration

- Sara Alouf
  - is a member of the Colloquium Jacques Morgenstern Committee of Inria center at Université Côte d'Azur, since March 2023;
  - is a member of NICE, the Invited Researchers Committee of Inria center at Université Côte d'Azur, since June 2020;
  - is vice-head of project-team Neo since January 2017;
  - is a member of the Selection Committee of the dAIEGDE European project.
- Konstantin Avrachenkov is an alternate representative (suppléant, collège A) on the Center Committee of Inria center at Université Côte d'Azur.
- Louis Hauseux is
  - a representative (titulaire, collège C) on the Center Committee of Inria center at Université Côte d'Azur;
  - a representantive on the catering commission of Inria center at Université Côte d'Azur.
- Alain Jean-Marie
  - has been leader of project-team Neo from January 2017 until March 2025.
- Samir Medina Perlaza
  - is an alternate representative (suppléant, collège A) on the Center Committee of Inria center at Université Côte d'Azur;
  - is a Member of the Bureau of the Réseaux, Information et Société Numérique Excellence Academy of Université Côte d'Azur;
  - represents Inria at the Conseil du Département Disciplinaire Informatique of Université Côte d'Azur.
- Giovanni Neglia
  - was a member of the competitive exam jury for a researcher position (CRCN/ISFP) at Inria Lille in 2025 (but did not participate to the interviews because of a conflict of interest);
  - is an elected member of Inria evaluation committee, since September 2023. In this role, he also participated in a working group that proposed a set of reflections, guidelines, and recommendations on the responsible use of generative AI in research professions. To ensure wider dissemination, the corresponding document is available both in English [49] and in French [50].
  - is a member of the steering commitee of Université Côte d'Azur Graduate School of Digital Systems for Humans (DS4H) since September 2022.
  - has been leader of project-team Neo since April 2025.

11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

Master

- Sara Alouf, Alain Jean-Marie, "Performance Evaluation of Networks," 24H, M2 Ubinet, Université Côte d'Azur, France.
- Louis Hauseux, Konstantin Avrachenkov, "Statistical Analysis of Networks," 21H, M2 Data Science and Artificial Intelligence, Université Côte d'Azur, France.
- Louis Hauseux, "Statistical Inference," 30H, M1 Data Science and Artificial Intelligence, Université Côte d'Azur, France.

- Giovanni Neglia, "Machine Learning: Theory and Algorithms," 24H, M2 Ubinet, Université Côte d'Azur, France.
- Giovanni Neglia together with Francesco Diana and Chuan Xu (Coati Inria team) "Federated Learning & Data Privacy," 24H, M2 Data Science and Artificial Intelligence, Université Côte d'Azur, France.
- Adrien Sardi "Optimisation Différentiable: Théorie et Algorithmes," 11H ETD, M1, ENSTA - IPP Paris, France.
- Xufeng Zhang, "Mathématiques pour l'IA," 24H, Bac+5 Formations expertes, CentraleDigitalLab @LaPlateforme_, Centrale Méditerranée, France.
- Xufeng Zhang, "Machine Learning," 22H, Bac+5 Formations expertes, CentraleDigitalLab @LaPlateforme_, Centrale Méditerranée, France.

11.2.1 Supervision

PhD defended

- José Francisco Daunas Torres, "Empirical Risk Minimization with f-divergence Regularizations in Statistical Learning," Univ. of Sheffield, defended October 29, 2025, advisors: Iñaki Esnaola and Samir Medina Perlaza.

PhD in progress

- Yaiza Bermudez, "Data Integrity in Distributed Learning Systems," Université Côte d'Azur, since November 1, 2024, advisor: Samir Medina Perlaza.
- Gaspard Gerard Philippe Berthelier, "Federated Learning for time series in the energy domain," Université Côte d'Azur, Cifre thesis with EDF, since February 1, 2025, advisors: Giovanni Neglia and Samir Medina Perlaza.
- Ibtihal El Mimouni, "Automated and responsible recommendation systems for digital marketing," Université Côte d'Azur, Cifre thesis with NSP SmartProfile, since October 1, 2022, advisor: Konstantin Avrachenkov.
- Louis Hauseux, "Classifiers on Random Graphs with applications to Social Networks and Image Processing," Université Côte d'Azur, since October 1, 2023, advisors: Konstantin Avrachenkov and Josiane Zerubia (Ayana Inria team).
- Ahmad Nasser, "Distributed training of heterogeneous architectures," Université Côte d'Azur, Cifre thesis with Nokia, since April 1, 2024, advisor: Giovanni Neglia.
- Isidoor Pinillo Esquivel, "Online Learning with Limited Resources," Université Côte d'Azur, since September 1, 2025, advisors: Sara Alouf and Giovanni Neglia.
- Julian Alfonso Santos Bustos, "Modélisation analytique de réseaux sans fils grande échelle intégrant les RIS - Optimisation de l'allocation dynamique des ressources," Université Côte d'Azur, Cifre thesis with Orange, since September 1, 2023, advisors: Eitan Altman and Konstantin Avrachenkov.
- Adrien Sardi, "Generative artificial intelligence models and resource energy management in 6G distributed networks," Université Côte d'Azur, Cifre thesis with Nokia, since January 1, 2025, advisors: Sara Alouf, Frederic Giroire (Coati Inria team), Joanna Moulierac (Coati Inria team).
- Kyrylo Tymchenko, "Enhancing Large-Scale Distributed Caching Systems with Erasure Coding," Université Côte d'Azur, since October 1, 2025, advisors: Sara Alouf and Frederic Giroire (Coati Inria team).
- Jingye Wang, "Robust Federated Learning," Université Côte d'Azur, since September 1, 2025, advisors: Sara Alouf, Giovanni Neglia, Chuan Xu (Coati Inria team).

- Xufeng Zhang, "Incentives for Federated Learning," Université Côte d'Azur, since December 1, 2023, advisors: Giovanni Neglia and Sara Alouf.
- Xinying Zou, "Generalization Capabilities of Machine Learning Algorithms," Université Côte d'Azur, since December 1, 2022, advisors: Eitan Altman and Samir Medina Perlaza.

11.2.2 Juries

PhD

- Mustapha Bounoua, "Harnessing Multimodality : Diffusion based Generative Modeling and Information Estimation," Sorbonne université, July 11, 2025 (Giovanni Neglia, jury member).
- Romain Chor, "Distributed and Federated Learning Systems: Information-Theoretic Generalization Bounds and Algorithms," Université Gustave Eiffel, September 11, 2025 (Giovanni Neglia, reviewer).
- Reiza Deylam Salehi, "Fundamental Limits of Distributed Non-Linear Function Computation in Several Multi-User Network Models," Sorbonne Université, December 5, 2025 (Giovanni Neglia, jury member).
- Yicheng Gao, "Stochastic Performance Modeling of Distributed Data Processing Systems," Imperial College London, February 19, 2025 (Sara Alouf, reviewer).
- Soumyajit Guin, "Algorithms for Various Cost Criteria in Reinforcement Learning," Indian Institute of Science Bangalore, April 16, 2025 (Konstantin Avrachenkov, reviewer).
- Jun Ju, "Reinforcement Learning for Partially Observable Environments," University of Queensland, April 28, 2025 (Konstantin Avrachenkov, reviewer).
- Lukas Stippel, "Privacy and confidentiality preserving data sharing methods for the optimization of multi-actor energy systems," Université Paris sciences et lettres, December 12, 2025 (Giovanni Neglia, jury member).
- Lucas Weber, "Exploiting Partial System Knowledge in Reinforcement Learning for Admission Control and Electricity Storage Optimization," Université Paris sciences et lettres, January 17, 2025 (Giovanni Neglia, reviewer, Alain Jean-Marie, jury president).
- Lotte Weedage, "Resilience of Cellular Networks," Twente University, January 17, 2025 (Konstantin Avrachenkov, reviewer).

12 Scientific production

12.1 Major publications

[1] K. Avrachenkov, A. Bobu and M. Dreveton. 'Higher-Order Spectral Clustering for Geometric Graphs'. In: *Journal of Fourier Analysis and Applications* 27 (15th Mar. 2021). DOI: 10.1007/s00041-021-09825-2. URL: https://inria.hal.science/hal-03169834.

[2] K. Avrachenkov and M. Dreveton. *Statistical Analysis of Networks*. Now Publishers, 6th Oct. 2022. DOI: 10.1561/9781638280514. URL: https://inria.hal.science/hal-03932416.

[3] K. E. Avrachenkov and V. Borkar. 'Whittle index based Q-learning for restless bandits with average reward'. In: *Automatica* 139 (May 2022), p. 110186. DOI: 10.1016/j.automatica.2022.110186. URL: https://inria.hal.science/hal-03582664.

[4] Y. Ben Mazziane, S. Alouf, G. Neglia and D. S. Menasche. 'TTL model for an LRU-based similarity caching policy'. In: *Computer Networks* 241 (Mar. 2024), p. 110206. DOI: 10.1016/j.comnet.2024.110206. URL: https://inria.hal.science/hal-04746044.

[5] V. Bucarey López, E. Della Vecchia, A. Jean-Marie and F. Ordoñez. 'Stationary Strong Stackelberg Equilibrium in Discounted Stochastic Games'. In: *IEEE Transactions on Automatic Control* 68.9 (2023), pp. 5271–5286. DOI: 10.1109/TAC.2022.3220512. URL: https://inria.hal.science/hal-03934114.

[6] M. Datar, E. Altman and H. Le Cadre. 'Strategic Resource Pricing and Allocation in a 5G Network Slicing Stackelberg Game'. In: *IEEE Transactions on Network and Service Management* 20.1 (2023), pp. 502–520. DOI: 10.1109/TNSM.2022.3216588. URL: https://inria.hal.science/hal-03824540.

[7] V. Kavitha and E. Altman. 'Controlling Packet Drops to Improve Freshness of information'. In: Netgcoop 2020 - International Conference on NETwork Games, Control and Optimisation. Cargese, France, 22nd Sept. 2021. DOI: 10.1007/978-3-030-87473-5_7. URL: https://inria.hal.science/hal-02931314.

[8] O. Marfoq, G. Neglia, A. Bellet, L. Kameni and R. Vidal. 'Federated Multi-Task Learning under a Mixture of Distributions'. In: NeurIPS 2021 - 35th Conference on Neural Information Processing Systems. Sydney / Virtual, Australia, 6th Dec. 2021. URL: https://hal.science/hal-03406994.

[9] S. M. Perlaza, G. Bisson, I. Esnaola, A. Jean-Marie and S. Rini. 'Empirical Risk Minimization with Relative Entropy Regularization'. In: *IEEE Transactions on Information Theory* 70.7 (2024), pp. 5122–5161. DOI: 10.1109/TIT.2024.3365728. URL: https://hal.science/hal-03849748.

[10] T. Si Salem, G. Neglia and S. Ioannidis. 'No-regret Caching via Online Mirror Descent'. In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 8.4 (11th Aug. 2023), pp. 1–32. DOI: 10.1145/3605209. URL: https://hal.science/hal-04181387.

[11] A. Tajer, S. M. Perlaza and H. Vincent Poor, eds. *Advanced Data Analytics for Power Systems*. Cambridge University Press, 1st Jan. 2021. DOI: 10.1017/9781108859806. URL: https://hal.science/hal-03128425.

[12] G. Vardoyan, P. Nain, S. Guha and D. Towsley. 'On the Capacity Region of Bipartite and Tripartite Entanglement Switching'. In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 8.1-2 (June 2023), pp. 1–18. DOI: 10.1145/3571809. URL: https://inria.hal.science/hal-04018593.

12.2 Publications of the year

International journals

[13] K. Agarwal, K. Avrachenkov, R. Vyas and V. Kavitha. 'Two Choice Behavioral Game Dynamics with Myopic-Rational and Herding Players'. In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 9.1 (2025), pp. 1–26. DOI: 10.1145/3711706. URL: https://inria.hal.science/hal-04883482 (cit. on pp. 20, 26, 36).

[14] K. Avrachenkov and F. Spieksma. 'Deviation matrix in denumerable Markov processes: Characterisation, rank-one perturbation and restart'. In: *The Annals of Applied Probability* 35.6 (1st Dec. 2025), pp. 4354–4380. DOI: 10.1214/25-AAP2223. URL: https://inria.hal.science/hal-05382306 (cit. on p. 11).

[15] Y. Ben Mazziane, F. Faticanti, S. Alouf and G. Neglia. 'Efficient and Optimal No-Regret Caching Under Partial Observation'. In: *IEEE Transactions on Networking* (3rd Dec. 2025), pp. 1–11. DOI: 10.1109/TON.2025.3636929. URL: https://inria.hal.science/hal-05400041 (cit. on p. 21).

[16] D. Carra, G. Neglia and X. Zhang. 'Low-Complexity online learning for caching'. In: *Computer Networks* 273 (Dec. 2025), p. 111743. DOI: 10.1016/j.comnet.2025.111743. URL: https://hal.science/hal-05422240 (cit. on pp. 22, 34, 35).

[17] J. Clarkson, K. Avrachenkov and E. Altman. 'Queues with inspection cost: To see or not to see?' In: *Queueing Systems* 109.13 (Mar. 2025). DOI: 10.1007/s11134-025-09940-7. URL: https://inria.hal.science/hal-04994271 (cit. on p. 20).

[18] F. Daunas, I. Esnaola, S. M. Perlaza and H. V. Poor. 'Asymmetry of the Relative Entropy in the Regularization of Empirical Risk Minimization'. In: *IEEE Transactions on Information Theory* 71.8 (2025), pp. 6198–6226. DOI: 10.1109/TIT.2025.3565706. URL: https://inria.hal.science/hal-04719783 (cit. on pp. 14, 36).

[19] L. L. Felipe, K. Avrachenkov and D. S. Menasché. 'From Leiden to Pleasure Island: The Constant Potts Model for Community Detection as a Hedonic Game'. In: *Physica A: Statistical Mechanics and its Applications* 680 (Dec. 2025), p. 130989. DOI: 10.1016/j.physa.2025.130989. URL: https://inria.hal.science/hal-05458053 (cit. on p. 16).

[20] L. Hauseux, K. Avrachenkov and J. Zerubia. 'Generalization of Single-Linkage with Higher-Order Interactions'. In: *Applied Network Science* (5th Nov. 2025). DOI: 10.1007/s41109-025-00756-1. URL: https://inria.hal.science/hal-05369659. In press (cit. on p. 16).

[21] P. Howlett, B. K. Beare, M. Franchi, J. Boland and K. Avrachenkov. 'The Granger–Johansen representation theorem for integrated time series on Banach space'. In: *Journal of Time Series Analysis* 46 (2025), pp. 432–457. DOI: 10.1111/jtsa.12766. URL: https://inria.hal.science/hal-04755738 (cit. on p. 13).

[22] H. Le Cadre, M. Datar, M. Guckert and E. Altman. 'Learning Market Equilibria Preserving Statistical Privacy Using Performative Prediction'. In: *IEEE Transactions on Automatic Control* 70.11 (1st Nov. 2025), pp. 7125–7140. DOI: 10.23919/ECC57647.2023.10178247. URL: https://inria.hal.science/hal-04343535 (cit. on p. 21).

[23] S. M. Perlaza and G. Bisson. 'Variations on the Expectation Due to Changes in the Probability Measure'. In: *Entropy* 27.8 (14th Aug. 2025), p. 865. DOI: 10.3390/e27080865. URL: https://hal.science/hal-05211105 (cit. on pp. 14, 36).

[24] A. R. Ramtin, P. Nain and D. Towsley. 'Quickest Change Detection in Continuous-Time in Presence of a Covert Adversary'. In: *IEEE Signal Processing Letters* 32 (Nov. 2025), pp. 4299–4303. DOI: 10.1109/LSP.2025.3628793. URL: https://inria.hal.science/hal-05357817 (cit. on p. 22).

[25] S. Ul Zuhra, S. M. Perlaza, H. V. Poor and M. Skoglund. 'Simultaneous Information and Energy Transmission with Short Packets and Finite Constellations'. In: *IEEE Transactions on Communications* 73.9 (24th Mar. 2025), pp. 7505–7516. DOI: 10.1109/TCOMM.2025.3554115. URL: https://inria.hal.science/hal-03848708 (cit. on pp. 22, 36).

International peer-reviewed conferences

[26] K. Agarwal, K. Avrachenkov, V. Kavitha and R. Vyas. 'Balancing rationality and social influence: Alpha-rational Nash equilibrium in games with herding'. In: GameNets 2025 - 13th EAI International Conference on Game Theory for Networks. Cambridge, United Kingdom, 17th Mar. 2025. DOI: 10.1007/978-3-032-12915-4_6. URL: https://inria.hal.science/hal-04955252 (cit. on pp. 19, 26, 36).

[27] K. Avrachenkov, K. Das, V. Kavitha and V. Singh. 'Stability of Polling Systems for a Large Class of Markovian Switching Policies'. In: CDC 2025 - IEEE Conference on Decision and Control. Rio de Janeiro, Brazil, 9th Dec. 2025. URL: https://inria.hal.science/hal-05458169 (cit. on pp. 12, 26).

[28] K. Avrachenkov, M. Dhiman and V. Kavitha. 'Constrained Average-Reward Intermittently Observable MDPs'. In: CDC 2025 - IEEE Conference on Decision and Control. Rio de Janeiro, Brazil, 10th Dec. 2025. URL: https://inria.hal.science/hal-05458150 (cit. on pp. 12, 26).

[29] Y. Ben Mazziane and O. Marfoq. 'Universal and Tight Bounds on Counting Errors of Count-Min Sketch with Conservative Updates'. In: International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS). Vol. 9. 2. Stony brook, NY, United States, 3rd June 2025, pp. 1–32. DOI: 10.1145/3727135. URL: https://hal.science/hal-05377989 (cit. on p. 13).

[30] Y. Ben Mazziane, C.-M. Mboulou-Moutoubi, F. De Pellegrini and E. Altman. 'Learning to Bid in Proportional Allocation Auctions with Budget Constraints'. In: 23rd International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt). Linkoping, Sweden, 26th May 2025, pp. 1–8. DOI: 10.23919/WiOpt66569.2025.11123293. URL: https://hal.science/hal-05377998 (cit. on p. 20).

[31] F. Daunas, I. Esnaola and S. M. Perlaza. 'A Dual Optimization View to Empirical Risk Minimization with f -Divergence Regularization'. In: IEEE-ITW 2025 - Information Theory Workshop. Sidney, Australia, 29th Sept. 2025. DOI: 10.1109/ITW62417.2025.11240444. URL: https://inria.hal.science/hal-04969347 (cit. on pp. 14, 36).

[32] F. Diana, O. Marfoq, C. Xu, G. Neglia, F. Giroire and E. Thomas. 'Attribute Inference Attacks for Federated Regression Tasks'. In: 39th Annual AAAI Conference on Artificial Intelligence - AAAI 2025. Vol. 39. 15. Philadelphia (Pennsylvania), United States, 25th Feb. 2025, pp. 16271–16279. DOI: 10.1609/aaai.v39i15.33787. URL: https://inria.hal.science/hal-04878082 (cit. on pp. 19, 34).

[33] F. Diana, A. Nusser, C. Xu and G. Neglia. 'Cutting Through Privacy: A Hyperplane-Based Data Reconstruction Attack in Federated Learning'. In: 41st Conference on Uncertainty in Artificial Intelligence (UAI '25). Rio de Janeiro, Brazil, 21st July 2025, pp. 959–980. DOI: 10.5555/3762387.3762431. URL: https://hal.science/hal-05318034 (cit. on pp. 19, 34, 35, 37).

[34] I. El Mimouni and K. Avrachenkov. 'Deep Q-Learning with Whittle Index for Contextual Restless Bandits: Application to Email Recommender Systems'. In: NLDL 2025 - 6th Northern Lights Deep Learning Conference. PMLR-265. Tromso, Norway, 6th Jan. 2025, pp. 176–183. URL: https://inria.hal.science/hal-05311893 (cit. on pp. 15, 25).

[35] I. El Mimouni and K. Avrachenkov. 'Fostering Responsibility in Email Marketing: A Contextual Restless Bandit Framework'. In: ECML PKDD 2025 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases. LNAI-16020. Porto, Portugal, 15th Sept. 2025, pp. 366–383. DOI: 10.1007/978-3-662-72243-5_21. URL: https://inria.hal.science/hal-05311872 (cit. on pp. 15, 25).

[36] A. Fox, F. De Pellegrini, E. Altman, A. Ghosh and N. Shroff. 'Performing Load Balancing under Constraints'. In: WiOpt 2025 - 23rd International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks. Linkoping, Sweden, 26th May 2025, pp. 103–110. DOI: 10.23919/WiOpt66569.2025.11123242. URL: https://hal.science/hal-05296436 (cit. on pp. 12, 36).

[37] F. Galante, G. Neglia and E. Leonardi. 'Scalable Decentralized Algorithms for Online Personalized Mean Estimation'. In: 39th Annual AAAI Conference on Artificial Intelligence - AAAI 2025. Philadelphia (Pennsylvania), United States, 25th Feb. 2025. DOI: 10.1609/aaai.v39i16.33835. URL: https://hal.science/hal-04878986 (cit. on pp. 18, 34, 37).

[38] L. Hauseux, N. Soprano-Loto and K. Avrachenkov. 'Higher-order Monte Carlo cluster dynamics for community detection in Euclidean graphs'. In: Asilomar Conference on Signals, Systems, and Computers 2025. Pacific Grove (CA), United States, 26th Oct. 2025. URL: https://inria.hal.science/hal-05267074 (cit. on p. 16).

[39] T.-K. Huynh, M. Egan, G. Neglia and J.-M. Gorce. 'Streaming Federated Learning with Markovian Data'. In: NeurIPS 2025 - Thirty-Ninth Annual Conference on Neural Information Processing Systems. San Diego, United States, 2nd Dec. 2025. URL: https://inria.hal.science/hal-05422204 (cit. on pp. 17, 24, 34, 35, 37).

[40] L. Lopes Felipe, K. Avrachenkov and D. Sadoc Menasché. 'Robustness against Frustration in Community Detection'. In: SIGMETRICS 2025 - ACM Sigmetrics 2025 (Student Competition). Vol. 53. 2. Stony Brook, United States, 27th Aug. 2025, pp. 128–130. DOI: 10.1145/3764944.3764977. URL: https://inria.hal.science/hal-05458107 (cit. on p. 16).

[41] C. Marlain Mboulou-Moutoubi, Y. Ben Mazziane, F. De Pellegrini and E. Altman. 'Best-Response Learning in Budgeted α-Fair Kelly Mechanisms'. In: NETGCOOP 2025 - 12th International Conference of Networks, Games, Control and Optimization. Bilbao, Spain, 8th Oct. 2025. DOI: 10.1007/978-3-032-09315-8_9. URL: https://hal.science/hal-05329798 (cit. on p. 20).

[42] D. Richards Arputharaj, C. Rodriguez, A. Rodio and G. Neglia. 'Green Federated Learning via Carbon-Aware Client and Time Slot Scheduling'. In: MASCOTS 2025 - 33rd International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication System. Paris, France, 21st Oct. 2025. DOI: 10.1109/MASCOTS67699.2025.11283314. URL: https://inria.hal.science/hal-05423023 (cit. on pp. 10, 18, 34, 35, 37).

[43] A. Rodio, G. Neglia, Z. Chen and E. G. Larsson. 'A Unified Convergence Analysis for Semi-Decentralized Learning: Sampled-to-Sampled vs. Sampled-to-All Communication'. In: AAAI-26 - 40th Annual AAAI Conference on Artificial Intelligence. Singapore, Singapore, 20th Jan. 2026. URL: https://inria.hal.science/hal-05423080.

[44] J. Talpini, M. Savi and G. Neglia. 'FedBEns: One-Shot Federated Learning based on Bayesian Ensemble'. In: ICML 2025 - 42nd International Conference on Machine Learning. Vol. PMLR 267. Vancouver, Canada, 13th July 2025. URL: https://inria.hal.science/hal-05429191 (cit. on pp. 18, 34, 35, 37).

[45] M. van Waerebeke, M. Lorenzi, G. Neglia and K. Scaman. 'When to Forget? Complexity Trade-offs in Machine Unlearning'. In: International Conference in Machine Learning (ICML 2025). Vancouver (BC), Canada, 13th July 2025. URL: https://hal.science/hal-05464347 (cit. on p. 15).

[46] X. Zhang, S. Alouf and G. Neglia. 'Memory-efficient Online Caching Policies with Regret Guarantees'. In: IFIP NETWORKING 2025. Limassol, Cyprus, May 2025. URL: https://inria.hal.science/hal-05326754 (cit. on pp. 22, 34, 35).

Conferences without proceedings

[47] Y. Bermudez, G. Bisson, I. Esnaola and S. Perlaza. 'Démonstrations des théorèmes folkloriques sur la dérivée de Radon-Nikodym'. In: GRETSI'25 - XXXe Colloque Francophone de Traitement du Signal et des Images. Strasbourg, France, 25th Aug. 2025. URL: https://hal.science/hal-05157449 (cit. on pp. 14, 36, 37).

Edition (books, proceedings, special issue of a journal)

[48] S. Alouf, O. Hohlfeld and Z. Jiang, eds. *Preface: Special issue on ITC 2023* 167 (Mar. 2025). DOI: 10.1016/j.peva.2024.102462. URL: https://inria.hal.science/hal-04731314 (cit. on p. 40).

Reports & preprints

[49] S. Arias, M. Bergmann, F. Campillo, M.-A. Enard, C. Fabre, F. Garcia, B. Guedj, E. Jeannot, G. Neglia, D. Peurichard, D. Racoceanu, B. Sagot and G. Tworkowski. *Reflections on the Use of Generative AI for Research Professions*. Inria, 9th July 2025. URL: https://inria.hal.science/hal-05188001 (cit. on p. 42).

[50] S. Arias, M. Bergmann, F. Campillo, M.-A. Enard, C. Fabre, F. Garcia, B. Guedj, E. Jeannot, G. Neglia, D. Peurichard, D. Racoceanu, B. Sagot and G. Tworkowski. *Réflexions sur l'usage de l'IA générative pour les métiers de la recherche*. Inria, 2025, pp. 1–10. URL: https://inria.hal.science/hal-05187992 (cit. on p. 42).

[51] K. Avrachenkov, L. Gamertsfelder and B. Mourrain. *Weighted Moment-SoS approach to POMDPs with polynomial data*. 12th Nov. 2025. URL: https://inria.hal.science/hal-05355845.

[52] Y. Bermudez, G. Bisson, I. Esnaola and S. M. Perlaza. *Proofs for Folklore Theorems on the Radon-Nikodym Derivative*. RR-9591. Inria, 10th July 2025. URL: https://inria.hal.science/hal-04969319 (cit. on pp. 14, 36, 37).

[53] Y. Bermudez, S. M. Perlaza, I. Esnaola and H. V. Poor. *Statistical Federated Learning and Generalization*. RR-9599. Inria, 31st Oct. 2025. URL: https://inria.hal.science/hal-05355756 (cit. on pp. 18, 36, 37).

[54] F. Daunas, I. Esnaola and S. M. Perlaza. *A Dual Optimization View to Empirical Risk Minimisation with f-Divergence Regularization*. RR-9594. Inria, 25th July 2025. URL: https://inria.hal.science/hal-05200897 (cit. on pp. 14, 36).

[55] M. Nicaise, Y. Bermudez and S. M. Perlaza. *On optimal transport with f-divergence regularization*. RR-9607. Inria, Jan. 2026. URL: https://inria.hal.science/hal-05454647.

[56] S. M. Perlaza and G. Bisson. *Variations on the Expectation due to Changes in the Probability Measure*. RR-9592. Inria, 10th July 2025, p. 28. URL: https://inria.hal.science/hal-04969296 (cit. on pp. 14, 36).

[57] A. R. Ramtin, P. Nain and D. Towsley. *Quick Change Detection in Continuous-Time in Presence of a Covert Adversary*. 19th Sept. 2025. URL: https://inria.hal.science/hal-05268481 (cit. on p. 22).

[58] X. Zou, S. M. Perlaza and I. Esnaola. *A Minimax Approach for Robustifying Statistical Learning Algorithms*. 9600. Inria, Dec. 2025. URL: https://inria.hal.science/hal-05393845 (cit. on pp. 14, 36).

Other scientific publications

[59] H. Le Cadre, M. Datar, M. Guckert and E. Altman. *Supplementary Material to "Learning Market Equilibria Preserving Statistical Privacy Using Performative Prediction"*. 30th Apr. 2025. URL: https://inria.hal.science/hal-05052889 (cit. on p. 21).

12.3 Cited publications

[60] M. Tidball, A. Jean-Marie and T. Jiménez. 'Theory of Conjectural Learning'. In: *Advances in Dynamic Games: Theory, Experiments, and Applications*. Ed. by G. Martín-Herrán, A. Traulsen and F. Wagener. Annals of the International Society of Dynamic Games. In press. Springer Nature Switzerland AG, 2026 (cit. on p. 16).