RESEARCH CENTRE

Inria Saclay Centre at Université Paris-Saclay

IN PARTNERSHIP WITH:

CNRS, Université Paris-Saclay

2024

ACTIVITY REPORT

Project-Team

TAU

TAckling the Underspecified

IN COLLABORATION WITH: Laboratoire Interdisciplinaire des Sciences du Numérique

DOMAIN

Applied Mathematics, Computation and Simulation

THEME

Optimization, machine learning and statistical methods

Contents

Project-Team TAU

Creation of the Project-Team: 2019 July 01

Keywords

Computer sciences and digital sciences

A3.3.3. – Big data analysis

A3.4. – Machine learning and statistics

A3.5.2. – Recommendation systems

A6.2. – Scientific computing, Numerical Analysis & Optimization

A8.2. – Optimization

A8.6. – Information theory

A8.12. – Optimal transport

A9.2. – Machine learning

A9.3. – Signal analysis

Other research topics and application domains

B1.1.4. – Genetics and genomics

B4. – Energy

B9.1.2. – Serious games

B9.5.3. – Physics

B9.5.5. – Mechanics

B9.5.6. – Data science

B9.6.10. – Digital humanities

1 Team members, visitors, external collaborators

Research Scientists

- Marc Schoenauer [Team leader, INRIA, Senior Researcher, until May 2024, HDR]
- Cyril Furtlehner [Team leader, INRIA, Researcher, HDR]
- Guillaume Charpiat [INRIA, Researcher]
- Alessandro Ferreira Leite [Safran, from Feb 2024 until Nov 2024]
- Alessandro Ferreira Leite [INRIA, Advanced Research Position, until Jan 2024]
- Flora Jay [CNRS, Researcher]
- Marc Schoenauer [INRIA, Emeritus, from May 2024, HDR]
- Martine Michele Sebag [CNRS, Senior Researcher, HDR]
- Beatriz Seoane Bartolomé [UNIV PARIS SACLAY, until Sep 2024]

Faculty Members

- Philippe Caillou [UNIV PARIS SACLAY, Associate Professor]
- Sylvain Chevallier [UNIV PARIS SACLAY, Professor, HDR]
- Cécile Germain [UNIV PARIS SACLAY, Emeritus]
- Isabelle Guyon [UNIV PARIS SACLAY, Professor]
- Matthieu Kowalski [UNIV PARIS SACLAY, Associate Professor, HDR]
- François Landes [UNIV PARIS SACLAY, Associate Professor]

Post-Doctoral Fellows

- Shuyu Dong [INRIA, until Sep 2024]
- Matthieu Nastorg [INRIA, Post-Doctoral Fellow, from Aug 2024]
- Leo Benoit Planche [UNIV PARIS SACLAY, Post-Doctoral Fellow, until Sep 2024]
- Stephane Rivaud [INRIA, from Sep 2024]
- Sara Sedlar [Universite Paris-Saclay, Post-Doctoral Fellow, until Jul 2024]

PhD Students

- Anaclara Alvez Canepa [UNIV PARIS SACLAY]
- Bruno Aristimunha Pinto [INRIA]
- Nicolas Atienza [Thalès, CIFRE]
- Nicolas Bereux [UNIV PARIS SACLAY]
- Guillaume Bied [UNIV PARIS SACLAY]
- Eva Boguslawski [RTE, CIFRE]
- Styliani Douka [INRIA]

- Romain Egele [INRIA, until Jun 2024]
- Emmanuel Goutierre [Université Paris-Saclay]
- Alice Lacan [UNIV PARIS SACLAY]
- Armand Lacombe [LISN, until Aug 2024]
- Jean-Baptiste Malagnoux [CentraleSupelec]
- Apolline Mellot [Universite Paris-Saclay, until Nov 2024]
- Thibault Monsel [CNRS]
- Matthieu Nastorg [INRIA, until Mar 2024]
- Solal Nathan [LISN]
- Francesco Saverio Pezzicoli [UNIV PARIS SACLAY]
- Audrey Poinsot [EKIMETRICS]
- Arnaud Quelin [SORBONNE UNIVERSITE]
- Cyriaque Rousselot [UNIV PARIS SACLAY]
- Theo Rudkiewicz [ENS PARIS-SACLAY, from Sep 2024]
- Nilo Elias Schwencke [Univ. Paris-Saclay]
- Haozhe Sun [UNIV PARIS SACLAY, until Jan 2024]
- Antoine Szatkownik [UNIV PARIS SACLAY]
- Sébastien Velut [ISAE SupAero]
- Manon Verbockhaven [UNIV PARIS SACLAY]
- Mathurin Videau [Meta, CIFRE]
- Assia Wirth [UNIV PARIS SACLAY]
- Maria Sayu Yamamoto [INRIA, until Aug 2024]
- Badr Youbi [Meta, CIFRE]

Technical Staff

- Hande Gozukan [INRIA, Engineer]
- Adrien Pavao [Inria until Oct. 24, Univ. Paris-Saclay since Nov. 2024]
- Dylan Sechet [CENTRALESUPELEC, Engineer, from Aug 2024]
- Sebastien Treguer [INRIA, Engineer, until Aug 2024]

Interns and Apprentices

- Michel Doroch [INRIA, Intern, from May 2024 until Aug 2024]
- Gianlucca Fiori Oliveira [INRIA, Intern, from Jun 2024 until Aug 2024]
- Chloe Godet [INRIA, Intern, from Jun 2024 until Aug 2024]
- Barbara Hajdarevic [UNIV PARIS SACLAY, until May 2024]
- Quang Phuoc Ho [INRIA, Intern, from May 2024 until Aug 2024]
- Gloire Linvani [INRIA, Intern, from May 2024 until Jul 2024]
- Theo Marchetta [INRIA, Intern, from Apr 2024 until Jul 2024]
- Matias Nicolas Ortiz Angel [INRIA, Intern, until Apr 2024]
- Andrei-Tiberiu Pantea [INRIA, Intern, from Apr 2024 until Aug 2024]
- Karim Rochd [INRIA, Intern, from May 2024 until Aug 2024]
- Abel Rohic Collard [INRIA, Intern, from May 2024 until Jun 2024]

Administrative Assistant

- Julienne Moukalou [INRIA]

External Collaborators

- Michele Bucci [Safran tech]
- Sergio Chibbaro [UNIV PARIS SACLAY, from Nov 2024]
- Aurelien Decelle [UNIV COMPUT MADRID, until Jun 2024]
- Anne-Catherine Letournel [Université Paris-Saclay]
- Burak Yelmen [UNIV PARIS SACLAY]

2 Overall objectives

2.1 Presentation

Building upon the expertise in machine learning (ML) and stochastic optimization, and statistical physics of the former TAO project-team, the TAU team aims to tackle **the vagueness of the Big Data purposes**. Based on the claim that (sufficiently) big data can to some extent compensate for the lack of knowledge, Big Data is hoped to fulfill all Artificial Intelligence commitments.

This makes Big Data under-specified in three respects:

- A first source of under-specification is related to **common sense**, and the gap between observation and interpretation. The acquired data do not report on “obvious” issues; still, obvious issues are not necessarily so for the computer. Providing the machine with common sense is a many-faceted, AI hard, challenge. A current challenge is to **interpret the data** and cope with its blind zones (e.g., missing values, contradictory examples, ...).
- A second source of under-specification regards the **steering of a Big Data system**. Such systems commonly require lifelong learning in order to deal with open environments and users with diverse profiles, expertises and expectations. A Big Data system thus is a dynamic process, whose behavior will depend in a cumulative way upon its future environment. The challenge regards **the control of a lifelong learning system**.

- A third source of under-specification regards its social acceptability. There is little doubt that Big Data can pave the way for Big Brother, and ruin the social contract through modeling benefits and costs at the individual level. What are the **fair trade-offs between safety, freedom and efficiency** ? We do not know the answers. A first practical and scientific challenge is to first assess, and then enforce, **the trustworthiness of solutions**.

However, several concerns have emerged in the last years regarding Big Data models. First, in industrial context, data is now always big, and many practical problems are relevant to **small data**. On the opposite, when big data is available, the arms race around LLMs has given birth to increasingly big models, involving hundreds of billions of parameters, and environmental concerns are becoming increasingly high, for their training, but even for their use and the inference process.

Our initial overall under-specification considerations, mitigated with the concerns above, have lead the team to align its research agenda along four pillars:

- Frugal Learning, addressing the environmental concerns, in terms of deep network architecture and considering the small data regimes;
- Causal Learning, a grounded way to address the trustworthiness issue by improving explainability of the results;
- Bidirectional links with Statistical Physics, to better understand very large systems and improve their performances, both in terms of accuracy of the models and energy consumption in their use;
- Hybridization of Machine Learning with Numerical Simulations, again aiming to reach better efficiency while decreasing the computing needs.

Last but not least, the organization of challenges and the design of benchmarks, a cornerstone of Machine Learning nowadays, remains an active thread of the team activity, in particular through the Codalab platform and its new version Codabench.

3 Research program

3.1 Frugal Learning

Frugality is a must for machine learning: because of scientific concerns (monster models imply non-reproducible science); because of sustainability concerns (energy consumption to train and use models); because of applicability concerns: in most non-GAFAM/GAMAM settings, we deal with small data, and PhD students not infrequently receive the promised data in the last months of their PhDs.

We target in particular three domains: data frugality, computational complexity at test time (to minimize environmental footprint when using the trained network at large scales), and computational complexity of neural architecture search (i.e. of the automatically finding of neural architectures suitable for a given machine learning task at hand, at training time). The mainstream strategy suggests finding a model in a large (overparameterized) model space, in order to avoid optimization and expressivity issues, and then pruning it [117]. An alternative to the above strategy, named neural network growth, consists in starting from a tiny architecture and grafting additional neurons or layers to extend its representation power on demand, on the fly during training. This raises interesting mathematical questions regarding optimization, generalization, and statistical significance.

An approach we are currently developping follows the preliminary proof of concept in M. Verbockhaven's PhD where we seek to adapt the neural tangent kernel to the directions desired by the functional gradient descent. This kind of approach could be useful not only to automatically (and frugally) design from scratch a neural network architecture suitable for a new task, but could also be of prime interest in classical Neural Architecture Search to provide directly optimal architecture variations instead of searching for them in a computationally-heavy trial-and-error fashion.

A nice byproduct is that by building smaller models, one potentially requires smaller data, and is potentially less prone to overfit. This opens interesting questions regarding regularization in deep learning and advocates for a more reasonable, guided use of combinatorics, that appear through traditional random initializations of numerous neurons (lottery ticket hypothesis [104]).

3.2 Causal Learning

The rise of causal modelling (CM) has an impact on the general agenda of machine learning, more aware of, and more robust w.r.t. the potential and usual differences of data distributions between training and testing times or along lifelong learning. This new agenda focuses on sorting out distribution-independent relations (hopefully causal ones) among the observed features, and other relations, possibly reflecting spurious correlations. The expected benefits of this causality-inspired focus is to deliver learned models that are: i) more robust w.r.t. the non iid setting; ii) more interpretable; iii) possibly humanly verified. The last two properties only hold, naturally, if the features are expressed at a sufficient level of generality.

A key scientific question is whether and how the main lesson of Deep Learning (It's the representation, stupid !) can be ported to causal modelling, particularly so when dealing with raw, redundant and/or high dimensional data. The use of latent variables and structures in e.g. [106, 126, 128] has shown its potential to disentangle root causes (sources of the observed data) and cope with hidden confounders. However, causal modelling comes with the key requirement of identifiability/uniqueness of the learned causal models, that is in general *not* satisfied in mainstream machine learning.

A promising research direction toward model identifiability is to investigate the stability of causal discovery. Formally, one might want that, if data $\mathscr{D}$ yields model $\mathscr{G}$, then data $\mathscr{D}'$ generated after $\mathscr{G}$ yields a model $\mathscr{G}'$ that is in essence same as $\mathscr{G}$. This direction opens to two strategies: i) observing the differences between $\mathscr{G}$ and $\mathscr{G}'$ sheds some light about the diversity in the data with some/no impact on the causal modelling output, i.e. the biases of the causal discovery algorithms; ii) and more deeply, the issue of stability can inspire new learning criteria, enforcing the stability of the causal models under such changes of distribution. Another hot research direction investigates how to improve the interpretability of a model, without degrading too seriously its accuracy. Let us focus on the task of interpreting hidden variables and their interactions. A possible strategy – at the core of the AI2 French-German proposal, 2023-2026; coll. Fraunhofer Bonn – takes inspiration from the Multi-Criteria Decision Aid literature (and the lessons learned in R. Bresson's PhD [78, 79]). The idea is that i) if the last say two layers of a deep net were structured as a hierarchical choquet integral (HCI); ii) and if their input (the nodes in the layer before) were interpretable (giving a feature name to each node), then the black box could be made transparent, expressing sparse hierarchical interactions (HCI) of these features. The first condition can be handled by retraining a trained efficient deep net, and imposing HCI constraints on the last two layers. A pending question is how these constraints would degrade the loss accuracy (depending on the number of would-be features). The second condition will be met by associating a supervised binary learning problem to each node, and involving the expert in the loop (or possibly exploiting textual information about the samples) to solve it.

3.3 Machine Learning with/for Statistical Physics

Concerning the links between statistical physics and machine learning, we are working on both aspect of ML with Statistical Physics and Statistical Physics for ML.

1- The first line of research, based on our expertise on generative models, will be headed toward efficient methods for frugal and interpretable generative models, typically energy based models (EBM)[113]. In particular concerning explainability we will look for physically-inspired interpretable feature extraction processes, exploring the possibilities of using EBMs as data-driven fitness landscapes.

2- This explainability aspect will be actually important for our second axes concerning applications of EBMs in bioinformatics. For instance, given data of protein's families with common ancestors, we expect to be able to learn a model describing the statistics of the family, and then use this model to predict the mutation of the amino-acid. More broadly we will develop methods for direct coupling extraction with RBMs, clustering of data in families and subfamilies, semi-supervised strategies and use EBM for pattern extraction in genomics/proteomics sequence datasets.

3- Our third axes will focus on symmetries both for methods and applications. "It is only slightly overstating the case to say that physics is the study of symmetry" (Philip Anderson 1972), and enforcing symmetries into models or finding symmetries in the data[116] is also key to ML. CNNs can enforce translation equivariance, GNNs enforce permutation equivariance, and more recently, rules for building roto-translation-equivariant networks have been devised[92]. The importance of symmetries has been acknowledged in [90], coining the term "geometric deep learning" to refer to group-invariance aware neural networks. We are working on pushing roto-translation equivariance further, with application to molecular systems or amorphous materials. Furthermore, from statistical physics we know that sytems display scale-invariant distributions at their critical point. Starting from simple avalanche models as benchmarks, we want to design networks that would be genuinely scale-equivariant (or invariant). Applications range from seismic hazard to solar wheather forecast, i.e. any area where large events-related data are scarce. Such networks would de facto perform extrapolation, a rare feature in Machine Learning. This avenue of research is being studied within Anaclara Alvez' PhD (co-supervised by Cyril Furtlehner and François Landes) and the ANR Scalp (2025-2028) to extend this to mult-fractal data.

4- Our last axes deals with fundamental properties of ML like for instance neural scaling laws[84] and is based on recent theoretical progresses like the formulation of the neural tangent kernel[108] and the lazy regime[91]. Various asymptotic results can be obtained thanks to random matrix theory or replica approaches. Equipped with such tools we would like to explore for instance the learning dynamics beyond the lazy regime, the out-of-equilibrium regimes of EBMs via dynamical mean field theory but also the utility-privacy trade-off with solvable models.

3.4 Machine Learning for Numerical Simulations

Until recently, applying off-the-shelf neural nets to numerical simulations (e.g., approximating the solution of PDEs) could only compete with numerical solvers in a few situations: when the problem is simple and of reasonable size, and when a limited accuracy, that does not need to be guaranteed, is sufficient. For instance, in cases involving chaotic behaviors (e.g. turbulent flows in fluid dynamics), current models fail to fit the target trajectory in the mid to long term. The situation is rapidly evolving (see e.g., GraphCast, by DeepMind [111]), but there remains a need for tighter coupling between ML and simulations.

Building upon TAU expertise in numerical engineering, it is suggested that the diversity of use cases tackled in applications (recent and on-going PhDs of W. Liu, E. Menier, M. Nastorg, E. Goutierre; T. Monsel, and collaboration with the IRT SystemX IA2 program as well as with IFPEN) can lead to formulating general principles and methodology.

One research direction is to consider more structured losses/architectures. This research direction evolves at a rapid pace: from convolutional architectures, to distributional architectures enforcing invariance or equivariance properties [83], to optimal transport based embeddings [123]. It is believed that new losses, aimed at preserving statistical quantities (e.g. high order moments; extreme value exponents), might help to learn and reproduce chaotic data trajectories, better than MSE losses. Nevertheless, until theoretical guidelines are available to the practionner, it is important to be able to experimentally guide and validate users' choices in terms of architecture/loss, for any new use-case. There is today a lack of well-grounded and widely accepted benchmarks, and we contribute to the IRT SystemX LIPS platform (Learning Industrial Physical Simulation benchmark suite) [115], lead by our collaborators from IRT (Mouadh Yagoubi) and RTE (Benjamin Donnot and Antoine Marot).

Another direction of research concerns how the domain know-how can best be conveyed to the learning process: through priors; or warm-starting the solution; or enforcing the required solution properties through specific loss terms; or maybe simply choosing the right training samples.

A theoretically and practically important domain concerns the coupling of an ML model and a numerical simulator, with mutual benefits (compensating for insufficient data; adjusting the simulator hyper-parameters; prioritizing new experiments toward optimal design or model identification; providing a fast sampler; addressing inverse problems). Mimicking the structure of the simulator/the physical phenomenon through the neural architecture helps to guide the optimization, all the more so as it supports the definition of auxiliary losses (e.g. based on internal states of the simulator). Again, the use of auxiliary losses can be very useful, *if* an appropriate learning schedule has been defined (controlling the impact/weight of each auxiliary loss depending on the current state of the model and of the learning trajectory).

Last but not least, unleashing the power of the recently emerged Foundation Models and Transformers resulted in low hanging fruits (e.g., more powerful surrogate models) that have not yet been picked up, and will also open new avenues for hybrid/multidisciplinary research.

3.5 Challenge Organization

In the rapidly evolving field of machine learning (Data-Driven Artificial Intelligence) empirical evaluations of new algorithms to confirm their effectiveness and reliability is even more essential. This trend is intensifying with the increasing complexity of methods, particularly with the emergence of deep neural networks, generative AI, and large language models, which are difficult to explain and interpret. Empirical evaluation is essential, in particular because of the complexity of the algorithms and the unpredictable nature of the data.

The approach taken in ths pillar is that of organizing scientific competitions (also called "challenges"). Scientific competitions systematize large-scale experiments and show the effectiveness of participants in solving complex problems. Annual competitions, organized on the Codalab competition platform, address various scientific or industrial questions, evaluating the automatic algorithms submitted by participants. The newer version of Codalab, called Codabench, extends the capabilities of Codalab to benchmarks.

Both challenges and benchmarks are crucial for comparing models and understanding their behavior. Recent applications include: improving decision-making, particularly useful in fields like finance and medicine; helping to combat climate change by optimizing the use of resources; personalizing the customer experience in e-commerce, banking, and other industries; improving security and preventing fraud; and improving accessibility for people with disabilities, for example, through voice recognition systems, visual aids for the visually impaired, and other assistive technologies.

The importance of impartial evaluations of algorithms is constantly increasing with the acceleration of progress in Artificial Intelligence. According to David Donoho: "The emergence of Frictionless Reproducibility flows from 3 data science principles that matured together after decades of work by many technologists and numerous research communities. The mature principles involve data sharing, code sharing, and competitive challenges, however implemented in the particularly strong form of frictionless open services." He cites the Codalab project as being exemplary in this area [99].

4 Application domains

4.1 Computational Social Sciences

Computational Social Sciences (CSS) studies social and economic phenomena, ranging from technological innovation to politics, from media to social networks, from human resources to education, from inequalities to health. It combines perspectives from different scientific disciplines, building upon the tradition of computer simulation and modeling of complex social systems [105] on the one hand, and data science on the other hand, fueled by the capacity to collect and analyze massive amounts of digital data.

The emerging field of CSS raises formidable challenges along three dimensions. Firstly, the definition of the research questions, the formulation of hypotheses and the validation of the results require a tight pluridisciplinary interaction and dialogue between researchers from different backgrounds. Secondly, the development of CSS is a touchstone for ethical AI. On the one hand, CSS gains ground in major, data-rich private companies; on the other hand, public researchers around the world are engaging in an effort to use it for the benefit of society as a whole [112]. The key technical difficulties relate to data and model biases, and to self-fulfilling prophecies. Thirdly, CSS does not only regard scientists: it is essential that the civil society participate in the science of society [127].

TAO/TAU was involved in CSS for the last five years, and its activities had been strengthened thanks to P. Tubaro's and I. Guyon's expertises respectively in sociology and economics, and in causal modeling. Their departures has negatively impacted the team activities in this domain, but many projects are still on-going and CSS remains a domain of choice (see Section 8.6).

4.2 Energy Management

Energy Management has been an application domain of choice for TAO since the mid 2000s, with main partners SME Artelys (METIS Ilab INRIA; ADEME projects POST and NEXT), RTE (three CIFRE PhDs), and IFPEN (bilateral contract, DATAIA project ML4CFD). The goals concern i) optimal planning over several spatio-temporal scales, from investments on continental Europe/North Africa grid at the decade scale (POST), to daily planning of local or regional power networks (NEXT); ii) monitoring and control of the French grid enforcing the prevention of power breaks (RTE); iii) improvement of house-made numerical methods using data-intense learning in all aspects of IFPEN activities (Section 8.4.2).

The daily maintainance of power grids requires the building of approximate predictive models on the top of any given network topology. Deep Networks are natural candidates for such modelling, considering the size of the French grid (~ 10000 nodes), but the representation of the topology is a challenge when, e.g. the RTE goal is to quickly ensure the "n-1" security constraint (the network should remain safe even if any of the 10000 nodes fails). Existing simulators are too slow to be used in real time, and the size of actual grids makes it intractable to train surrogate models for all possible (n-1) topologies (see Section 8.5 for more details).

Another aspect of Power Grid management regards the real-time control of the grid topology, man-made at the moment. Its automation is yet a difficult challenge, but results on the L2RPN challenge have demonstrated its feasibility with Reinforcement Learning, opening the way to more ambitious goals (e.g., decentralized control via multi-agent Reinforcement Learning, see Section 8.5).

4.3 Data-driven Numerical Modeling

In domains where both first principle-based models and equations, and empirical or simulated data are available, their combined usage can support more accurate modelling and prediction, and when appropriate, optimization, control and design, and help improving the time-to-design chain through fast interactions between the simulation, optimization, control and design stages. The expected advances regard: i) the quality of the models or simulators (through data assimilation, e.g. coupling first principles and data, or repairing/extending closed-form models); ii) the exploitation of data derived from different distributions and/or related phenomenons; and, most interestingly, iii) the task of optimal design and the assessment of the resulting designs.

A first challenge regards the design of the model space, and the architecture used to enforce the known domain properties (symmetries, invariance operators, temporal structures). When appropriate, data from different distributions (e.g. simulated vs real-world data) will be reconciled, for instance taking inspiration from real-valued non-volume preserving transformations [95] in order to preserve the natural interpretation.

Another challenge regards the validation of the models and solutions of the optimal design problems. The more flexible the models, the more intensive the validation must be. Along this way, generative models will be used to support the design of "what if" scenarios, to enhance anomaly detection and monitoring via refined likelihood criteria.

In the application domains described by Partial Differential Equations (PDEs), the goal of incorporating machine learning into classical simulators is to speed up the simulations while maintaining as much as possible the accuracy ad physical relevance of the proposed solutions. Many possible tracks are possible for this; one can build surrogate models, either of the whole system, or of its most computationaly costly parts; one can search to provide better initialization heuristics to numerical solvers, which make sure that physical constraints are satisfied. Or one can inject physical knowledge/constraints at different stages of the numerical solver.

5 Social and environmental responsibility

5.1 Footprint of research activities

The Laboratory (LISN) is currently actively re-thinking its carbon footprint, being part of the Labo1.5 initiative. We participate in working groups about GreenAI (being able to measure, compare and mitigate the negative impact of training and inference for large models). To start changing practices, the simple fact of reporting the cost of training one's model in publications has been spotted as en efficient tool. Ideally, the development cost (all the trainings performed during the research, not just the training of the model presented in the paper) should also be mentionned.

Another axis studied by the lab is the limitation of (aerial) transport, keeping in mind that the younger members should be allowed to build their own research network and foreign experiences.

5.2 Impact of research results

All our work on Energy (see Sections 4.2) is ultimately targeted toward optimizing the distribution of electricity, be it in planning the investments in the power network by more accurate previsions of user consumption, or helping the operators of RTE to maintain the French Grid in optimal conditions.

A collaboration with IDEEV has just started, with the idea of leveraging Deep Learning as a tool to help unlock agro-ecological research. In particular, we aim to help measure the yields in mixed cropping (requiring to be able to classify grains of a given species but of different varieties – something impossible to the naked eye) and detect pollinators on video footage taken outside (including wind, change in light conditions, etc).

6 Highlights of the year

6.1 Prestigous Publications

In 2024, the team has successfully submitted papers in the most prestigious ML venues:

- Two papers at IJCAI 2024, one in the main track [27] and one in the survey track [40];
- Three papers at NeurIPS 2024, one selected as spotlight in the main track [37], one in the main track [29] and one in the competition track [48];
- Three papers at ICLR 2025 [42, 30, 28]
- One paper at AAAI 2025 [32]
- One paper has been published in TMLR [26]

6.2 Awards

Isabelle Guyon (who left the team in 2023) was elected at the French *Académie des technologies*

6.3 Spin-off

Three former PhD students of the team, Emmanuel Menier, Matthieu Nastorg and Alice Lacan, launched the startup company AUGUR, proposing to build foundational models for numerical simulations.

7 New software, platforms, open data

7.1 New software

7.1.1 Codalab

Keywords: Benchmarking, Competition

Functional Description: Challenges in machine learning and data science are competitions running over several weeks or months to resolve problems using provided datasets or simulated environments. Challenges can be thought of as crowdsourcing, benchmarking, and communication tools. They have been used for decades to test and compare competing solutions in machine learning in a fair and controlled way, to eliminate "inventor-evaluator" bias, and to stimulate the scientific community while promoting reproducible science. Current production infrastructure has been consolidated in 2021 (sovereign distributed storage, 20 GPU workers) thanks to the sponsorship of Région Ile-de-France, ANR, Université Paris-Saclay, CNRS, INRIA, and ChaLearn, to support 20,000 new users (2024), organizing or participating each year to hundreds of competitions. Some of the areas in which Codalab is used include Computer vision and medical image analysis, natural language processing, time series prediction, causality, and automatic machine learning. Codalab has been selected by the Région Ile de France to organize industry-scale challenges. Codalab has been ranked first on scientific criteria, in an independent international study: https://mlcontests.com/state-of-competitive-machine-learning-2023/. TAU continues expanding Codalab to accommodate new needs, including teaching. Link to the historical server (read-only) https://competitions.codalab.org.

@article{codalab_competitions_JMLR, author = {Adrien Pavao and Isabelle Guyon and Anne-Catherine Letournel and Dinh-Tuan Tran and Xavier Baro and Hugo Jair Escalante and Sergio Escalera and Tyler Thomas and Zhen Xu}, title = {CodaLab Competitions: An Open Source Platform to Organize Scientific Challenges}, journal = {Journal of Machine Learning Research}, year = {2023}, volume = {24}, number = {198}, pages = {1–6}, url = {http://jmlr.org/papers/v24/21-1436.html} }

URL: http://competitions.codalab.org

Contact: Isabelle Guyon

7.1.2 Cartolabe

Name: Cartolabe

Keyword: Information visualization

Functional Description: The goal of Cartolabe is to build a visual map representing the scientific activity of an institution/university/domain from published articles and reports. Using the HAL Database, Cartolabe provides the user with a map of the thematics, authors, and articles. ML techniques are used for dimensionality reduction, cluster, and topic identification, visualization techniques are used for a scalable 2D representation of the results.

Cartolabe has, in particular, been applied to the Grand Debat dataset (3M individual propositions from French Citizen, see https://cartolabe.fr/map/debat). The results were used to test both the scaling capabilities of Cartolabe and its flexibility to non-scientific and non-English corpora. We also added sub-map capabilities to display the result of a year/lab/word filtering as an online generated heatmap with only the filtered points to facilitate the exploration. Cartolabe has also been applied in 2020 to the COVID-19 Kaggle publication dataset (Cartolabe-COVID project) to explore these publications.

URL: http://www.cartolabe.fr/

Publication: hal-02499006

Contact: Philippe Caillou

Participants: Philippe Caillou, Jean Daniel Fekete, Michèle Sebag, Anne-Catherine Letournel, Hande Gozukan

Partners: LRI - Laboratoire de Recherche en Informatique, CNRS

7.1.3 DeepHyper

Keywords: Deep learning, Autotuning, HPC

Functional Description: Machine learning algorithms are continually evolving to serve diverse applications, yet their development often entails a significant trial-and-error process to identify optimal learning pipelines.This is compounded by the multitude of data preprocessing techniques, prediction (or generative) models, and learning procedures available, each offering a range of configurable parameters, also referred to as hyperparameters. DeepHyper addresses this challenge by automating the selection and configuration of algorithms and their corresponding hyperparameters, facilitating a streamlined approach for engineers and scientists to comprehend and optimize the learning pipeline. At its core, DeepHyper employs parallel Bayesian optimization, validated through rigorous testing involving up to 8,000 parallel tasks. This methodology is adaptable for both single and multi-objective tasks, enabling efficient early discarding of costly training steps. Furthermore, DeepHyper seamlessly integrates with various parallel backends, including multi-threading, multi-processing, Clouds (via the Ray library), and MPI-based schedulers on supercomputers, enhancing its scalability and versatility across different computing environments. The development of DeepHyper is supported by the TAU-team through advances in learning theory for improving and explaining its core algorithms.

URL: https://github.com/deephyper/deephyper

Contact: Romain Egele

7.1.4 OmniPrint

Keyword: Open data

Functional Description: Benchmarks and shared datasets have been fostering progress in deep learning. While there is an increasing number of available datasets, there is a need for larger ones. However, collecting and labeling data is time-consuming and expensive, and systematically varying environmental conditions is difficult and necessarily limited. Therefore, resorting to artificially generated data is helpful to drive fundamental research in deep learning. OmniPrint is geared to generating an unlimited amount of printed characters.

Character images provide excellent benchmarks for deep learning problems because of their relative simplicity and visual nature while opening the door to high-impact real-life applications. A conjunction of technical features is required to meet our specifications: pre-rasterization manipulation of anchor points, post-rasterization distortions, natural background and seamless blending, foreground filling, anti-aliasing rendering, and importing new fonts and styles. Modern fonts such as TrueType or OpenType are made of straight line segments and quadratic Bezier curves, connecting anchor points. Thus, it is easy to modify characters by moving anchor points. This allows users to perform vectors-space pre-rasterization geometric transforms (rotation, shear, etc.) as well as distortions (e.g., modifying the length of ascenders of descenders) without incurring aberrations due to aliasing when transformations are done in pixel space (post-rasterization).

The key technical contributions include implementing transformations and styles such as elastic distortions, natural background, foreground filling, and so on, selecting characters from the Unicode standard to form alphabets from more than 20 languages around the world, further grouped into partitions, to facilitate creating meta-learning tasks, identifying fonts, implementing character rendering with a low-level FreeType font rasterization engine, which enables direct manipulation of anchor points, adding anti-aliasing rendering, implementing and optimizing utility code to facilitate dataset formatting. To our knowledge, OmniPrint is the pioneering text image synthesizer geared toward ML research, supporting pre-rasterization transforms, which allows Omniprint to imitate handwritten characters to some degree. More details can be found in the paper (https://openreview.net/forum?id=R07XwJPmgpl, https://arxiv.org/abs/2201.06648).

URL: https://github.com/SunHaozhe/OmniPrint

Contact: Haozhe Sun

7.1.5 codabench

Keywords: Competition, Benchmarking

Functional Description: Obtaining standardized crowdsourced benchmark of computational methods is a major issue in data science communities. Dedicated frameworks enabling fair benchmarking in a unified environment are yet to be developed. Here we introduce Codabench, an open-source, community-driven platform for benchmarking algorithms or software agents versus datasets or tasks. Codabench, released in summer 2023, is the follower of Codalab, enabling the same features and more: inverted data challenges, better user experience, easier platform administration and robustness. Competition design is backward compatible allowing an easy migration from Codalab to Codabench. A public instance of Codabench (https://codabench.org) is open to everyone, free of charge, and allows benchmark organizers to compare fairly submissions, under the same setting (software, hardware, data, algorithms), with custom protocols and data formats. Codabench has unique features facilitating the organization of benchmarks flexibly, easily and reproducibly, such as the possibility of re-using templates of benchmarks, and supplying compute resources on-demand. In 2024, Codabench has registered more than 10,000 new users and computed near 80,000 participants submissions.

@article{codabench, title = {Codabench: Flexible, easy-to-use, and reproducible meta-benchmark platform}, author = {Zhen Xu and Sergio Escalera and Adrien Pavão and Magali Richard and Wei-Wei Tu and Quanming Yao and Huan Zhao and Isabelle Guyon}, journal = {Patterns}, volume = {3}, number = {7}, pages = {100543}, year = {2022}, issn = {2666-3899}, doi = {https://doi.org/10.1016/j.patter.2022.100543}, url = {https://www.sciencedirect.com/science/article/pii/S2666389922001465} }

URL: https://www.codabench.org/

Contact: Isabelle Guyon

Partner: Région Île-de-France

7.1.6 pyriemann-qiskit

Keywords: Quantum programming, Riemannian geometry, Symmetric positive definite matrices

Functional Description: Literature on quantum computing suggests it may offer an advantage compared with classical computing in terms of computational time and outcomes, such as for pattern recognition or when using limited training sets. Bulding on the Qiskit library on quantum computing, pyriemann-qiskit implements a wrapper around quantum-enhanced support vector classifiers (QSVCs) and variational quantum classifiers (VQCs), to use quantum classification with Riemannian geometry. It also introduces a quantum version of the MDM algorithm, a classifier operating on the manifold of symmetric positive definite matrices.

URL: https://pyriemann-qiskit.readthedocs.io/en/latest/

Publication: https://hal.science/hal-04040814v1

Contact: Sylvain Chevallier

Partner: IBM

7.1.7 pyriemann

Keywords: Riemannian geometry, Hermitian positive definite matrices, Symmetric positive definite matrices

Functional Description: Pyriemann is a Python machine learning package based on scikit-learn API. It provides a high-level interface for processing and classification of real (resp. complex)-valued multivariate data through the Riemannian geometry of symmetric (resp. Hermitian) positive definite (SPD) (resp. HPD) matrices.

pyRiemann aims at being a generic package for multivariate data analysis but has been designed around biosignals (like EEG, MEG or EMG) manipulation applied to brain-computer interface (BCI), estimating covariance matrices from multichannel time series, and classifying them using the Riemannian geometry of SPD matrices. It is widely used in the scientific community with more than one million download.

URL: https://pyriemann.readthedocs.io

Contact: Sylvain Chevallier

7.1.8 braindecode

Keywords: Brain-Computer Interface, Deep learning

Functional Description: BrainDecode is an open-source Python toolbox for decoding raw electrophysiological brain data with deep learning models. It includes dataset fetchers, data preprocessing and visualization tools, as well as implementations of several deep learning architectures and data augmentations for analysis of EEG, ECoG and MEG. It is design for neuroscientists who want to work with deep learning and deep learning researchers who want to work with neurophysiological data.

URL: https://braindecode.org/stable/index.html

Contact: Sylvain Chevallier

Partner: Roche

7.1.9 MOABB

Name: Mother of all BCI Benchmarks

Keywords: Brain-Computer Interface, Open data, Benchmarking

Functional Description: Mother of all BCI Benchmarks (MOABB) allows to build a comprehensive benchmark of popular brain-computer interface algorithms applied on an extensive list of freely available EEG datasets. This is an open science initiative, serving as a reference point for the future algorithmic developments. Build on reference libraries like scikit-learn and MNE-python, machine learning pipelines can be ranked and promoted on a website, providing a clear picture of the different solutions available in the field. This software has 80k downloads and an active international development community.

URL: https://neurotechx.github.io/moabb/

Contact: Sylvain Chevallier

7.1.10 dnadna

Name: Deep Neural Architectures for DNA

Keywords: Deep learning, Population genetics

Functional Description: DNADNA provides utility functions to improve development of neural networks for population genetics and is currently based on PyTorch. In particular, it already implements several neural networks that allow inferring demographic and adaptive history from genetic data. Pre-trained networks can be used directly on real/simulated genetic polymorphism data for prediction. Implemented networks can also be optimized based on user-specified training sets and/or tasks. Finally, any user can implement new architectures and tasks, while benefiting from DNADNA input/output, network optimization, and test environment.

URL: https://mlgenetics.gitlab.io/dnadna/

Contact: Flora Jay

7.2 New platforms

Participants: Isabelle Guyon, Anne-Catherine Letournel, Adrien Pavao, Hande Gozukan.

- **CODALAB:** The TAU group is community lead (under the leadership of Isabelle Guyon) of the open-source Codalab project, hosted by Universite‘ Paris-Saclay, whose goal is to host competitions and benchmarks in machine learning [120]. We have replaced the historical server by a dedicated server hosted in our lab. Since inception in December 2021, over 40000 participants entered 640 public competitions (see statistics). The engineering team, overseen by Anne-Catherine Letournel (CNRS engineer) includes two engineers dedicated full time to administering the platform and developing challenges: Adrien Pavao, financed by a project started in 2020 with the Re‘gion Ile-de-France, et Dinh-Tuan Tran, financed by the ANR AI chaire of Isabelle Guyon, Ihsan Ullah, financed by a collaboration with LBNL/CERN and IJCLAB, and Benjamn Bearce financed by the ANR AI chaire of Isabelle Guyon. Several other engineers are engaged as contractors on a needs-be basis. The rapid growth in usage led us to put in place a new infrastructure. We have migrated the storage over a distributed Minio (4 physical servers, each with 12 disks of 16 TB) spread over 2 buildings for robustness, and added 10 more GPUs to the existing 10 previous ones in the backend. A lot of horsepower to suport Industry-strength challenges, thanks for the sponsorship of re‘gion Ile-de-France, ANR, Universite‘ Paris-Saclay, CNRS, INRIA, and ChaLearn.

- **CODABENCH:** Codabench [129] is a new version of Codalab emphasizing the orgnization of benchmarks, which can be thought of as ever-lasting challenges, de-emphasizing competiton, and favoring the comparison between algorithms. Codabench has also all the capabilities of Codalab and will progressively replace it. When Codabench is fully stable, we will retire Codalab.

  The V1 of Codabench was launched in August 2023. The user base is rapidly growing (over 3000 users, 67 public competitions, and 25000 submissions)

8 New results

8.1 Frugal Learning

Participants: Guillaume Charpiat, Isabelle Guyon, Alessandro Ferreira Leite, Marc Schoenauer, Michèle Sebag, Sylvain Chevallier, Alice Lacan, Romain Egele, Manon Verbockhaven, François Landes, Maria Sayu Yamamoto, Bruno Aristimunha Pinto, Blaise Hanczar(Univ. Evry).

8.1.1 Model frugality

In Manon Verbockhaven's PhD thesis, we study how to optimally grow a neural network architecture, to increase the performance (in terms of loss) while keeping the network as small as possible (in particular, avoiding redundancy). We showed [26] how to formulate the notion of "expressivity bottleneck" in an easily computable manner, and obtain optimal neuron weights as the result of a small SVD. We showed that the approach can scale up, with an experiment using ResNet18 on CIFAR-100. With Barbara Hajdarevic's internship, we had also started to extend the addition of neurons to existing layers to the addition of layers to an existing computation graph. Thanks to the European project MANOLO and to an ENS grant (CDSN), this work has been continued by new PhD students, Styliani Douka and Théo Rudkiewicz, and a post-doc, Stéphane Rivaud. In particular, we now allow the architecture to grow as an arbitrary DAG (Directed Acyclic Graph) [33] (paper accepted at ESANN 2025).

Within the context of Alice Lacan's PhD (defended on Feb 4., 2025; coll. U. Evry), we applied generative modelling for data augmentation in transcriptomics [110]. The high computational requirements of mainstream generative models (GAN, WGAN, diffusion model) for such high dimensional domains led to the design of a new frugal generative modelling approach, based on density alignment [34].

Within the context of Nicolas Atienza's PhD (to be defended in March 2025; Cifre Thales, coll. GALAC@LISN), the latent space of a trained teacher is decomposed using Information Bottleneck principles. The core latent space is exploited to learn a frugal student, by distillation of the teacher.

8.1.2 Data frugality

Frugal learning is also investigated along the line of the design of data frugal algorithms. It is a major challenge in the context of domain adaptation and transfer learning, to ensure that relevant representation are learned from limited data accessibility. In the context of time serie prediction and brain signal decoding, [10] reduces the number of sensors required to obtain state-of-the-art results and [46] further explore Riemannian deep learning for domain adaptation. Riemannian models could also encompass more complex representations, using trajectory in the space of spectral frequency, to generate more robust decoding of time series [44].

Another line of research concerns the dataset alignment methods: [17] proposes a systematic exploration to avoid negative transfer effects, [41] adds data augmentation in the alignment process, [36] integrates physics-informed constraints for aligning datasets while mitigating the heterogeneity in the dimension of the data, or fighting dataset shift both in time series and labels [37].

Also, within the context of Vincenzo Schimmenti's PhD [124], we demonstrate [25] that the information contained in GPS stations, that monitor the Earth surface deformation, can be leveraged to predict aftershocks to large earthquakes, reaching a balanced accuracy of 70%. This is true both for a very robust model (logistic regression, 2 parameters) and for our proof-of-concept CNN, which we manage to avoid to overfit using a robust ensembling technique, despite the very small number of training samples (48 spatial maps only).

8.2 Toward Good AI

Participants: Philippe Caillou, Isabelle Guyon, Alessandro Leite, Michèle Sebag, Sylvain Chevallier, Flora Jay, Cyril Furtlehner, Guillaume Charpiat, Aurelien Decelle, Armand Lacombe, Cyriaque Rousselot, Nicolas Atienza, Romain Egele, Haozhe Sun, Shuyu Dong, Antoine Szatkownik, Olivier Allais (INRAE), Julia Mink (Univ. Bonn), Jean Pierre Nadal (CAMS EHESS), Annick Vignes (CAMS EHESS), David Lopez-Paz (Facebook/Meta), Burak Yelmen(U Tartu) .

8.2.1 Causal Learning

Causal learning is commonly regarded as a key research direction to enforce the properties of good AI models in terms of explainability, verifiability and fairness. Its importance is acknowledged through the PEPR-IA-Causalit-AI, starting in 2024, gathering four French laboratories/teams (Loria at Nancy; TAU and CELESTE at UPSaclay; LIG at Grenoble)

In the last year of Shuyu Dong's postdoc (partnership Fujitsu), building upon previous results [96, 97], we have tackled the notorious lack of scalability of causal graph learning from observational data. The proposed approach, called DCILP, is a Divide-and-Conquer approach. Formally, sub-problems involving the Markov blanket of each variable are defined and solved (thus with moderate complexity); the reconciliation of these partial solutions is formulated as an integer linear programming problem, reaching a very good trade-off between time-complexity and accuracy [32], recently accepted at AAAI'25.

Audrey Poinsot's PhD (Cifre Ekimetrics) is concerned with counter-factual reasoning in the context of strategic and marketing decisions. Because data in that area does not pertain to Big Data, the PhD first focused on Data Augmentation in Causal context [121], as known causality links can be leveraged to ease learning. But the lack of recognized benchmarks in that area led us to propose a comprehensive survey of deep structural causal models, focusing on their ability to answer counterfactual queries using observational data within known causal structures [40].

Armand Lacombe's PhD, defended on March 5th, 2024 [52], tackles the identification of the conditional average treatment effet in the case where the control and treated distributions are different (which is the usual case in practice). A thorough theoretical and empirical analysis have been proposed for the original approach, relying on the design of two latent representations.

A collaboration with Inria Nerv team, from Paris-Brain Institute, investigates the causal chains found in the brain activity. Motor imagery and preparation of movement executions is build on sequential activations of neural assemblies within the brain. In [13], an avalanche-based approach is designed to reconstruct those causal chain of activations in brain imaging.

Causality is also at the core of TAU participation in the INRIA Challenge *OceanIA*, that started in 2021 [82]. The main challenge is related to out-of-distribution learning, motivated by the analysis of the TARA images to identify the ecosystems in the diverse sites of the data collection. The high imbalance of the data among the classes, the prevalence of outliers, suggest that the use of multi-modal embeddings as explored in N. Atienza's PhD [27] (see below) might support the design of relevant metrics in the considered space. A post-doc has been hired, who will start in February 2025.

8.2.2 Explainable Learning

Nicolas Atienza's PhD (Cifre Thales, co-supervised with Johanne Cohen, LISN; 2 patents Thales pending) is to tackle the main three goals of Trustable AI, i.e. explainability, reliability and frugality. Toward explainability, a new approach to build conceptual model explanations, called CB2 (Cut the Black Box) proceeds by combining multi-modal embeddings and Multi-Criteria Decision Aid [27]. Toward reliability, the new Sample-efficient Probabilistic Detection using Extreme Value Theory (SPADE) transforms a classifier into an abstaining classifier, offering provable protection against out-of-distribution and adversarial samples.

During the last year of the Horizon Europe project TRUST-AI, and even though our work within this project was completed, we extended the results obtained on Memetic Semantic Genetic Programming (MSGP) [114]. MSGP is able to generate short, and hence hopefully easily explainable expressions for Symbolic Regression problems. We implemented a boosting procedure around MSGP, which improved the performances without degrading too much the interpretability [19].

Also, note that the work described in Section 8.4.1 about Interpretable Learning Effective Dynamics (iLED) framework [21] also contributes to this line of research, adding interpretability to the Deep Learning approach to dynamical systems simulations.

8.2.3 Improved Learning

The research in AutoML is gently fading out at TAU, with Isabelle Guyon's departure for Google Brain and Michèle Sebag's and Marc Schoenauer's retirements (even if emeritus, they cannot supervise new PhD students nor can they be PIs for new projects). On the other hand, the long-lasting expertise of the team in terms of black-box optimization proved to be useful in order to globaly improve the learning process, a priori or post-hoc.

Romain Egele's PhD (coll Argonne National Labs, USA), defended in June 2024 [51] was focused on Hyper-Parameter Optimization (HPO) and Neural Architecture Search (NAS), and the deployment of AI and HPC. The last contrbution in his thesis [60] complements his previous work in developing DeepHyper, a package allowing users to conduct NAS with genetic algorithms using TensorFlow or PyTorch [103]. A benchmark of early discarding strategies was conducted to compare state of the art algorithms. It was noticed that a very simple strategy, dubbed 1-Epoch, performed significantly better when "computing duration" is the bottleneck. A method based on Bayesian regression (including both aleatoric and epistemic uncertainties) for learning curve extrapolation was also proposed and dubbed Robust Bayesian Early Rejection (RoBER) [102]; Best paper Award at IEEE International Conference on e-Science.

Mathurin Videau is a CIFRE PhD student with Meta (that started long before Meta had to obey Musk's worst practices), co-supervised by Olivier Teytaud (former member of the team before joining Meta). Mathurin's PhD focuses on post-hoc improvements of fully trained models by using black-box algorithms to optimize an impactful but small part of model. Using BBO allows us to optimize non-differentiable loss functions that match the user's true goals more efficiently than the loss used for the initial training of the model by standard backpropagation, loss that needs to be differentiable [71] (submitted). This allows for instance to exactly optimize the Word Error Rate in translation tasks, the number of deaths in Doom RL agents, or to let the user interactively guide generative processes by directly acting in the latent space [73] (submitted).

8.2.4 Towards high-quality and private genomes based on generative neural networks

In collaboration with the Institute of Genomics of Tartu, we have been leveraging two types of generative neural networks (Generative Adversarial Networks and Restricted Boltzmann Machines) to learn the high dimensional distributions of real genomic datasets and create artificial genomes [77]. These artificial genomes retain important characteristics of the real genomes (genetic allele frequencies and linkage, hidden population structure, ...) without copying them and have the potential to be valuable assets in future genetic studies by providing anonymous substitutes for private databases (such as the ones hold by companies or public institutes like the Institute of Genomics of Tartu).

The main challenges lie in scaling up to the full genome and in making sure that no personal genetic data is leaked. For this, we had developed various deep learning generative architectures, from plain GANs and RBMs, to convolutional GANs, with or without attention. Following this body of work where models are trained in the SNP data space (ie., the space of DNA sequences, removing sites that are constant accross the dataset), we propose in [45] a conceptually different approach. Our method combines dimensionality reduction, achieved by Principal Component Analysis (PCA), and a Generative Adversarial Network (GAN) learning in this reduced space, which is much smaller when facing datasets with fewer individuals (5 000) than SNP sequence length (60 000). Such a low ratio between number of samples and sample dimension makes the task proned to overfitting and calls for careful check of possible privacy leaks. We studied various privacy scores, including in particular the AATS metric (nearest neighbor adversarial accuracy) proposed by [130], and sorted the different models according to quality and privacy scores.

Furthermore, we proposed an alternative approach based on a diffusion model, which had never been investigated in the context of population genetics [68]. We tested how these synthetic genomes could replace or augment reference databases for local ancestry inference (LAI). LAI is a major task in population genetics which consists in identifying the origin of genomic regions along the chromosomes of a given indiviual. It relies on reference panels that should be as diverse as possible to avoid biaises and could thus highly benefit from accurate data augmentation, in particular for underrepresented populations.

A next challenging step is to design interpretable generative models capable of handling genotypes and phenotypes jointly. We have indeed recently shown, in a supervised setting, that classical neural networks combined with post-hoc interpretation techniques yielded insights on the relationships linking genetic loci and phenotypes [72]. In parallel, we aim to investigate the potential of generative modeling for this task (collaboration with U Tartu).

8.3 Machine Learning with/for Statistical Physics

Participants: Cyril Furtlehner, François Landes, Beatriz Seoane, Guillaume Charpiat, Michele Sebag, Anaclara Alvez-Canepa, Nicolas Béreux, Nilo Schwencke, Emmanuel Goutierre, Decelle Aurélien (UCM), Catania Giovanni (UCM), Rahul Chako (external post-doc), Andrea Liu (UPenn), David Reichman (Columbia), Johanne Cohen (LISN), Christelle Bruni (IJCLAB), Hayg Guler (IJCLAB).

Generative models constitute an important piece of unsupervised ML techniques, which is under rapid development. In this context, insights from statistical physics are important, especially for energy-based models such as restricted Boltzmann machines. Over the years we have contributed to build a global picture of the Restricted Boltzmann Machine (RBM) and to identify the main hurdles: the information content of a trained RBM and its learning dynamics can be precisely analyzed using ensemble averaging techniques [93, 94]. We have also described in great detail the effects of inadequate MCMC sampling on the quality and performance of RBMs [80]. The spectral dynamics reveals that learning materializes itself by the emergence of new modes in the weight matrix, each one being accompagnied by a second order phase transition which are further characterized in [29]. A second important observation made in [76] is that for structured data the learning process takes place along a first order transition line which renders sampling inefficient even for advanced methods like parallel tempering. The subject of Nicolas Béreux's

PhD is to address these difficulties, and in [30] an efficient approach has been proposed which rely on: (i) designing a pre-training strategy allowing us to bypass the most severe 2nd order phase transitions, based on the mapping between the RBM and the Coulomb machine proposed in [76]; (ii) introducing a novel framework for estimating log-likelihood (LL) by leveraging the learning trajectory's softness, rather than relying on temperature integration; (iii) setting a variation of the standard parallel tempering algorithm in which exchanges occur between the parameters of models trained at different stages, rather than across temperatures thereby avoiding to cross the first order transition line. Overall this allows us to train equilibrium models for a broad range of structured data. This, in combination with a second work [14], which determines how the parameters of a Bernoulli or Ising RBM can be mapped onto general Hamiltonian, opens the possibility for obtaining highly interpretable generative models well suited for scientific data. In complement of that we analyze cooperative and federated learning via the copycat perceptron model [11] finding that under the teacher-student framework learning is improved under some conditions characterized in terms of phase diagram.

Physics informed machine learning is also an important axis of research in the team with several avenues. One is concerned with learning critical phenomena, i.e. phenomena displaying specific scaling properties, where typically all scales contribute. This is the subject of Anaclara Alvez thesis, investigating the question of how to exploit scale invariance for processes dispaying statistical self-similarity, like avalanche processes which the objective of being able to extrapolate the predictions from small to large scales. The second avenue deals with physics informed neural networks (PINN's) [122], which is an appealing way to solve PDE with by inserting the physics into the loss function and which in principle is mesh-free. Unfortunately, this method which is still in its infancy, is plagued by many shortcomings and failures that remain not properly understood. In the context of Nilo Schwencke's thesis we have proposed ANaGRAM [42] which adresses some of these issues, in particular the one concerning spectral biases, with four contributions: (i) an extension of neural tangent kernel theory (NTK) [107] which introduce the notions of empirical tangent space and empirical natural gradient leading to a family of algorithms (Anagram) which depends on the way the projection of the functional gradient on the empirical tangent space is approximated; (ii) a key relation showing that our formulation of Natural Gradient for PINNs coincides with the operators Green function restricted to the tangent space; (iii) an efficient implementation of the simplest instantiation of Anagram, which can be seen as a combination of Gauss-Newton with SVD adapted to PINNs, with good scaling properties, showing robustness and superior empirical results to existing baseline; (iv) a new, simple and principled optimization criteria for the collocation point problem, which is a direct byproduct of our theoretical framework.

As mentioned earlier, the use of ML to address fundamental physics problems is quickly growing. A place where ML can help address fundamental physics questions is the domain of glasses (how the structure of glasses is related to their dynamics), which is one of the major problems in modern theoretical physics [85] and play a key role in Giorgio Parisi's career (2021 Nobel prize laureate). This year, with controlled numerical experiments, we clarified the important role of Dynamical Facilitation in the melting or equilibration of glasses, discarding first-order transition style analogies [12]. There are various ways in which ML can help address fundamental questions about the physics of glasses, that we reviewed in a Roadmap paper [16]. Our angle is to learn the hidden structures (features) that control the flowing or non-flowing state of matter, discriminating liquid from solid states, using rotation-equivariant neural networks. We prove that rotation-equivariant GNNs outperform other approaches in terms of generalization power, displaying especially good generalization to unseen temperatures [24]. Our approach was benchmarked against other recent works in the roadmap [16], confirming that our approach is extremely promising; we currently are actively exploring this avenue of research. The main PhD student carrying out this research defended in 2024 [56].

A parallel line of research consists in using replica computation (a tool from statistical physics) to compute the whole statistics of possibly learned models, in simple settings. This has been used to show that the optimal training imbalance is different from 0.5, in a simplified setup of Anomaly Detection [56] (there is also a corresponding AISTATS paper that was just accepted, and is not yet published).

8.4 Machine Learning for Numerical Simulations

8.4.1 ML and Reduced Order Models for Dynamical Systems

Participants: Michele Alessandro Bucci, Marc Schoenauer, Emmanuel Menier, Thibault Monsel, Mouadh Yagoubi (IRT-SystemX), Lionel Mathelin (DATAFLOT team, LISN), Onofrio Semeraro (DATAFLOT team, LISN), Petros Koumoutsakos (Harvard SEAS), Sebastian Kaltenbach (Harvard SEAS).

Numerical simulations of fluid dynamics in industrial applications require the spatial discretization of complex 3D geometries with consequent demanding computational operations for the PDE integration. The computational cost is mitigated by the formulation of Reduced Order Models (ROMs) aiming at describing the flow dynamics in a low dimensional feature space. The Galerkin projection of the driving equations onto a meaningful orthonormal basis speeds up the numerical simulations but introduces numerical errors linked to the underrepresentation of dissipative mechanisms.

Emmanuel Menier's PhD, defended in January 2024 [54] trained a DNN to compensate missing information in the projection basis. By exploiting the projection operation, the ROM correction consists in a forcing term in the reduced dynamical system which has to: i) recover the information living in the subspace orthonormal to the projection one; ii) ensure that its dynamics is dissipative. A constrained optimization is then employed to minimize the ROM errors but also to ensure the reconstruction and the dissipative nature of the forcing improving the prediction while preserving the guarantees of the ROM. The approach was extended on Michelin use case of rubber calendering process.

During his PhD thesis, Emmanuel Menier also spent 3 months in Spring 2023 in Prof. Petros Koumoutsakos' group at SEAS - Harvard, John A. Paulson School of Engineering and Applied Sciences. It was a perfect match between his previous work and the group's expertise in high dimensional dynamical complex system (e.g., CFD), and resulted in the Interpretable Learning Effective Dynamics (iLED) framework, a novel framework based on nonlinear dimension reduction thanks to deep neural networks, that offers comparable accuracy to state-of-the-art recurrent neural network-based approaches while providing the added benefit of interpretability [21]. The basic idea of iLED is grounded on the Mori-Zwanzig formalism, an approach that has been later generalized to other dynamical systems [119].

Thibault Monsel's PhD has been indeed focusing on the learning of dynamical sytems involving delays, i.e. Delayed Differential Equations (DDE). While Neural ODE was a conceptual breakthrough, it cannot learn partially-observable dynamical systems. Inspired by the Mori-Zwanzig formalism and Takens theorem, we develop another way to extract this information, using delays, i.e. using past observable states of the system [119, 64]. These delays may depend on the current state. Thibault contributed to implement efficiently Delayed Differential Equations in deep learning frameworks (Ajax, PyTorch) and showed the advantage of DDEs over (Augmented) Neural DDEs and recurrent networks (LSTM) under certain circonstances, in particular in the case of constant delays, that can now be learned during training [38].

8.4.2 Graph Neural Networks for Numerical Simulations

Participants: Guillaume Charpiat, Michele Alessandro Bucci (Safran Tech), Marc Schoenauer, Matthieu Nastorg, Lionel Mathelin (DATAFLOT team, LISN), Thilbault Faney (IFPEN), Jean-Marc Gratien (IFPEN).

During the 2.5 years that he spent at TAU (2021-2023), Alessandro MIchele Bucci, now at Safran Tech, worked on several use cases of IFPEN, with the goal of accelerating some softwares that IFPEN uses daily. This IFPEN/TAU collaboration led to a successful DataIA proposal, ML4CFD, that funded Matthieu Nastorg's PhD, defended in March 2024 [55]. After making significant improvements on B. Donon's Deep Statistical Solvers (DSS) [101], replacing the arbitrary number of iterations of the message passing mechanism by the solution of a fixed point equation [22], the final part of his PhD by considered the DSS approach as a preconditionner for a domain decomposition method [39].

Alessandro also contributed to propose a new approach to maintain the physical consistancy of GNN based approach to data assimilation for the RANS (Reynolds-Averaged Navier Stokes) equations, by hybridization with the classical adjoint metho [66].

Last but not least, we successfull applied to an Action Exploratoire (PI Guillaume Charpiat) *Large Physics Models* to investigate the use of Transformers and Large Foundational Models for Numerical simulations. Matthieu Nastorg has been hired on this AeX and he will co-supervise a new PhD student. This AeX is also tightly linked with the startup company AUGUR, founded by Emmanuel Menier, Matthieu Nastorg and Alice Lacan, all former PhD students in the team.

8.4.3 Advances in sparse recovery for inverse problem and application in M/EEG

Participants: Matthieu Kowalski, Jean-Baptiste Malagnoux, Diego Delle Donne (Essec), Leo Liberti (LiX), Benoît Malézieux (Inria Mind), Pierre Barbault (CentraleSupelec), Thomas Moreau (Inria Mind), Charles Soussen (CentraleSupelec).

Inverse problems involve reconstructing underlying signals or images from indirect or incomplete measurements and often require additional constraints or regularization to ensure unique and stable solutions. Sparse coding addresses these challenges by representing signals with a small number of nonzero coefficients in a suitable basis or dictionary. Methods like Convolutional Dictionary Learning build on this principle and have been successfully applied in areas such as neuroimaging and audio signal analysis.

In [35], we have established a theoretical equivalence between Convolutional Dictionary Learning (CDL) and Non-Negative Matrix Factorization (NMF) methods for signal processing in the time-frequency domain. We show that signals represented using CDL, which relies on sparse coding, can also be synthesized using factorized time-frequency coefficients in semi-NMF or complex-NMF forms. This connection bridges two widely-used approaches, highlighting their potential for joint optimization in applications like music transcription and biomedical signal analysis.

In [15] we provide groundbreaking perspective on the L_2+L_0 sparse approximation problems, introducing a comprehensive, big-M independent integer linear programming formulation. This development effectively circumvents the shortcomings of existing methods, ensuring the accurate recovery of global minimizers without requiring predefined bounds.

[58] introduces LEMUR, an Expectation-Maximization (EM) framework for estimating Bernoulli-Gaussian model parameters in inverse problems. By jointly estimating both the signal and hyperparameters, LEMUR reduces the need for manual tuning and proves effective for structured inverse problems, particularly with correlated measurement operators. Complementing this practical approach, [57] provides a theoretical study on Maximum A Posteriori (MAP) estimation for sparse recovery using the Bernoulli-Gaussian model. It establishes connections between joint-MAP and marginal-MAP estimators and classical methods like the Lasso and Sparse Bayesian Learning, demonstrating how specific relaxations of the MAP problem can efficiently recover the support of sparse signals.

[63] explores the feasibility and limitations of prior learning in unsupervised inverse problems, particularly focusing on dictionary learning and structured priors. The study shows that recovering a dictionary from incomplete data is only possible when multiple measurement operators span the entire signal space or when weak priors, such as equivariance or group structures, are incorporated. It highlights that handcrafted priors can be outperformed by learned priors, but only when they are well-adapted to the specific inverse problem. Complementing this theoretical perspective, [62] introduces a practical approach to prior learning, leveraging structured optimization techniques to refine sparse signal representations. This work provides experimental validation of learned priors in various inverse problem settings, demonstrating their potential for improving reconstruction quality compared to traditional handcrafted priors.

8.4.4 Instrument recognition in MIR

Participants: Matthieu Kowalski, Dylan Sechet, Francesca Bugiotti (LISN), Edouard d'Hérouville (Linkaband), Filip Langiewicz (Linkaband).

Detecting musical instruments in audio recordings is a complex challenge in Music Information Retrieval (MIR), particularly in polyphonic settings where multiple instruments play simultaneously. While existing methods perform well for common instruments with abundant training data, they often fail to detect rare or underrepresented instruments, such as those found in orchestral, non-Western, or niche musical contexts. This issue is further compounded by the lack of labeled datasets for rare instruments, limiting the ability of models to generalize across diverse instruments and styles.

In [43] we have proposed a hierarchical deep learning approach to address the challenge of detecting minority instruments in polyphonic music recordings. By leveraging the MedleyDB dataset, the study introduces a hierarchical classification framework that integrates group-level and instrument-level predictions to enhance the detection of rare instruments. The proposed two-pass classification approach demonstrates improved performance for underrepresented instruments.

8.4.5 Simulations for evolutionary genomics

Participants: Guillaume Charpiat, Flora Jay, Léo Planche, Arnaud Quelin.

Collaboration: Bioinfo Team (LISN), MNHN (Paris), UNAM (Mexico), U Brown (USA), METU (Turkey).

In population genetics, simulators are a valuable resource, enabling the testing of tool robustness, comparing the outcomes of stochastic evolutionary models with real observations, and performing simulation-based inference. In the latter case, simulations allow us to work in a supervised setting to solve the inverse problem of inferring the simulator's evolutionary parameter inputs from genomic outputs. We previously demonstrated how machine learning and deep learning could contribute to this task. More recently, we have focused on using simulations to test and infer individuals' relatedness and population structure from ancient DNA [18, 118], two key questions frequently asked in paleogenomics, such as in the study of Neolithic population dynamics in Anatolia [61] (collaboration with METU, Turkey).

8.5 Energy Management

Participants: Isabelle Guyon, Alessandro Leite, Marc Schoenauer, Eva Boguslawski, Benjamin Donnot (RTE), Matthieu Dussartre (RTE).

Our collaboration with RTE has a long history, starting with Benjamin Donnot's (2016-2019) [98] and Balthazar Donon's [100] CIFRE PhDs, and is centered on the maintainance of the national French Power Grid. Eva Boguslawski's CIFRE PhD, co-supervised by Alessandro Leite and Marc Schoenauer, started in Sept. 2022, and will be defended in 2025. It addresses the control of the grid through decentralized decision process using multi-agent Reinforcement Learning, in the line of the LR2PN challenge that Eva contributed to organize during her Master internship [125]. During the second year of her PhD, she focused on the emulation of Zonal Controllers for the Power System Transport Problem [31].

8.6 Computational Social Sciences

Participants: Philippe Caillou, Michèle Sebag, Cyriaque Rousselot, Guillaume Bied, Armand Lacombe, Soal Nathan, Hande Gozukan.

8.6.1 Labor Studies

Participants: Philippe Caillou, Michèle Sebag, Guillaume Bied, Armand Lacombe, Solal Nathan, Hande Gozukan, Jean-Pierre Nadal (EHESS), Bruno Crépon (ENSAE).

Job markets The DATAIA project Vadore [88] (partners ENSAE and Pôle Emploi/France Travail) benefits from the sustained cooperation and from the wealth of data gathered by France Travail. The data management is regulated along a 3-partite convention (GENES-ENSAE, Univ Paris-Saclay, Pôle Emploi). Extensive efforts have been required to achieve the data pipelines required to enable learning recommendation models and exploiting them in a confidentiality preserving way (G. Bied's PhD, [50]). The acceptability of the algorithm for the job seekers has been investigated using large-scale (100,000 job seekers) in Feb. 2023 and June 2024.

An important criterion, besides the performance in terms of recall and the time-to-solution, regards the fairness of the recommendation model [87, 86]. A comprehensive study examining gender-related gap in several utilities (wages, types of contract, distance-to-job) has been conducted, comparing the gaps observed in actual hirings, in applications, and in recommendations. Interestingly, the gap in recommendations closely mimics that in actual hirings and in applications (if any, the recommendation algorithm tends to decrease the gap). Algorithmic fairness in domains as sensitive as employment is under scrutiny of French and European regulations. The difficulty is to decouple the biases observed in applications (thus reflecting job seekers' preferences, that should be respected) from those due to recruiters (that should not be perpetuated in the learned models).

Another criterion concerns the congestion of the job market (share of job offers paid attention to by job seekers). Recommender systems tend to increase the congestion due to the so-called popularity bias. Early attempts to prevent the congestion have been investigated in [89], using optimal transport.

Both fairness and congestion issues are at the core of S. Nathan's PhD (coll. Univ. Ghent, Belgium). A first research direction is concerned with integrating the congestion in the recommendation loss; this requires a global view of the market dynamics, and the difficulty is to design a loss term that is both computationally affordable and differentiable. A second research direction along this line is to integrate an estimation of job offer popularity within the recommendation system, enabling job seekers to anticipate and react to, competition. Interestingly, such compound architecture (integrating the job offer popularity estimate and its effects on the decision of applying) could enable to model competition-avoidance strategies, in particular in relation with gender effects.

A key difficulty for research on ML-based job recommendation is the lack of open and representative datasets, owing to the very sensitive nature of the data and the protection of vulnerable persons. We collaborate with U. Gent, Belgium, welcoming Guillaume Bied for his post-doc and Solal Nathan for a PhD visit, on this topic.

8.6.2 Health and practices

Participants: Philippe Caillou, Michèle Sebag, Armand Lacombe, Cyriaque Rousselot, Olivier Allais (INRA), Julia Mink (Univ. Bonn, DE), Florian Yger (INSA Rouen).

Continuing our former partnership with INRAE (in the context of the *Initiative de Recherche Stratégique* Nutriperso; [81]), we proposed the HORAPEST DATAIA project to uncover the potential causal relationships between pesticide dissemination and children's health (Cyriaque Rousselot's PhD). Medical data is accessed using the Health data Hub after the CNIL approval and proprietary data from INRAE is used for detailed pesticide purchase on every french parcel. Contacts have been taken with the CHU Toulouse for cooperation on complementary data.

8.6.3 Scientific Information System and Visual Querying

Participants: Philippe Caillou, Michèle Sebag, Anne-Catherine Letournel, Hande Gozukan, Jean-Daniel Fekete (AVIZ, Inria Saclay).

A third area of activity concerns the 2D visualisation and querying of a corpus of documents. Its initial motivation was related to scientific organisms, institutes or Universities, using their scientific production (set of articles, authors, title, abstract) as corpus. The Cartolabe project (see also Section 7) started as an

Inria ADT (coll. TAO and AVIZ, 2015-2017). It received a grant from CNRS (coll. TAU, AVIZ and HCC-LRI, 2018-2019).

The originality of the approach is to rely on the content of the documents (as opposed to, e.g. the graph of co-authoring and citations). This specificity allowed to extend Cartolabe to various corpora, such as Wikipedia, Bibliotheque Nationale de France, or the Software Heritage. Cartolabe was also applied in 2019 to the *Grand Debat* dataset: to support the interactive exploration of the 3 million propositions; and to check the consistency of the official results of the *Grand Debat* with the data. Cartolabe has also been applied in 2020 to the COVID-19 kaggle publication dataset (Cartolabe-COVID project) to explore these publications.

Among its intended functionalities are: the visual assessment of a domain and its structuration (who is expert in a scientific domain, how related are the domains); the coverage of an institute expertise relatively to the general expertise; the evolution of domains along time (identification of rising topics). A round of interviews with beta-user scientists has been performed in 2019-2020. Cartolabe usage raises questions at the crossroad of human-centered computing, data visualization and machine learning: i) how to deal with stressed items (the 2D projection of the item similarities poorly reflects their similarities in the high dimensional document space; ii) how to customize the similarity and exploit the users' feedback about relevant neighborhoods. A statement of the current state of the project was published in 2021 [75].

8.7 Organization of Challenges

Participants: Isabelle Guyon, Marc Schoenauer , Anne-Catherine Letournel, Sylvain Chevallier, Sébastien Tréguer, Adrien Pavao, Romain Egele, Mouadh Yagoubi (IRT SystemX), Antoine Marot (RTE), Benjamin Donnot (RTE), Bruno Aristimunha.

The TAU group uses challenges (scientific competitions) as a means of stimulating research in machine learning and engage a diverse community of engineers, researchers, and students to learn and contribute advancing the state-of-the-art. The TAU group is community lead of the open-source **Codalab** platform (see Section 7), hosted by Université Paris-Saclay. The project had grown since 2019 and included until last year an engineer dedicated full time to administering the platform and developing challenges, Adrien Pavao. **Codabench**, the new version of Codalab, was financed in 2021 by a 500k€ project with the Région Ile-de-France. This project also received the support of the Chaire Nationale d'Intelligence Artificielle of Isabelle Guyon (2020-2024), Lawrence Berkeley Labs (2022-2025, Fair Universe project), and TAILOR ICT48 Network of Excellence (2020-2024).

TAILOR also allowed us to hire part time Sébastien Treguer, who worked on the co-organizing challenges linked to TAILOR scientific interests, Trustworthy AI (TAI) and combinations of Learning, Optimisation and Reasoning (LOR). Eight challenges were run with TAILOR banner, and TAILOR Deliverable 2.4 [69] reports the lessons learned from these challenges. Furthermore, some of these challenges were reported in dedicated publications:

- The Smarter Mobility Data Challenge deals with Forecasting Electric Vehicle Charging Station Occupancy [9];
- Meta Learning from Learning Curves [23] was concerned, as its title says, with the prediction of Deep Learning runs from very few iterations, after learning from a large database of various runs;
- The ML4CFD Competition: Harnessing Machine Learning for Computational Fluid Dynamics in Airfoil Design was accepted as a Neurips 2024 Challenge [48];

Beyond the "TAILOR Challenges" listed above, we continued the work on the AutoSurvey [109] series of challenge, sponsored by Google and ChaLearn https://auto-survey.chalearn.org/.

The goal of these challenges is to advance the generation of systematic review reports, overview papers, white papers, and essays that synthesize on-line information. The coverage spans multiple domains including Literary or philosophical essays (LETTERS), Scientific literature (SCIENCES), and topics surrounding the United Nations Sustainable Development Goals (SOCIAL SCIENCES). The participants submit code (AI-agents) capable of composing survey papers, using internet resources. Such AI-agents will thus operate as AI-Authors.

A further step was made with the AutoSurvey competition presented at the AutoML 2023 conference [49], in which participants were tasked to both presenting stand-alone models able to author articles from designated prompts, and subsequently review them. Assessment criteria include clarity, reference appropriateness, accountability, and the substantive value of the content.

It should be noted that the team is also participating to dataset design, prior to organizing competitions, as shown in [47], where an international team gathered to build a competitive dataset for predicting images directly from brain signals, that was presented during a dedicated workshop in CVPR.

Last but not least, beyond his PhD, Adrien also wrote a survey of the different types of competitions in Machine Learning, detailing the recommended protocol for each of them [70]. He also published a genera tutorial to guide future challenge organisers [65].

Also, we continue using challenges in teaching. The masters students of the AI master designed several small challenges, which are then given to other students in labs, and both types of students seem to love it. In 2023, they organized biomedicine, particle physics and computer vision challenges, on the theme of bias in data and fairness.

9 Bilateral contracts and grants with industry

Participants: Whole team.

9.1 Bilateral contracts with industry

TAU continues its policy about technology transfer, accepting any informal meeting following industrial requests for discussion (and we are happy to be often solicited), and deciding about the follow-up based upon the originality, feasibility and possible impacts of the foreseen research directions, provided they fit our general canvas. This lead to the following 4 on-going CIFRE PhDs, with the corresponding side-contracts with the industrial supervisor, and the continuation until September 2023 of the bilateral contract with Fujitsu (within the national "accord-cadre" Inria/Fujitsu).

- **CIFRE RTE** 2021-2024 (72 kEuros), with RTE, related to Eva Boguslawski's CIFRE PhD *Decentralized Partially Observable Markov Decision Process for Power Grid Management*
  Coordinator: Marc Schoenauer and Matthieu Dussartre (RTE)
  Participants: Eva Boguslawski, Alessandro Leite

- **CIFRE Ekimetrics** 2022-2025 (45 kEuros), with Ekimetrics, related to Audrey Poinsot's CIFRE PhD *Causal incertainty quantification under partial knowledge and low data regimes*
  Coordinator: Marc Schoenauer and Nicolas Chesneau (Ekimetrics)
  Participants: Guillaume Charpiat, Alessandro Leite, Audrey Poinsot and Michèle Sebag

- **CIFRE MAIR** 2022-2025 (75 kEuros), with Meta (Facebook) AI Research, related to Mathurin Videau's CIFRE PhD *Reinforcement Learning: Sparse Noisy Reward*
  Coordinator: Marc Schoenauer and Olivier Teytaud (Meta)
  Participants: Alessandro Leite and Mathurin Videau

- **CIFRE MAIR** 2022-2025 (75 kEuros), with Meta (Facebook) AI Research, related to Badr Youbi's CIFRE PhD *Learning invariant representations from temporal data*
  Coordinator: Michèle Sebag and David Lopez-Paz (Meta)
  Participants: Badr Youbi

- **AI Verse**, related to Abir Affane's post-doc
  Coordinator: Pierer Alliez (INRIA Titane)
  Participant: Guillaume Charpiat

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Visits to international teams

Research stays abroad

- Audrey Poinsot was a Computer Science Visiting Student Intern at Columbia University from June 3 through August 30, 2024. Under the mentorship of Professor Elias Bareinboim, Audrey worked on novel uncertainty measures for identified causal graphs, considering their underlying assumptions.

10.2 European initiatives

10.2.1 Horizon Europe

Adra-e Adra-e project on cordis.europa.eu

Title: AI, Data and Robotics ecosystem

Duration: From July 1, 2022 to June 30, 2025

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- LINKOPINGS UNIVERSITET (LIU), Sweden
- UNIVERSITY OF GALWAY (OLLSCOIL NA GAILLIMHE), Ireland
- DUBLIN CITY UNIVERSITY (DCU), Ireland
- AI DATA AND ROBOTICS ASSOCIATION, Belgium
- TRUST-IT SERVICES SRL, Italy
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (CEA), France
- UNIVERSITEIT TWENTE (UNIVERSITEIT TWENTE), Netherlands
- DEUTSCHES FORSCHUNGSZENTRUM FUR KUNSTLICHE INTELLIGENZ GMBH (DFKI), Germany
- ATOS SPAIN SA, Spain
- HRVATSKA UDRUGA ZA UMJETNU INTELIGENCIJU (CROATIAN ARTIFICIAL INTELLIGENCE ASSOCIATION), Croatia
- COMMPLA SRL (Commpla Srl), Italy
- ATOS IT SOLUTIONS AND SERVICES IBERIA SL (ATOS IT), Spain
- SIEMENS AKTIENGESELLSCHAFT, Germany
- UNIVERSITEIT VAN AMSTERDAM (UvA), Netherlands

Inria contact: Marc Schoenauer

Coordinator: Marc Schoenauer (Inria)

Summary: AI, Data and Robotics (ADR) is omnipresent in our daily lives and key to addressing some of the most pressing challenges facing our society. Europe has excellent research centres, innovative start-ups, a world-leading position in robotics and competitive manufacturing and services sectors, from automotive to healthcare, energy, financial services and agriculture. While the essentials are present, European ADR is waiting for exploitation to achieve its full potential. The ADR ecosystem is inherently complex because many stakeholders at many different levels require a holistic strategy towards collaboration to be effective and efficient. The Adra Association, representing the private side of the ADR Partnership, leverages this diversity through its founding organisations (BDVA, euRobotics, CLAIRE, ELLIS, EurAI) and channels it to the benefit of the European ecosystem. The Adra-e CSA proposal is set up in close liaison with Adra Association and includes it as a partner, committed to sustain its outcomes. Adra-e should be seen as the operational arm of the partnership to foster collaboration, convergence and interoperability between communities and disciplines to advance European ADR while safeguarding the interest of European citizens. This is achieved by supporting the ADR Partnership in the update and implementation of the SRIDA, creating the conditions for an inclusive, sustainable, effective, multi-layered, and coherent European ADR ecosystem, leading to increased trust and adoption of ADR, a more competitive supply and demand sides in the EU and raising private investments at the same time.The consortium is composed of leading industry and research organisations with significant expertise in all three disciplines. All are involved in Adra and the associations and partnerships shaping European research. Many of them are supporting the Digitising European Industry initiative from the EC participating in the constitution of Digital Innovation Hubs Network and Digital platforms.

MANOLO MANOLO project on cordis.europa.eu

Title: Trustworthy Efficient AI for Cloud-Edge Computing

Duration: From January 1, 2024 to December 31, 2026

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- PAL ROBOTICS SLU (PAL ROBOTICS), Spain
- LAUREA-AMMATTIKORKEAKOULU OY (LAUREA UNIVERSITY OF APPLIED SCIENCES), Finland
- UNIVERSITY COLLEGE DUBLIN, NATIONAL UNIVERSITY OF IRELAND, DUBLIN (NUID UCD), Ireland
- Q-PLAN INTERNATIONAL ADVISORS PC (Q-PLAN INTERNATIONAL), Greece
- FRAUNHOFER GESELLSCHAFT ZUR FORDERUNG DER ANGEWANDTEN FORSCHUNG EV (Fraunhofer), Germany
- BIT & BRAIN TECHNOLOGIES SL (BIT&BRAIN TECHNOLOGIES), Spain
- YRKESHOGSKOLAN ARCADA AB (ARCADA UNIVERSITY OF APPLIED SCIENCES LTD), Finland
- TECHNISCHE UNIVERSITAET BRAUNSCHWEIG, Germany
- EVIDEN TECHNOLOGIES SRL, Romania
- FOUR DOT INFINITY INFORMATION AND TELECOMMUNICATIONS SOLUTIONS PRIVATE COMPANY (FOUR DOT INFINITY LYSEIS PLIROFORIKIS KAI EPIKOINONION IDIOTIKI KEFALAIOUCHIKI ETAIREIA), Greece
- UNIVERSITAT POLITECNICA DE CATALUNYA (UPC), Spain
- "NATIONAL CENTER FOR SCIENTIFIC RESEARCH ""DEMOKRITOS""" ("NCSR ""D"""), Greece
- UNIVERSITE PARIS-SACLAY, France
- ATOS IT SOLUTIONS AND SERVICES IBERIA SL (ATOS IT), Spain
- KATHOLIEKE UNIVERSITEIT LEUVEN (KU Leuven), Belgium
- EIT DIGITAL, Belgium
- ARX NET AE YPIRESIES KAI EPICHIRISIS DIADIKTYOU ANONIMI ETAIRIA (ARX.NET S.A.), Greece

Inria contact: Guillaume Charpiat

Coordinator: Ricardo Simon Carbajo (Dublin University)

Summary: MANOLO will deliver a complete stack of trustworthy algorithms and tools to help AI systems reach better efficiency and seamless optimization in their operations, resources and data required to train, deploy and run high-quality and lighter AI models in both centralised and cloud-edge distributed environments. It will push the state of the art in the development of a collection of complementary algorithms for training, understanding, compressing and optimising machine learning models by advancing research in the areas of: model compression, meta-learning (few-shot learning), domain adaptation, frugal neural network search and growth and neuromorphic models. Novel dynamic algorithms for data/energy efficient and policy-compliance allocation of AI tasks to assets and resources in the cloud-edge continuum will be designed, allowing for trustworthy widespread deployment.

To support these activities a data management framework for distributed tracking of assets and their provenance (data, models, algorithms) and a benchmark system to monitor, evaluate and compare new AI algorithms and model deployments will be developed. Trustworthiness evaluation mechanisms will be embedded at its core for explainability, robustness and security of models while using the Z-Inspection methodology for TrustworthyAI assesment, helping AI systems conform to the new AI Act regulation.

MANOLO will be deployed as a toolset and tested in lab environments via Use Cases with different distributed AI paradigms within cloud-edge continuum settings; it will be validated in verticals such as health, manufacturing, and telecommunications aligned with ADRA identified market opportunities, and with a granular set of embedded devices covering robotics, smartphones, IoT as well as using Neuromorphic chips. MANOLO will integrate with ongoing projects at EU level developing the next operating system for cloud-edge continuum, while promoting its sustainability via the AI-on-demand platform and EU portals.

10.2.2 H2020 projects

TRUST-AI

Participants: Marc Schoenauer, Alessandro Leite.

TRUST-AI project on cordis.europa.eu

Title: Transparent, Reliable and Unbiased Smart Tool for AI

Duration: From October 1, 2020 to March 31., 2025

Partners:

- INESC TEC - INSTITUTO DE ENGENHARIADE SISTEMAS E COMPUTADORES, TECNOLOGIA E CIENCIA, Portugal (coordinator)
- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- TARTU ULIKOOL, Estonia
- STICHTING NEDERLANDSE WETENSCHAPPELIJK ONDERZOEK INSTITUTEN, Netherlands
- APPLIED INDUSTRIAL TECHNOLOGIES (APINTECH), Cyprus
- LTPLABS, SA, Portugal
- TAZI BILISIM TEKNOLOJILERI ANONIM SIRKETI, Türkiye

Inria contact: Marc Schoenauer

Coordinator: Gonçalo Figueira (INESC)

Summary: Due to their black-box nature, existing artificial intelligence (AI) models are difficult to interpret, and hence trust. Practical, real-world solutions to this issue cannot come only from the computer science world. The EU-funded TRUST-AI project is involving human intelligence in the discovery process. It will employ 'explainable-by-design' symbolic models and learning algorithms and adopt a human-centric, 'guided empirical' learning process that integrates cognition. The project will design TRUST, a trustworthy and collaborative AI platform, ensure its adequacy to tackle predictive and prescriptive problems and create an innovation ecosystem in which academics and companies can work independently or together.

TAILOR

Participants: Marc Schoenauer, Sébastien Treguer.

TAILOR project on cordis.europa.eu

Title: Foundations of Trustworthy AI - Integrating Learning, Optimization, and Reasoning.

Duration: From September 1, 2020 to August 31, 2024

Partners:

- LINKOPINGS UNIVERSITET, Sweden (coordinator)
- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- and 56 other partners in Europe

Inria contact: Marc Schoenauer

Coordinator: Fredrik Heinz (Linköping U.)

Summary: Maximising opportunities and minimising risks associated with artificial intelligence (AI) requires a focus on human-centred trustworthy AI. This can be achieved by collaborations between research excellence centres with a technical focus on combining expertise in theareas of learning, optimisation and reasoning. Currently, this work is carried out by an isolated scientific community where research groups are working individually or in smaller networks. The EU-funded TAILOR project aims to bring these groups together in a single scientific network on the Foundations of Trustworthy AI, thereby reducing the fragmentation and increasing the joint AI research capacity of Europe, helping it to take the lead and advance the state-of-the-art in trustworthy AI. The four main instruments are a strategic roadmap, a basic research programme to address grand challenges, a connectivity fund for active dissemination, and network collaboration activities.

VISION VISION project on cordis.europa.eu

Title: Value and Impact through Synergy, Interaction and coOperation of Networks of AI Excellence Centres

Duration: From September 1, 2020 to August 31, 2024

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- UNIVERSITEIT LEIDEN (ULEI), Netherlands
- NEDERLANDSE ORGANISATIE VOOR TOEGEPAST NATUURWETENSCHAPPELIJK ONDERZOEK TNO (NETHERLANDS ORGANISATION FORAPPLIED SCIENTIFIC RESEARCH), Netherlands
- THALES SIX GTS FRANCE SAS (THALES SIX GTS France), France

- DEUTSCHES FORSCHUNGSZENTRUM FUR KUNSTLICHE INTELLIGENZ GMBH (DFKI), Germany
- CESKE VYSOKE UCENI TECHNICKE V PRAZE (CVUT), Czechia
- FONDAZIONE BRUNO KESSLER (FBK), Italy
- INTELLERA CONSULTING SPA (INTELLERA CONSULTING), Italy
- UNIVERSITY COLLEGE CORK - NATIONAL UNIVERSITY OF IRELAND, CORK (UCC), Ireland

Inria contact: Jozef Geurts

Coordinator: Holger Hoos (ULEI)

Summary: A broad and ambitious vision is needed for artificial intelligence (AI) research and innovation in Europe to thrive and remain internationally competitive. Building upon its strength across allareas of AI and commitment to its core values, Europe is excellently positioned to take a human-centred, ethical and trustworthy approach to AI. However, to establish itself as a powerhouse in AI, Europe needs to overcome the present fragmentation of the AI community, which creates inefficiencies and limits progress on next-generation, trustworthy AI tools and systems that draw from methods across a broader spectrum of AI techniques. Europe also needs to improve interaction between AI researchers, innovators and users. Following on from the European Commission's Communication on AI for Europe and the Coordinated Action Plan between the European Commission and the Member States, European efforts in AI should be strongly coordinated to be internationally competitive. Europe must scale up existing research capacities and reach critical mass through tighter networks of European AI excellence centres. Towards this end, effective coordination between the four networks of AI excellence centres to be established under ICT-48-2020 (Research and Innovation Actions) is of crucial importance. VISION will reinforce and build on Europe's assets in AI, including its world-class community of researchers, and thus enable Europe to stay at the forefront of AI developments, which is widely recognised as critical in maintainingEurope's strategic autonomy in AI. VISION will achieve this in the most efficient and effective manner possible, by strongly building on the success and organisation of CLAIRE (the Confederation of Laboratories for AI Research inEurope, claire-ai.org) as well as on AI4EU, and by leveraging the expertise and connections of several of Europe's leading institutions in AI research and innovation.

10.3 National initiatives

10.3.1 ANR

- Chaire IA **HUMANIA** 2020-2024 (600kEuros), *Democratizing Artificial Intelligence.*
  Coordinator: Isabelle Guyon (TAU)
  Participants: Marc Schoenauer, Michèle Sebag, Anne-Catherine Letournel, François Landes.

- PEPR IA **SAIF** (400k€) *Safe AI through Formal methods*
  Coordinator: Caterina Urban (INRIA Antique)
  Participant: Guillaume Charpiat.

- PEPR IA **CAUSALI-T-AI** (400k€) *CAUSALIty Teams up with Artificial Intelligence*
  Coordinator: Marianne Clausel (Université de Loraine)
  Participant: Michèle Sebag, Alessandro Leite.

- **RoDAPoG** 2021-2025 (302k€) *Robust Deep learning for Artificial genomics and Population Genetics*
  Coordinator: Flora Jay,
  Participants: Cyril Furtlehner, Guillaume Charpiat.

- **SPEED** 2021-2024 (49k€) *Simulating Physical PDEs Efficiently with Deep Learning*
  Coordinator: Lionel Mathelin (LISN (ex-LIMSI))
  Participants: Michele Alessandro Bucci, Guillaume Charpiat, Marc Schoenauer.

10.3.2 Others

- **Inria Challenge OceanAI** 2021-2024, *AI, Data, Models for a Blue Economy*
  Coordinator: Nayat Sanchez Pi (Inria Chile)
  Participants: Marc Schoenauer, Michèle Sebag and Shiyang

- **DATAIA YARN** 2022-2025, *Automatic Processing of Messy Brain Data with Robust Methods and Transfer Learning*
  Coordinator: Sylvain Chevallier
  Participants: Florent Bouchard (L2S), Fredéric Pascal (L2S), Alexandre Gramfort (Meta), Sara Sedlar

- **Fair Universe** 2022-2025, We received with Lawrence Berkeley Labs a grant of 6.4 million USD to develop benchmarks in High Energy Physics and implement them on Codabench. Colaboration with David Rousseau of IJCLAB.
  Coordinator: Isabelle Guyon
  Participants: David Rousseau, Ragansu Chakkappai, Ihsan Ullah

- **Action Exploratoire** 2024-2026, *Large Physics Models*
  Coordinator: Guillaume Charpiat
  Participants: Mac Schoenauer, Matthieu Nastorg (post-doc), Theofanis Ifaistos (PhD student)

11 Dissemination

Participants: Whole team.

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair

- Flora Jay - Organizer of the first edition of the international conference LEGEND Machine Learning for Evolutionary Genomics, Greece 13-15/05/24

- Flora Jay - Organizer of GT LEGO Machine Learning for genomics Research Day, Paris 9/12/24

11.1.2 Scientific events: selection

Reviewer All TAU members are reviewers of the main conferences in their respective fields of expertise.

11.1.3 Journal

Member of the editorial boards

- Marc Schoenauer - Action editor, Journal of Machine Learning Research (JMLR); Advisory Board, Evolutionary Computation Journal, MIT Press, and Genetic Programming and Evolutionary Machines, Springer Verlag.

- Michèle Sebag - Editorial Board, ACM Transactions on Evolutionary Learning and Optimization.

Reviewer All members of the team reviewed numerous articles for the most prestigious journals in their respective fields of expertise.

11.1.4 Invited talks

- Guillaume Charpiat *An example of AI4Science: ML4CFD* at InPEx workshop (NumPEx), Sitges (Spain) June 18th 2024
- Guillaume Charpiat *ML4CFD: Deep learning for numerical simulations* at Dassault Systèmes & Inria Scientific Day, November 29th 2024
- Cyril Furtlehner *Online feature learning in terms of spectral flow processes* workshop "Complex systems, statistical mechanics and machine learning crossover (in memory of Giovanni Paladin)" in Les Houches (March 24-29 2024)
- Cyril Furtlehner *Bypassing first order phase transition for RBM training* workshop "From Machine-Learning Theory to Driven Complex Systems and back" Lausanne May 22-24 2024
- François Landes *Learning representations of glassy liquids with roto-translation equivariant Graph Neural Networks* workshop "Complex systems, statistical mechanics and machine learning crossover (in memory of Giovanni Paladin)" in Les Houches (March 24-29 2024)
- Marc Schoenauer, *Intelligence artificielle générale : utopie ou dystopie ?*, séminaire de philosophie "Les formes de l'intelligence" de l'université Paris Est Créteil;
- Sylvain Chevallier, *Riemannian geometry applied to time series*, UMR MIA, AgroParisTech, lab seminar, 19/12/2024.
- Sylvain Chevallier *Open science tools for reproducibility*, Open science week, Paris-Saclay, 5/11/2024
- Michèle Sebag: *Machine Learning and Interventions*, Harvest Alliance, AgroParisTech, 3/2/2025
- Michèle Sebag: *Job Recommendation for all: Challenges, Biases and Next*, From physics to neurosciences and social sciences, Conférence en l'honneur de J. P. Nadal, Ulm, 29/9/2024
- Michèle Sebag: *Causal Discovery : A Divide and Conquer approach using Integer Linear Programming*, Data conference, Isite; https://cap2025.fr/developpement-instrumental), 20/6/2024
- Michèle Sebag, *Causal models are generative models. What is special about them ?*, séminaire Collège de France, cours S. Mallat, 14/2/2024
- Michèle Sebag,*IA: 10 minutes* congrès de la SIF, 6-7/6/24.

11.1.5 Leadership within the scientific community

- Sylvain Chevallier: President of the academic society CORTICO, promoting the research in brain-computer interface; Executive Committee, Institut de Convergence DataIA; Research Committee, IA Cluster Paris-Saclay; Head of MSCA-Horizon Europe Cofund DeMythif.AI
- Flora Jay: member of GDR BiM science board (23-now)
- Marc Schoenauer: Advisory Board, ACM-SIGEVO, Special Interest Group on Evolutionary Computation; Chair of Advisory Board (Founding President 2015-2022), SPECIES, Society for the Promotion of Evolutionary Computation In Europe and Surroundings, that organizes the yearly series of conferences EvoStar. Senior Reviewer, IJCAI.
- Michèle Sebag: Member of scientific council of the AMIES Labex; Area Chair NeurIPS, ECML-PKDD; Senior Meta-Reviewer ECAI, nommée membre d'honneur de la SIF.

11.1.6 Scientific expertise

- Guillaume Charpiat : MdC hiring committees at DI ENS and at LIPN, USPN (Université Sorbonne Paris Nord, Villetaneuse)
- Guillaume Charpiat: Jean Zay (GENCI/IDRIS) committee member for resource allocation (GPU) demand expertise
- Guillaume Charpiat: expertise for post-doctoral grant allocation at MIAI (Institut Interdisciplinaire en Intelligence Artificielle de Grenoble)
- Sylvain Chevallier, "Conseil Scientifique", Inclusive Brain
- Sylvain Chevallier, PR hiring committees: LS2N, Nantes 19/04/2024; IUT Orsay, 26/04/2024
- Cyril Furtlehner, Inria representative at the Dataia COMP
- Marc Schoenauer, Scientific Advisory Board, BCAM, Bilbao, Spain
- Marc Schoenauer, "Conseil Scientifique", IFPEN
- Marc Schoenauer, "Conseil Scientifique", Mines Paritech
- Marc Schoenauer, "Conseil Scientifique", ADEME
- Marc Schoenauer, scientific coordinator of the IRT SystemX IA2 program (Artificial Intelligence for Augmented Engineering)
- Michele Sebag, "Conseil scientifique", IRSN
- Michele Sebag, FNRS (PhDs and Post-docs)
- Michele Sebag, hiring committee, Univ. Grenoble

11.1.7 Research administration

- Guillaume Charpiat: head of the Data Science department at LISN, Université Paris-Saclay.
- Michele Sebag, Member of Council: Institut Pascal, IRSN, ISC-PIF.
- Sylvain Chevallier, co-chair of the Scientific Council of Computer Science dept. from Universite Paris-Saclay; elected member of executive committee of University Institute

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Licence : Philippe Caillou, Computer Science for students in Accounting and Management, 192h, L1, IUT Sceaux, Univ. Paris Sud.
- Licence : François Landes, Introduction to Statistical Learning, 51h, L3, Univ. Paris-Sud.
- Licence : Matthieu Kowalski, Signal Processing, L3, 25h, Univ. Paris-Saclay
- Master : Guillaume Charpiat, Deep Learning in Practice, 24h, M2 Recherche, MVA / Centrale-Supelec / DSBA / MscIA
- Master : Guillaume Charpiat, Information Theory, 14h, M1 IA Paris-Sud.
- Master: Sylvain Chevallier, Machine learning algorithms, 12h, M1, Univ. Paris-Saclay.
- Master : Isabelle Guyon, Project: Creation of mini-challenges, M2, Univ. Paris-Sud.
- Master : Flora Jay, Population genetics inference, 4h, M2, U PSaclay.

- Master: Matthieu Kowalski, Signal Processing, 25h, M2, Univ. Paris-Saclay
- Master: Matthieu Kowalski, Sparse Coding, 36h, M2, Univ. Paris-Saclay
- Master : François Landes, Foundational Principles of Machine Learning, 25h, M1 Recherche (AI track), U. Paris-Sud.
- Master : François Landes, Machine Learning, 42h, M2 Recherche, Univ. Paris-sud, physics department (PCS international Master)
- Master : Michèle Sebag, Deep Learning, 4h; Reinforcement Learning, 12h; M2 Recherche, U. Paris-Sud.
- Master : Beatriz Seoane, Applied Statistics, 25h, M1 Recherche (AI track), U. Paris-Saclay.
- Tutorial for PhDs (and others) : Guillaume Charpiat, Deep Learning for Physics, 3h
- Continuing education (ie teaching in companies): Guillaume Charpiat, Machine Learning and Deep Learning, 6 days.

11.2.2 Supervision

- PhD Emmanuel MENIER, *Deep Learning for Reduced Order Modeling*, from 1/9/2020, Michele Alessandro Bucci, Marc Schoenauer, and Mouadh Yagoubi (IT SystemX), defended 25/01/2024.
- PhD Armand LACOMBE, *Changes of representation for counter-factual inference*, Michele Sebag and Philippe Caillou, defended 05/03/2024.
- PhD Romain EGELE, *Optimization of Learning Workflows at Large Scale on High-Performance Computing Systems*, Isabelle Guyon/Michèle Sebag and Prasanna Balaprakash (Argonne), defended 17/06/2024.
- PhD Mathieu NASTORG, *Scalable GNN Strategies to Solve Poisson Pressure Problems in CFD Simulations*, Guillaume Charpiat, Marc Schoenauer and Michele Alessandro Bucci, defended 15/04/2024.
- PhD Guillaume BIED, *Concevoir et évaluer les algorithmes de recommandation pour le marché du travail*, 1/10/2019, Bruno Crepon (CREST-ENSAE) and Philippe Caillou, defended 10/07/2024.
- PhD Appoline MELLOT, *Machine learning and domain adaptation for enhancing the measure of brain health with MEG and EEG signals*, Alexandre Gramfort (Inria MIND) and Sylvain Chevallier, defended 08/11/2024.
- PhD Maria Sayu YAMAMOTO, *Addressing the Large Variability of EEG Data with Riemannian Geometry: Toward Designing Reliable Brain-Computer Interfaces* , Fabien Lotte (Inria Potioc) and Sylvain Chevallier, defended 02/12/2024.
- PhD Emmanuel GOUTIERRE, *Machine learning-based particle accelerator modeling*, Johanne Cohen (LISN/Galac) and Michèle Sebag, defended 19/12/2024
- PhD Francisco PEZZICOLI *Statistical Physics - Machine Learning Interplay: from Addressing Class Imbalance with Replica Theory to Predicting Dynamical Heterogeneities with SE(3)-equivariant Graph Neural Networks*, François Landes and Guillaume Charpiat, defended 19/12/2024.
- PhD Thibault MONSEL, *Active Deep Learning for Complex Physical Systems*, Alexandre Allauzen (LAMSADE), Guillaume Charpiat, Lionel Mathelin (LISN), Onofrio Semeraro (LISN), defended 20/12/2024.
- PhD in progress - Anaclara ALVEZ *Scale-Equivariant Neural Networks* from 1/11/2023, François Landes and Cyril Furtlehner.

- PhD in progress - Bruno ARISTIMUNHA PINTO *deep learning for decoding electroencephalography* from 01/06/2023, Raphael Y de Camargo (UFABC Brazil), Marie-Constance Corsi (Inria Nerv), Sylvain Chevallier
- PhD in progress - Nicolas ATIENZA, *Towards Reliable ML: Leveraging Multi-Modal Representations, Information Bottleneck and Extreme Value Theory*, from 1/4/21, Michèle Sebag and Johanne Cohen.
- PhD in progress - Nicolas BÉREUX *interpretability and pattern extraction in Restricted Boltzmann Machines* from 1/11/2023, Beatriz Seoane Bartolome, Cyril Furtlehner.
- PhD in progress - Eva BOGUSLAWSKI *Congestion handling on Power Grid governed by complex automata*, from 1/05/22, Alessandro Leite, Mathieu Dussartre (RTE) and Marc Schoenauer
- PhD in progress - Thibault DE SURREL *Learning context invariant representations for EEG data*, from 1/11/2023, Florian Yger (ENSICaen), Fabien Lotte (Inria Potioc), Sylvain Chevallier
- PhD in progress - Styliani DOUKA *Growth strategies for neural architectures* from 01/01/2024, Guillaume Charpiat and Sylvain Chevallier
- PhD in progress - Badr Youbi IDRISSI *Learning an invariant representation through continuously evolving data*, from 01/10/22, David Lopez-Paz (Meta) and Michèle Sebag
- PhD in progress - Jean-Baptiste MALAGNOUX *Convolutional Dictionary Learning and time-frequency Nonmegative Matrix Factorization*, from 1/10/2022, Matthieu Kowalski
- PhD in progress - Florent MICHEL *Deep Learning for Dictionary Learning*, from 1/10/2022, Matthieu Kowalski and Thomas Moreau (Inria Mind)
- PhD in progress - Solal NATHAN, *Job recommendation, AI Ethics and Optimal Transport.*, 1/1/2023, Michèle Sebag and Philippe Caillou.
- PhD in progress - Audrey POINSOT, *Causal Uncertainty Quantification under Partial Knowledge and Low Data Regimes*, from 1/03/22, Nicolas Chesneau (Ekimetrics), Guillaume Charpiat, Alessandro Leite, and Marc Schoenauer
- PhD in progress - Arnaud QUELIN, *Infering Human population history with approximated Bayesian computation and machine learning, from ancient and recent genomes' polymorphism data*, from 1/10/22, Frédéric Austerlitz (MNHN), Flora Jay
- PhD in progress - Cyriaque ROUSSELOT, *Spatio-temporal causal discovery – Application to modeling pesticides impact*, from 1/10/22, Philippe Caillou
- PhD in progress - Théo RUDKIEWICZ*Growing neural networks for frugal learning* from 01/10/2024, Guillaume Charpiat and Sylvain Chevallier
- PhD in progress: Nilo SCHWENKE *Modélisation des batteries Lithium-Ion par Physics-Informed Neural Networks* from 1/09/2023, Cyril Furtlehner
- PhD in progress - Antoine SZATKOWNIK, *Deep learning for population genetics*, from 1/10/22, Flora Jay, Burak Yelmen, Cyril Furtlehner and Guillaume Charpiat
- PhD in progress - Manon VERBOCKHAVEN, *Strategies for Neural Architecture Growth*, from 11/2021, Sylvain Chevallier and Guillaume Charpiat
- PhD in progress - Sebastien VELUT *Understanding and addressing within-user variability in reactive and passive Brain-Computer Interfaces* since 13/11/2023, Frédéric Dehais (ISAE SupAero), Marie-Constance Corsi (Inria Nerv), Sylvain Chevallier
- PhD in progress, Mathurin VIDEAU, *Reinforcement Learning with sparse reward*, from 01/10/2021, Alessandro Leite, Marc Schoenauer and Olivier Teytaud (Meta).

11.2.3 Juries

- Flora Jay, PhD jury member: Guillaume Lan-Fong 13/12/24 , Aurélien Beaude 6/12/24, Letizia Lamperti 22/11/24, Luca Nesterenko 18/11/24, Margaux Lefebvre 16/09/24
- Marc Schoenauer, PhD jury member: Thimotée Anne, LARSEN team at Inria Nancy, 6/6/24; Claire Bizon Monroc, DYOGENE team, Inria Paris, 12/11/24; Michele Quattromini, LISN, 13/12/24; Abdelkader Dib, IFPEN, 14/1/25;
- Sylvain Chevallier, PhD reviewer: Mohammad Javad DARVISHI BAYAZI, MILA Montréal, 29/10/2024; Mathieu SERAPHIM, ENSICaen, 11/12/2024 Alexandre BLEUZE, GIPSA-lab, Grenoble, 8/12/2023, PhD jury president: Maxime TOQUEBIAU, ISIR, Paris, 16/12/2024; Ahmad CHAMMA, Inria Mind, 14/06/2024; Juan Jesús TORRE TRESOLS, ISAE SupAero, 21/10/2024, PhD jury member: Igor CARRARA, Inria Cronos, 12/10/2024; Fernando GONZALEZ, Coria, Rouen 29/03/2023; Armita KHAJEH NASSIRI, Team LaHDAK, LISN, 13/07/2023
- François Landes: head of the M1 and M2 AI track selection comittee (M1 and M2 combined: 1000+ applicants per year). Also head of the scholarship short-listing comittee.
- Matthieu Kowalski, PhD reviewer
- Cyril Furtlehner, PhD reviewer: Maciej KARCZ, CEA Cadarache 1/10/2024
- Michele Sebag, PhD reviewer: F. Jourdan, U. Toulouse; Simo Alami (X); Yoosof Mashayekhi, U. Gent, Belgium; Lei Zan, U. Grenoble; Colin Troisemaine, IMT Altlantique.
- Guillaume Charpiat, PhD jury member: Alexandre Vérine (Lamsade Dauphine PSL) 01/07/2024, Daniel Zyss (École des Mines Paris PSL) 17/10/2024, Elouan Argouarc'h (TelecomSudParis/CEA) 11/12/2024

11.3 Popularization

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

- Alessandro Leite and Marc Schoenauer produced the Webinar *Memetic Semantic Genetic Programming*, within the TRUST-AI Horizon Europe project
- Michèle Sebag participated in *Ce qui échappe à l'intelligence artificielle*, eds. François Levin, Étienne Ollion. ISBN : 9791037038449
- Michèle Sebag: conf.de presse *IA et mésinformation* (with Nicolas Curien); report on *IA et mésinformation*, ed. N. Curien, Académie des Technologies, 13/12/24; report on *Prouesses et limites de l'imittion artificielle de langages*, ed G. Roucayrol, https://www.academie-technologies.fr/publications/prouesses-et-limites-de-limitation-artificielle-de-langages-avis/

12 Scientific production

12.1 Major publications

[1] N. Atienza, R. Bresson, C. Rousselot, P. Caillou, J. Cohen, C. Labreuche and M. Sebag. 'Cutting the Black Box: Conceptual Interpretation of a Deep Neural Net with Multi-Modal Embeddings and Multi-Criteria Decision Aid'. In: *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. IJCAI-24, Thirty-Third International Joint Conference on Artificial Intelligence. Jeju, South Korea: International Joint Conferences on Artificial Intelligence Organization, 2024, pp. 3669–3678. DOI: 10.24963/ijcai.2024/406. URL: https://hal.science/hal-04728875.

[2] N. Atienza, C. Labreuche, J. Cohen and M. Sebag. 'Provably Safeguarding a Classifier from OOD and Adversarial Samples: an Extreme Value Theory Approach'. In: *Proc. ICLR'25*. ICLR 2025 - The Thirteenth International Conference on Learning Representations. Singapore (SG), Singapore, 17th Jan. 2025. URL: https://inria.hal.science/hal-04922382.

[3] D. Bachtis, G. Biroli, A. Decelle and B. Seoane. 'Cascade of phase transitions in the training of Energy-based models'. In: NeurIPS 2024 - 38th Annual Conference on Neural Information Processing Systems. Vancouver, Canada, 10th Dec. 2024. DOI: 10.48550/arXiv.2405.14689. URL: https://hal.science/hal-04897693.

[4] N. Bereux, A. Decelle, C. Furtlehner, L. Rosset and B. Seoane. 'Fast training and sampling of Restricted Boltzmann Machines'. In: 13th International Conference on Learning Representations - ICLR 2025. Singaour, Malaysia, 15th Mar. 2025. URL: https://inria.hal.science/hal-04885777.

[5] S. Dong, M. Sebag, K. Uemura, A. Fujii, S. Chang, Y. Koyanagi and K. Maruhashi. 'DCDILP: a distributed learning method for large-scale causal structure learning'. In: *Proc. AAAI 2025*. AAAI 25 - The 39th Annual AAAI Conference on Artificial Intelligence. Philadelphia (PA), United States, 25th Feb. 2025. URL: https://hal.science/hal-04710846.

[6] A. Mellot, A. Collas, S. Chevallier, A. Gramfort and D. Engemann. 'Geodesic optimization for predictive shift adaptation on EEG data'. In: *Proc. NeuIPS'24*. NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver, Canada, 10th Dec. 2024. URL: https://hal.science/hal-04902523.

[7] N. Schwencke and C. Furtlehner. 'ANAGRAM: a natural gradient relative to adapted model for efficient PINNS learning'. In: *In proceeding of ICLR 2025*. ICLR 2025 - International Conference on Learning Representations. 13th International Conference on Learning Representations (ICLR 2025). Singapour, Malaysia, 24th Apr. 2025. URL: https://inria.hal.science/hal-04918272.

[8] M. Verbockhaven, T. Rudkiewicz, S. Chevallier and G. Charpiat. 'Growing Tiny Networks: Spotting Expressivity Bottlenecks and Fixing Them Optimally'. In: *Transactions on Machine Learning Research Journal* (28th Oct. 2024). URL: https://hal.science/hal-04591472.

12.2 Publications of the year

International journals

[9] Y. Amara-Ouali, Y. Goude, N. Doumèche, P. Veyret, A. Thomas, D. Hebenstreit, T. Wedenig, A. Satouf, A. Jan, Y. Deleuze, P. Berhaut, S. Treguer and T. Phe-Neau. 'Forecasting Electric Vehicle Charging Station Occupancy: Smarter Mobility Data Challenge'. In: *Journal of Data-centric Machine Learning Research* 1 (1st July 2024). URL: https://hal.science/hal-04119408 (cit. on p. 24).

[10] I. Carrara, B. Aristimunha, M.-C. Corsi, R. Y. de Camargo, S. Chevallier and T. Papadopoulo. 'Geometric Neural Network based on Phase Space for BCI decoding'. In: *Journal of Neural Engineering* (18th Oct. 2024). DOI: 10.1088/1741-2552/ad88a2. URL: https://inria.hal.science/hal-04500580 (cit. on p. 16).

[11] G. Catania, A. Decelle and B. Seoane. 'The Copycat Perceptron: Smashing Barriers Through Collective Learning'. In: *Physical Review E* 109.6 (2024), p. 065313. DOI: 10.1103/PhysRevE.109.065313. URL: https://hal.science/hal-04918895 (cit. on p. 19).

[12] R. Chacko, F. Landes, G. Biroli, O. Dauchot, A. Liu and D. Reichman. 'Dynamical Facilitation Governs the Equilibration Dynamics of Glasses'. In: *Physical Review X* 14.3 (19th July 2024), p. 031012. DOI: 10.1103/PhysRevX.14.031012. URL: https://hal.science/hal-04749267 (cit. on p. 19).

[13] M.-C. Corsi, P. Sorrentino, D. P. Schwartz, N. George, L. L. Gollo, S. Chevallier, L. Hugueville, A. E. Kahn, S. Dupont, D. S. Bassett, V. Jirsa and F. de Vico Fallani. 'Measuring Neuronal Avalanches to inform Brain-Computer Interfaces'. In: *iScience* 27.1 (2024), p. 108734. DOI: 10.1016/j.isci.2023.108734. URL: https://inria.hal.science/hal-04345847 (cit. on p. 17).

[14] A. Decelle, C. Furtlehner, A. d. J. Navas Gómez and B. Seoane. 'Inferring effective couplings with Restricted Boltzmann Machines'. In: *SciPost Physics*. Scipost Physics 16.4 (5th Apr. 2024), p. 095. DOI: 10.21468/SciPostPhys.16.4.095. URL: https://inria.hal.science/hal-04885720 (cit. on p. 19).

[15] D. Delle Donne, M. Kowalski and L. Liberti. 'A Novel Integer Linear Programming Approach for Global L0 Minimization'. In: *Journal of Machine Learning Research* 24.1 (6th Mar. 2024), Art. no. 382, p.18322–18349. DOI: 10.5555/3648699.3649081. URL: https://hal.science/hal-04372309 (cit. on p. 21).

[16] G. Jung, R. M. Alkemade, V. Bapst, D. Coslovich, L. Filion, F. P. Landes, A. Liu, F. S. Pezzicoli, H. Shiba, G. Volpe, F. Zamponi, L. Berthier and G. Biroli. 'Roadmap on machine learning glassy liquids'. In: *Nature Rev.Phys.* 7.2 (2025), pp. 91–104. DOI: 10.1038/s42254-024-00791-4. URL: https://hal.science/hal-04344224 (cit. on p. 19).

[17] B. Junqueira, B. Aristimunha, S. Chevallier and R. de Camargo. 'A systematic evaluation of Euclidean alignment with deep learning for EEG decoding'. In: *Journal of Neural Engineering* 21.3 (23rd May 2024), p. 036038. DOI: 10.1088/1741-2552/ad4f18. URL: https://hal.science/hal-04677699 (cit. on p. 16).

[18] M. Lefeuvre, M. D. Martin, F. Jay, M.-C. Marsolier and C. Bon. 'GRUPS-rs, a high-performance ancient DNA genetic relatedness estimation software relying on pedigree simulations'. In: *Human Population Genetics and Genomics* 4.1 (4th Jan. 2024), p. 0001. DOI: 10.47248/hpgg2404010001. URL: https://hal.science/hal-04709778 (cit. on p. 22).

[19] A. Leite and M. Schoenauer. 'Memetic Semantic Boosting for Symbolic Regression'. In: *Genetic Programming and Evolvable Machines* (2025). URL: https://hal.science/hal-04911540. In press (cit. on p. 17).

[20] C. Loot, G. A. Millot, E. Richard, E. Littner, C. Vit, F. Lemoine, B. Néron, J. Cury, B. Darracq, T. Niault, D. Lapaillerie, V. Parissi, E. P. C. Rocha and D. Mazel. 'Integron cassettes commonly integrate into bacterial genomes via widespread non-classical attG sites'. In: *Nature Microbiology* 9.1 (3rd Jan. 2024), pp. 228–240. DOI: 10.1038/s41564-023-01548-y. URL: https://pasteur.hal.science/pasteur-04384854.

[21] E. Menier, S. Kaltenbach, M. Yagoubi, M. Schoenauer and P. Koumoutsakos. 'Interpretable learning of effective dynamics for multiscale systems'. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 481.2306 (31st Jan. 2025), p. 20240167. DOI: 10.1098/rspa.2024.0167. URL: https://inria.hal.science/hal-04392797 (cit. on pp. 17, 20).

[22] M. Nastorg, M. A. Bucci, T. Faney, J.-M. Gratien, G. Charpiat and M. Schoenauer. 'An implicit GNN solver for Poisson-like problems'. In: *Computers & Mathematics with Applications* 176 (15th Dec. 2024), pp. 270–288. DOI: 10.1016/j.camwa.2024.10.036. URL: https://hal.science/hal-04919611 (cit. on p. 20).

[23] M. H. Nguyen, L. Sun-Hosoya and I. Guyon. 'Meta-Learning from Learning Curves for Budget-Limited Algorithm Selection'. In: *Pattern Recognition Letters* 185 (1st Sept. 2024), pp. 225–231. DOI: 10.1016/j.patrec.2024.08.010. URL: https://inria.hal.science/hal-04719035 (cit. on p. 24).

[24] F. S. Pezzicoli, G. Charpiat and F. P. Landes. 'Rotation-equivariant graph neural networks for learning glassy liquids representations'. In: *SciPost Physics* 16.5 (28th May 2024), p. 136. DOI: 10.21468/SciPostPhys.16.5.136. URL: https://inria.hal.science/hal-03868206 (cit. on p. 19).

[25] V. M. Schimmenti, G. Petrillo, A. Rosso and F. P. Landes. 'Assessing the predicting power of GPS data for aftershocks forecasting'. In: *Seismological Research Letters* 95.6 (Nov. 2024), pp. 3243–3249. DOI: 10.1785/0220240008. URL: https://hal.science/hal-04352596 (cit. on p. 16).

[26] M. Verbockhaven, T. Rudkiewicz, S. Chevallier and G. Charpiat. 'Growing Tiny Networks: Spotting Expressivity Bottlenecks and Fixing Them Optimally'. In: *Transactions on Machine Learning Research Journal* (28th Oct. 2024). URL: https://hal.science/hal-04591472 (cit. on pp. 10, 15).

International peer-reviewed conferences

[27] N. Atienza, R. Bresson, C. Rousselot, P. Caillou, J. Cohen, C. Labreuche and M. Sebag. 'Cutting the Black Box: Conceptual Interpretation of a Deep Neural Net with Multi-Modal Embeddings and Multi-Criteria Decision Aid'. In: *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. IJCAI-24, Thirty-Third International Joint Conference on Artificial Intelligence. Jeju, South Korea: International Joint Conferences on Artificial Intelligence Organization, 2024, pp. 3669–3678. DOI: 10.24963/ijcai.2024/406. URL: https://hal.science/hal-04728875 (cit. on pp. 10, 17).

[28] N. Atienza, C. Labreuche, J. Cohen and M. Sebag. 'Provably Safeguarding a Classifier from OOD and Adversarial Samples: an Extreme Value Theory Approach'. In: *Proc. ICLR'25*. ICLR 2025 - The Thirteenth International Conference on Learning Representations. Singapore (SG), Singapore, 17th Jan. 2025. URL: https://inria.hal.science/hal-04922382 (cit. on p. 10).

[29] D. Bachtis, G. Biroli, A. Decelle and B. Seoane. 'Cascade of phase transitions in the training of Energy-based models'. In: NeurIPS2024. Vancouver, United States: arXiv, 2024. DOI: 10.48550/arXiv.2405.14689. URL: https://hal.science/hal-04844953 (cit. on pp. 10, 18).

[30] N. Bereux, A. Decelle, C. Furtlehner, L. Rosset and B. Seoane. 'Fast training and sampling of Restricted Boltzmann Machines'. In: 13th International Conference on Learning Representations - ICLR 2025. Singaour, Singapore, 15th Mar. 2025. URL: https://inria.hal.science/hal-04885777 (cit. on pp. 10, 19).

[31] E. Boguslawski, A. Leite, B. Donnot, M. Dussartre and M. Schoenauer. 'Emulation of Zonal Controllers for the Power System Transport Problem'. In: ML4SPS 2024 - Machine Learning for Sustainable Power Systems ECML 2024 Workshop. Vilnius, Lithuania, 9th Sept. 2024. URL: https://hal.science/hal-04924271 (cit. on p. 22).

[32] S. Dong, M. Sebag, K. Uemura, A. Fujii, S. Chang, Y. Koyanagi and K. Maruhashi. 'DCDILP: a distributed learning method for large-scale causal structure learning'. In: *Proc. AAAI 2025*. AAAI 25 - The 39th Annual AAAI Conference on Artificial Intelligence. Philadelphia (PA), United States, 25th Feb. 2025. URL: https://hal.science/hal-04710846 (cit. on pp. 10, 16).

[33] S. Douka, M. Verbockhaven, T. Rudkiewicz, S. Rivaud, F. P. Landes, S. Chevallier and G. Charpiat. 'Growth strategies for arbitrary DAG neural architectures'. In: *Proc. ESANN'25*. ESANN 2025 - 33th European Symposium on Artificial Neural Networks. Bruges, Belgium, 23rd Apr. 2025. URL: https://hal.science/hal-04902059 (cit. on p. 15).

[34] A. Lacan, B. Hanczar and M. Sebag. 'Frugal Generative Modeling for Tabular Data'. In: *Lecture Notes in Computer Science*. ECML PKDD 2024 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases. Vol. 14948. Lecture Notes in Computer Science. Vilnius, Lithuania: Springer Nature Switzerland, 22nd Aug. 2024, pp. 55–72. DOI: 10.1007/978-3-031-70371-3_4. URL: https://univ-evry.hal.science/hal-04705131 (cit. on p. 15).

[35] J.-B. Malagnoux and M. Kowalski. 'From Convolutional Sparse Coding To *-NMF Factorization of Time-Frequency Coefficients'. In: *2024 IEEE International Conference on Acoustics, Speech and Signal Processing*. ICASSP 2024 - IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). Séoul, South Korea: IEEE, 2024, pp. 5530–5534. DOI: 10.1109/ICASSP48485.2024.10447466. URL: https://hal.science/hal-04608586 (cit. on p. 21).

[36] A. Mellot, A. Collas, S. Chevallier, D. Engemann and A. Gramfort. 'Physics-informed and Unsupervised Riemannian Domain Adaptation for Machine Learning on Heterogeneous EEG Datasets'. In: EUSIPCO 2024 - 32nd European Signal Processing Conference. Lyon, France: arXiv, 2024. DOI: 10.48550/arXiv.2403.15415. URL: https://hal.science/hal-04716030 (cit. on p. 16).

[37] A. Mellot, A. Collas, S. Chevallier, A. Gramfort and D. Engemann. 'Geodesic optimization for predictive shift adaptation on EEG data'. In: *Proc. NeuIPS'24*. NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver, Canada, 10th Dec. 2024. URL: https://hal.science/hal-04902523 (cit. on pp. 10, 16).

[38] T. Monsel, O. Semeraro, L. Mathelin and G. Charpiat. 'Time and State Dependent Neural Delay Differential Equations'. In: *PMLR*. ML-DE@ECAI 2024 : Machine Learning Meets Differential Equations: From Theory to Applications. Santiago de compostela, Galicia, Spain, Oct. 2024. URL: https://hal.science/hal-04794450 (cit. on p. 20).

[39] M. Nastorg, J.-M. Gratien, T. Faney, M. A. Bucci, G. Charpiat and M. Schoenauer. 'Multi-Level GNN Preconditioner for Solving Large Scale Problems'. In: 2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). San Francisco (CA, USA), United States: IEEE Computer Society, 27th May 2024, pp. 486–495. DOI: 10.1109/IPDPSW63119.2024.00101. URL: https://inria.hal.science/hal-04447099 (cit. on p. 20).

[40] A. Poinsot, A. Leite, N. Chesneau, M. Sebag and M. Schoenauer. 'Learning Structural Causal Models through Deep Generative Models: Methods, Guarantees, and Challenges'. In: *Proc. Thirty-Third International Joint Conference on Artificial Intelligence*. IJCAI 2024 - Survey Track. Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence. Jeju, South Korea, 3rd Aug. 2024, pp. 8207–8215. DOI: 10.24963/ijcai.2024/907. URL: https://inria.hal.science/hal-04706985 (cit. on pp. 10, 16).

[41] G. H. Rodrigues, B. Aristimunha, S. Chevallier and R. Y. de Camargo. 'Combining Euclidean Alignment and Data Augmentation for BCI decoding'. In: EUSIPCO 2024 - 32nd European Signal Processing Conference. Lyon, France, 2024. DOI: 10.48550/arXiv.2405.14994. URL: https://inria.hal.science/hal-04743791 (cit. on p. 16).

[42] N. Schwencke and C. Furtlehner. 'ANAGRAM: a natural gradient relative to adapted model for efficient PINNS learning'. In: *In proceeding of ICLR 2025*. ICLR 2025 - International Conference on Learning Representations. 13th International Conference on Learning Representations (ICLR 2025). Singapour, Singapore, 24th Apr. 2025. URL: https://inria.hal.science/hal-04918272 (cit. on pp. 10, 19).

[43] D. Sechet, F. Bugiotti, M. Kowalski, E. D'hérouville and F. Langiewicz. 'A Hierarchical Deep Learning Approach for Minority Instrument Detection'. In: DAFx 2024 - the 27th International Conference on Digital Audio Effects. Guildford Surrey, United Kingdom, 3rd Sept. 2024. URL: https://hal.science/hal-04793160 (cit. on p. 22).

[44] T. de Surrel, S. Chevallier, F. Lotte and F. Yger. 'Averaging trajectories on the manifold of symmetric positive definite matrices'. In: EUSIPCO 2024 - 32nd European Signal Processing Conference. Lyon, France, 26th Aug. 2024. URL: https://hal.science/hal-04708878 (cit. on p. 16).

[45] A. Szatkownik, C. Furtlehner, G. Charpiat, B. Yelmen and F. Jay. 'Towards creating longer genetic sequences with GANs: Generation in principal component space'. In: *PMLR*. MLCB 2023 - 18th Conference on Machine Learning in Computational Biology. Seattle, United States, 2024. URL: https://hal.science/hal-04419057 (cit. on p. 18).

[46] S. Velut, S. Chevallier, M.-C. Corsi and F. Dehais. 'Deep Riemannian Neural Architectures for Domain Adaptation in Burst cVEP-based Brain Computer Interface'. In: ESANN 2024 - 32nd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges, Belgium: Ciaco - i6doc.com, 3rd Oct. 2024, pp. 571–576. DOI: 10.14428/esann/2024.ES2024-112. URL: https://hal.science/hal-04720928 (cit. on p. 16).

[47] J. Xu, B. Aristimunha, M. E. Feucht, E. Qian, C. Liu, T. Shahjahan, M. Spyra, S. Z. Zhang, N. Short, J. Kim, P. Perdomo, R. R. Mao, Y. Sabharwal, M. Ahedor, M. Shoura and A. Nestor. 'Alljoined1 -A dataset for EEG-to-Image decoding'. In: Workshop Data Curation and Augmentation in Medical Imaging at 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Settle, United States: arXiv, 14th May 2024. DOI: 10.48550/arXiv.2404.05553. URL: https://inria.hal.science/hal-04743819 (cit. on p. 25).

[48] M. Yagoubi, D. Danan, M. Leyli-Abadi, J.-P. Brunet, M. Gmati, A. Mazari, F. Bonnet, A. Farjallah, P. Cinnella, P. Gallinari, M. Schoenauer and A. Farjallah. 'NeurIPS 2024 ML4CFD Competition: Harnessing Machine Learning for Computational Fluid Dynamics in Airfoil Design'. In: NeurIPS 2024 - The 38th Annual Conference on Neural Information Processing Systems. Vancouver, Canada, 10th Dec. 2024. URL: https://hal.science/hal-04869170 (cit. on pp. 10, 24).

Conferences without proceedings

[49] P. H. Couto, Q. P. Ho, N. Kumari, B. K. Rachmat, T. G. H. Khuong, I. Ullah and L. Sun-Hosoya. 'RelevAI-Reviewer: A Benchmark on AI Reviewers for Survey Paper Relevance'. In: CAp 2024 - Conférence sur l'Apprentissage Automatique. Lille, France, 1st July 2024. URL: https://inria.hal.science/hal-04608255 (cit. on p. 25).

Doctoral dissertations and habilitation theses

[50] G. Bied. 'Designing Recommender Systems for the Labor Market'. Université Paris-Saclay, 10th July 2024. URL: https://theses.hal.science/tel-04740932 (cit. on p. 23).

[51] R. Egele. 'Optimization of Learning Workflows at Large Scale on High-Performance Computing Systems'. Université Paris-Saclay, 17th June 2024. URL: https://theses.hal.science/tel-04636586 (cit. on p. 17).

[52] A. Lacombe. 'Changes of representation for counter-factual inference'. Université Paris-Saclay, 5th Mar. 2024. URL: https://theses.hal.science/tel-04759706 (cit. on p. 16).

[53] A. Mellot. 'Machine learning and domain adaptation for enhancing the measure of brain health with MEG and EEG signals'. Université Paris-Saclay, 8th Nov. 2024. URL: https://theses.hal.science/tel-04906458.

[54] E. Menier. 'Deep Learning for Reduced Order Modeling'. Université Paris-Saclay, 25th Jan. 2024. URL: https://theses.hal.science/tel-04616516 (cit. on p. 20).

[55] M. Nastorg. 'Scalable GNN Solutions for CFD Simulations'. Université Paris-Saclay, 15th Apr. 2024. URL: https://theses.hal.science/tel-04590477 (cit. on p. 20).

[56] F. Pezzicoli. 'Statistical Physics - Machine Learning Interplay : from Addressing Class Imbalance with Replica Theory to Predicting Dynamical Heterogeneities with SE(3)-equivariant Graph Neural Networks'. Université Paris-Saclay, 19th Dec. 2024. URL: https://theses.hal.science/tel-04910839 (cit. on p. 19).

Reports & preprints

[57] P. Barbault, M. Kowalski and C. Soussen. *Bernoulli and Gauss Take a Look at the MAP*. 14th Nov. 2024. URL: https://hal.science/hal-04782425 (cit. on p. 21).

[58] P. Barbault, M. Kowalski and C. Soussen. *LEMUR: Latent EM Unsupervised Regression for Sparse Inverse Problems*. Apr. 2024. URL: https://hal.science/hal-04542061 (cit. on p. 21).

[59] S. Chevallier, I. Carrara, B. Aristimunha, P. Guetschel, S. Sedlar, B. Junqueira Lopes, S. Velut, S. Khazem and T. Moreau. *The largest EEG-based BCI reproducibility study for open science: the MOABB benchmark*. 8th Apr. 2024. URL: https://universite-paris-saclay.hal.science/hal-04537061.

[60] R. Egele, F. Mohr, T. Viering and P. Balaprakash. *The Unreasonable Effectiveness Of Early Discarding After One Epoch In Neural Network Hyperparameter Optimization*. 5th Apr. 2024. URL: https://hal.science/hal-04537565 (cit. on p. 17).

[61] D. Koptekin, A. Aydoğan, C. Karamurat, N. E. Altınışık, K. B. Vural, D. D. Kazancı, A. K. Doğu, D. Kaptan, H. C. Gemici, E. Yüncü, G. Umurtak, R. Duru, E. Fidan, Ö. Çevik, B. Erdoğu, T. Korkut, C. Knüsel, S. Haddow, C. S. Larsen, R. Özbal, F. Gerritsen, E. Özdoğan, U. O. Usanmaz, Y. C. Derici, M. Uçmazoğlu, F. Jay, M. Özdoğan, A. Götherström, Y. S. Erdal, A.-S. Malaspinas, Ç. Atakuman, F. Özer and M. Somel. *Out-of-Anatolia: cultural and genetic interactions during the Neolithic expansion in the Aegean*. 28th June 2024. DOI: 10.1101/2024.06.23.599747. URL: https://hal.science/hal-04801125 (cit. on p. 22).

[62] M. Kowalski, B. Malézieux, T. Moreau and A. Repetti. *Analysis and synthesis denoisers for forward-backward plug-and-play algorithms*. 11th Dec. 2024. URL: https://hal.science/hal-04786802 (cit. on p. 21).

[63] B. Malézieux, F. Michel, T. Moreau and M. Kowalski. *Where prior learning can and can't work in unsupervised inverse problems.* 14th Nov. 2024. URL: https://hal.science/hal-04782335 (cit. on p. 21).

[64] T. Monsel, E. Menier, L. Mathelin, O. Semeraro and G. Charpiat. *Neural DDEs with Learnable Delays for Partially Observed Dynamical Systems.* 1st Oct. 2024. URL: https://hal.science/hal-04715748 (cit. on p. 20).

[65] A. Pavão. *Hands-on tutorial on how to create your own challenge or benchmark.* 25th Jan. 2024. URL: https://inria.hal.science/hal-04415747 (cit. on p. 25).

[66] M. Quattromini, M. A. Bucci, S. Cherubini and O. Semeraro. *Graph Neural Networks and Differential Equations: A hybrid approach for data assimilation of fluid flows.* 15th Nov. 2024. URL: https://hal.science/hal-04794444 (cit. on p. 20).

[67] A. Szatkownik, C. Furtlehner, G. Charpiat, B. Yelmen and F. Jay. *Latent generative modeling of long genetic sequences with GANs.* 7th Aug. 2024. DOI: 10.1101/2024.08.07.607012. URL: https://hal.science/hal-04801130.

[68] A. Szatkownik, L. Planche, M. Demeulle, T. Chambe, M. Ávila-Arcos, E. Huerta-Sanchez, C. Furtlehner, C. Furtlehner, F. Jay and B. Yelmen. *Diffusion-based artificial genomes and their usefulness for local ancestry inference.* 31st Oct. 2024. DOI: 10.1101/2024.10.28.620648. URL: https://hal.science/hal-04801121 (cit. on p. 18).

[69] S. Treguer and M. Schoenauer. *Lessons learned from TAILOR benchmarks/challenges.* INRIA - team TAU, 2nd Dec. 2024. URL: https://inria.hal.science/hal-04923991 (cit. on p. 24).

[70] W.-W. Tu and A. Pavão. *Special designs and competition protocols.* 25th Jan. 2024. URL: https://inria.hal.science/hal-04416565 (cit. on p. 25).

[71] M. Videau, M. Zameshina, A. Leite, L. Najman, M. Schoenauer and O. Teytaud. *Evolutionary Retrofitting.* Oct. 2024. URL: https://hal.science/hal-04733511 (cit. on p. 17).

[72] B. Yelmen, M. Alver, E. B. R. Team, F. Jay and L. Milani. *Interpreting artificial neural networks to detect genome-wide association signals for complex traits.* 2024. DOI: 10.48550/arXiv.2407.18811. URL: https://hal.science/hal-04801114 (cit. on p. 18).

[73] M. Zameshina, M. Videau, A. Leite, M. Schoenauer, L. Najman and O. Teytaud. *Agnostic latent diversity enhancement in generative modeling.* 28th Oct. 2024. URL: https://hal.science/hal-04661473 (cit. on p. 17).

Other scientific publications

[74] B. Aristimunha, T. Moreau, S. Chevallier, R. Y. de Camargo and M.-C. Corsi. 'What is the best model for decoding neurophysiological signals? Depends on how you evaluate'. In: CNS 2024 - 33rd Annual Computational Neuroscience Meeting. Natal, Brazil, 20th July 2024. URL: https://inria.hal.science/hal-04743845.

12.3 Cited publications

[75] P. Caillou, J. Renault, J.-D. Fekete, A.-C. Letournel and M. Sebag. 'Cartolabe: A Web-Based Scalable Visualization of Large Document Collections'. In: *IEEE Computer Graphics and Applications* 41.2 (Apr. 2021), pp. 76–88. DOI: 10.1109/MCG.2020.3033401. URL: https://hal.inria.fr/hal-02499006 (cit. on p. 24).

[76] A. Decelle and C. Furtlehner. 'Exact Training of Restricted Boltzmann Machines on Intrinsically Low Dimensional Data'. In: *Physical Review Letters* (6th Sept. 2021). URL: https://hal.inria.fr/hal-03432350 (cit. on pp. 18, 19).

[77] B. Yelmen, A. Decelle, L. Ongaro, D. Marnetto, C. Tallec, F. Montinaro, C. Furtlehner, L. Pagani and F. Jay. 'Creating artificial human genomes using generative neural networks'. In: *PLoS Genetics* (4th Feb. 2021). DOI: 10.1371/journal.pgen.1009303. URL: https://hal.archives-ouvertes.fr/hal-03149930 (cit. on p. 18).

[78] R. Bresson, J. Cohen, E. Hüllermeier, C. Labreuche and M. Sebag. 'Neural Representation and Learning of Hierarchical 2-additive Choquet Integrals'. In: IJCAI-PRICAI-20 - Twenty-Ninth International Joint Conference on Artificial Intelligence and Seventeenth Pacific Rim International Conference on Artificial Intelligence. Yokohama, France, 11th July 2020, pp. 1984–1991. DOI: 10.24963/ijcai.2020/275. URL: https://hal.archives-ouvertes.fr/hal-03017078 (cit. on p. 6).

[79] R. Bresson, J. Cohen, E. Hüllermeier, C. Labreuche and M. Sebag. 'On the Identifiability of Hierarchical Decision Models'. In: 18th International Conference on Principles of Knowledge Representation and Reasoning (KR-2021). Online, France: International Joint Conferences on Artificial Intelligence Organization, 12th Nov. 2021, pp. 151–162. DOI: 10.24963/kr.2021/15. URL: https://hal.archives-ouvertes.fr/hal-03453996 (cit. on p. 6).

[80] A. Decelle, C. Furtlehner and B. Seoane. 'Equilibrium and non-Equilibrium regimes in the learning of Restricted Boltzmann Machines'. In: NeurIPS 2021. Proceedings NeurIPS 2021. Vancouver, United States, 6th Dec. 2021. URL: https://hal.archives-ouvertes.fr/hal-03518796 (cit. on p. 18).

[81] K. Gasnikova, P. Caillou, O. Allais and M. Sebag. 'Towards causal modeling of nutritional outcomes'. In: Causal Analysis Workshop Series (CAWS) 2021. Vol. 5. 19. online, United States, 2021. URL: https://hal.inria.fr/hal-03620867 (cit. on p. 23).

[82] N. Sanchez-Pi, L. Marti, J. Salomon, J. Sainte-Marie, O. Bernard, M. Sebag, M. Schoenauer, A. Maass, D. Eveillard, A. Abreu, C. de Vargas and P. A. Marquet. *OcéanIA: AI, Data, and Models for Understanding the Ocean and Climate Change*. 1st July 2021, pp. 1–64. URL: https://hal.archives-ouvertes.fr/hal-03274323 (cit. on p. 17).

[83] G. De Bie, H. Rakotoarison, G. Peyré and M. Sebag. *Distribution-Based Invariant Deep Networks for Learning Meta-Features*. 26th Feb. 2021. URL: https://hal.inria.fr/hal-03153200 (cit. on p. 7).

[84] J. K. et al. *Scaling Laws for Neural Language Models*. arXiv:2001.08361. 2020 (cit. on p. 7).

[85] F. Arceri, F. P. Landes, L. Berthier and G. Biroli. 'Glasses and aging: A Statistical Mechanics Perspective'. In: *Encyclopedia of Complexity and Systems Science (Living Reference)*. 50 pages, 24 figs. This is an updated version of a chapter initially written in 2009 for the Encyclopedia of Complexity and Systems Science (Springer). Mar. 2022. DOI: 10.1007/978-3-642-27737-5_248-2. URL: https://inria.hal.science/hal-02942375 (cit. on p. 19).

[86] G. Bied, C. Gaillac, M. Hoffmann, P. Caillou, B. Crépon, S. Nathan and M. Sebag. 'Fairness in job recommendations: estimating, explaining, and reducing gender gaps'. In: *Proceedings of the 1st Workshop on Fairness and Bias in AI co-located with 26th European Conference on Artificial Intelligence (ECAI 2023)*. Vol. 3523. Krakow, Poland: CEUR-WS.org, Oct. 2023. URL: https://inria.hal.science/hal-04438512 (cit. on p. 23).

[87] G. Bied, C. Gaillac, M. Hoffmann, S. Nathan, P. Caillou, B. Crépon and M. Sebag. 'Gender fairness in job recommendation: a case study'. In: *AI for HR and Public Employment Services*. Ghent (BE), Belgium, Feb. 2023. URL: https://inria.hal.science/hal-04025006 (cit. on p. 23).

[88] G. Bied, S. Nathan, E. Perennes, C. Gaillac, P. Caillou, B. Crépon and M. Sebag. 'Using Data from job seekers, job offers and past hirings to learn a Job Recommender System: the VADORE Project'. In: *AI for HR and Public Employment Services*. Ghent (BE), Belgium, Feb. 2023. URL: https://inria.hal.science/hal-04025000 (cit. on p. 23).

[89] G. Bied, E. Perennes, S. Nathan, V. Naya, P. Caillou, B. Crépon, C. Gaillac and M. Sebag. 'RECTO : REcommandation diminuant la Congestion par Transport Optimal'. In: *Proc. APIA 2023*. Ed. by C. Launois and C. Roussey. Vol. APIA. 2023. AFIA and ICube. Strasbourg, France: AFIA, July 2023, pp. 89–98. URL: https://hal.science/hal-04158566 (cit. on p. 23).

[90] M. M. Bronstein, J. Bruna, T. Cohen and P. Veličković. 'Geometric deep learning: Grids, groups, graphs, geodesics, and gauges'. In: *arXiv preprint arXiv:2104.13478* (2021) (cit. on p. 6).

[91] L. Chizat, E. Oyallon and F. Bach. 'On Lazy Training in Differentiable Programming'. In: *NeurIPS*. Ed. by H. Wallach et al. Vol. 32. 2019 (cit. on p. 7).

[92] T. Cohen and M. Welling. 'Group Equivariant Convolutional Networks'. In: *Proc. ICML*. Ed. by M. F. Balcan and K. Q. Weinberger. Vol. 48. PMLR. 2016, pp. 2990–2999 (cit. on p. 6).

[93] A. Decelle, G. Fissore and C. Furtlehner. 'Spectral dynamics of learning in restricted Boltzmann machines'. In: *EPL (Europhysics Letters)* 119.6 (2017), p. 60001 (cit. on p. 18).

[94] A. Decelle, G. Fissore and C. Furtlehner. 'Thermodynamics of Restricted Boltzmann Machines and Related Learning Dynamics'. In: *J. Stat. Phys.* 172 (2018), pp. 1576–1608 (cit. on p. 18).

[95] L. Dinh, J. Sohl-Dickstein and S. Bengio. 'Density estimation using Real NVP'. In: *Int. Conf. on Learning Representations (ICLR)*. 2017 (cit. on p. 9).

[96] S. Dong and M. Sebag. 'From graphs to DAGs: a low-complexity model and a scalable algorithm'. In: *ECML-PKDD 2022 - European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Grenoble, France, Sept. 2022. URL: https://inria.hal.science/hal-03798964 (cit. on p. 16).

[97] S. Dong, K. Uemura, A. Fujii, S. Chang, Y. Koyanagi, K. Maruhashi and M. Sebag. 'Learning Large Causal Structures from Inverse Covariance Matrix via Matrix Decomposition'. working paper or preprint. Oct. 2023. URL: https://inria.hal.science/hal-03885791 (cit. on p. 16).

[98] B. Donnot. 'Deep learning methods for predicting flows in power grids : novel architectures and algorithms'. Theses. Université Paris Saclay (COmUE), Feb. 2019. URL: https://tel.archives-ouvertes.fr/tel-02045873 (cit. on p. 22).

[99] D. Donoho. *Data Science at the Singularity*. 2023. arXiv: 2310.00865 [stat.OT] (cit. on p. 8).

[100] B. Donon. 'Deep statistical solvers & power systems applications'. Theses. Université Paris-Saclay, Mar. 2022. URL: https://theses.hal.science/tel-03624628 (cit. on p. 22).

[101] B. Donon, W. Liu, A. Marot, Z. Liu, I. Guyon and M. Schoenauer. 'Deep Statistical Solvers'. In: *NeurIPS 2020 - 34th Conference on Neural Information Processing Systems*. Vancouver / Virtuel, Canada, Dec. 2020. URL: https://hal.inria.fr/hal-02974541 (cit. on p. 20).

[102] R. Egele, I. Guyon, Y. Sun and P. Balaprakash. 'Is One Epoch All You Need For Multi-Fidelity Hyperparameter Optimization?' In: *ESANN 2023 - 31th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*. 5 pages, with extended appendices. Bruges / Hybrid, Belgium, Oct. 2023. URL: https://hal.science/hal-04219301 (cit. on p. 17).

[103] R. Egele, R. Maulik, K. Raghavan, B. Lusch, I. Guyon and P. Balaprakash. 'AutoDEUQ: Automated Deep Ensemble with Uncertainty Quantification'. In: *26TH International Conference on Pattern Recognition*. Montréal, Canada: IEEE, Aug. 2022, pp. 1908–1914. URL: https://hal.science/hal-03518597 (cit. on p. 17).

[104] J. Frankle and M. Carbin. 'The lottery ticket hypothesis: Finding sparse, trainable neural networks'. In: *arXiv preprint arXiv:1803.03635* (2018) (cit. on p. 5).

[105] J. Giles. 'Computational social science: Making the links'. In: *Nature - News* 488.7412 (2012), pp. 448–450 (cit. on p. 8).

[106] A. Hyvarinen, H. Sasaki and R. Turner. 'Nonlinear ICA Using Auxiliary Variables and Generalized Contrastive Learning'. In: *Proc. AISTAT*. Ed. by K. Chaudhuri and M. Sugiyama. Vol. 89. PMLR. 2019, pp. 859–868. URL: https://proceedings.mlr.press/v89/hyvarinen19a.html (cit. on p. 6).

[107] A. Jacot, F. Gabriel and C. Hongler. 'Neural Tangent Kernel: Convergence and Generalization in Neural Networks'. In: *Advances in Neural Information Processing Systems 31*. 2018, pp. 8571–8580 (cit. on p. 19).

[108] A. Jacot, F. Gabriel and C. Hongler. 'Neural Tangent Kernel: Convergence and Generalization in Neural Networks'. In: *NeurIPS*. Ed. by S. Bengio et al. Vol. 31. 2018 (cit. on p. 7).

[109] T. G. H. Khuong and B. K. Rachmat. 'Auto-survey Challenge'. In: *JDSE 2023 - 8th Junior Conference on Data Science and Engineering*. Orsay, France, Sept. 2023. URL: https://inria.hal.science/hal-04206578 (cit. on p. 24).

[110] A. Lacan, M. Sebag and B. Hanczar. 'GAN-based data augmentation for transcriptomics: survey and comparative assessment'. In: *Bioinformatics* 39.Supplement_1 (June 2023), pp. i111–i120. DOI: 10.1093/bioinformatics/btad239. URL: https://univ-evry.hal.science/hal-04263316 (cit. on p. 15).

[111] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirnsberger, M. Fortunato, A. Pritzel, S. Ravuri, T. Ewalds, F. Alet, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, J. Stott, O. Vinyals, S. Mohamed and P. Battaglia. 'GraphCast: Learning skillful medium-range global weather forecasting'. In: *Science* 382.6677 (2022). cite arxiv:2212.12794, pp. 1416–142. URL: http://arxiv.org/abs/2212.12794 (cit. on p. 7).

[112] D. Lazer, A. Pentland, L. Adamic, S. Aral, A. L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy and M. Van Alstyne. 'Life in the network: the coming age of computational social science'. In: *Science* 323.5915 (2009), pp. 721–723 (cit. on p. 8).

[113] Y. Lecun, S. Chopra, R. Hadsell, M. Ranzato and F. Huang. 'A tutorial on energy-based learning'. In: *Predicting structured data*. Ed. by G. Bakir, T. Hofman, B. Scholkopt, A. Smola and B. Taskar. MIT Press, 2006 (cit. on p. 6).

[114] A. Leite and M. Schoenauer. 'Memetic Semantic Genetic Programming for Symbolic Regression'. In: *Lecture Notes in Computer Science*. Ed. by G. Pappa, M. Giacobini and Z. Vasicek. Vol. LNCS-13986. Genetic Programming. Species Society. Brno, Czech Republic: Springer Nature Switzerland, Apr. 2023, pp. 198–212. DOI: 10.1007/978-3-031-29573-7_13. URL: https://hal.science/hal-04563511 (cit. on p. 17).

[115] M. Leyli-Abadi, D. Danan, M. Yagoubi, S.-E. Attoui, A. Marot, J. Picault, B. Donnot, P. Dimitrov, C. Etienam and A. Farjallah. 'LIPS - Learning Industrial Physical Simulation benchmark suite'. In: *NeurIPS - Data & Benchmark Track*. New Orleans, United States, Nov. 2022. URL: https://hal.science/hal-03780816 (cit. on p. 7).

[116] Z. Liu and M. Tegmark. 'Machine Learning Hidden Symmetries'. In: *Phys. Rev. Lett.* 128 (18 2022), p. 180201 (cit. on p. 6).

[117] C. Louizos, K. Ullrich and M. Welling. 'Bayesian compression for deep learning'. In: *Advances in neural information processing systems* 30 (2017) (cit. on p. 5).

[118] J. Medina-Tretmanis, F. Jay, M. C. Ávila-Arcos and E. Huerta-Sanchez. 'Simulation-based Benchmarking of Ancient Haplotype Inference for Detecting Population Structure'. In: *Human Population Genetics and Genomics* (Sept. 2023), pp. 1–25. DOI: 10.1101/2023.09.28.560049. URL: https://hal.science/hal-04244011 (cit. on p. 22).

[119] T. Monsel, L. Mathelin, O. Semeraro and G. Charpiat. 'End-to-end learning of dynamical systems with the Mori-Zwanzig formalism'. In: *CSE 2023 - SIAM Computational Science and Engineering*. Amsterdam, Netherlands, Feb. 2023. URL: https://hal.science/hal-04406551 (cit. on p. 20).

[120] A. Pavao, I. Guyon, A.-C. Letournel, X. Baró, H. Escalante, S. Escalera, T. Thomas and Z. Xu. *CodaLab Competitions: An open source platform to organize scientific challenges*. Technical Report. Université Paris-Saclay, FRA., Apr. 2022. URL: https://inria.hal.science/hal-03629462 (cit. on p. 15).

[121] A. Poinsot and A. Leite. 'A Guide for Practical Use of ADMG Causal Data Augmentation'. Workshop on the pitfalls of limited data and computation for Trustworthy ML, ICLR 2023, Kigali, Rwanda. Mar. 2023. URL: https://hal.science/hal-04250612 (cit. on p. 16).

[122] M. Raissi, P. Perdikaris and G. E. Karniadakis. 'Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations'. In: *Journal of Computational physics* 378 (2019), pp. 686–707 (cit. on p. 19).

[123] H. Rakotoarison, L. Milijaona, A. Rasoanaivo, M. Sebag and M. Schoenauer. 'Learning Meta-features for AutoML'. In: *ICLR 2022 - International Conference on Learning Representations (spotlight)*. Virtual, United States, Apr. 2022. URL: https://inria.hal.science/hal-03583789 (cit. on p. 7).

[124] V. Schimmenti. 'Temporal and spatial correlations in earthquake dynamics : physical modeling and data analysis'. Theses. Université Paris-Saclay, Nov. 2023. URL: https://theses.hal.science/tel-04344778 (cit. on p. 16).

[125] G. Serré, E. Boguslawski, B. Donnot, A. Pavão, I. Guyon and A. Marot. 'Reinforcement learning for Energies of the future and carbon neutrality: a Challenge Design'. In: *SSCI 2022 - IEEE Symposium Series on Computational Intelligence*. IEEE. Singapour, Singapore, Dec. 2022. URL: https://hal.science/hal-03726294 (cit. on p. 22).

[126] C. Shi, D. M. Blei and V. Veitch. 'Adapting Neural Networks for the Estimation of Treatment Effects'. In: *NeurIPS 2019*. 2019, pp. 2503–2513. URL: https://proceedings.neurips.cc/paper/2019/hash/8fb5f8be2aa9d6c64a04e3ab9f63feee-Abstract.html (cit. on p. 6).

[127] I. Stengers. *Another Science Is Possible*. Open Humanities Press, 2013 (cit. on p. 8).

[128] Y. Wang and D. Blei. 'A Proxy Variable View of Shared Confounding'. In: *Proc. 38th ICML*. Ed. by M. Meila and T. Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 2021, pp. 10697–10707. URL: https://proceedings.mlr.press/v139/wang21c.html (cit. on p. 6).

[129] Z. Xu, S. Escalera, A. Pavao, M. Richard, W.-W. Tu, Q. Yao, H. Zhao and I. Guyon. 'Codabench: Flexible, Easy-to-Use and Reproducible Meta-Benchmark Platform'. In: *Patterns* (July 2022). URL: https://hal.science/hal-03374222 (cit. on p. 15).

[130] A. Yale, S. Dash, R. Dutta, I. Guyon, A. Pavao and K. P. Bennett. 'Generation and Evaluation of Privacy Preserving Synthetic Health Data'. In: *Neurocomputing* 416 (Nov. 2020), pp. 244–255. DOI: 10.1016/j.neucom.2019.12.136. URL: https://hal.inria.fr/hal-03158544 (cit. on p. 18).