RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

2025 *Activity Report*

RESEARCH CENTRE: **Inria Branch at the University of Montpellier**

IN PARTNERSHIP WITH: **INSERM, Université de Montpellier**

Project-Team

PREMEDICAL

Precision Medicine by Data Integration and Causal Learning

In collaboration with Institut Desbrest d'Épidémiologie et de Santé Publique (IDESP)

Project-Team PREMEDICAL

Creation of the Project-Team: 2022 June 01

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

Keywords

Computer sciences and digital sciences

A3.4. – Machine learning and statistics

A4. – Security and privacy

A4.8. – Privacy-enhancing technologies

A6.1. – Methods in mathematical modeling

A9. – Artificial intelligence

A9.2. – Machine learning

A9.6. – Decision support

A9.9. – Distributed AI, Multi-agent

Other research topics and application domains

B2. – Digital health

B2.2. – Physiology and diseases

B2.3. – Epidemiology

Contents

Project-Team PREMEDICAL 1

1 Team members, visitors, external collaborators 5

2 Overall objectives 6

3 Research program 7
3.1 Research Axis 1: Personalized medicine by optimal prescription of treatment 7
3.2 Research axis 2: Personalized medicine by integration of different data sources 8
3.3 Research Axis 3: Personalized medicine with privacy and fairness guarantees 9

4 Application domains 10

5 Social and environmental responsibility 11
5.1 Impact of research results 11

6 Highlights of the year 11
6.1 Awards 11
6.2 Hackathon PREMEDICAL-CHU-CINES 11
6.3 Inria International Chair 12

7 Latest software developments, platforms, open data 12
7.1 Latest software developments 12
7.1.1 declearn 12
7.1.2 CaMeA 13
7.1.3 missMDA 13
7.1.4 factominer 13
7.2 New platforms 14

8 New results 14
8.1 Treatment effect estimation 14
8.2 Federated Learning 15
8.3 Learning with Privacy Guarantees 16
8.4 Handling missing data 17
8.5 Application domain 18

9 Bilateral contracts and grants with industry 19
9.1 Bilateral contracts with industry 19

10 Partnerships and cooperations 21
10.1 International research visitors 21
10.1.1 Visits of international scientists 21
10.1.2 Visits to international teams 22
10.2 National initiatives 23
10.2.1 PEPR Digital Health 23
10.2.2 PEPR Cybersecurity 23
10.2.3 Inria Challenge FedMalin 23
10.3 PANAME Project 24
10.3.1 ANR JCJC PRIDE 24
10.4 Regional initiatives 24

11 Dissemination **25**
11.1 Promoting scientific activities . . . 25
11.1.1 Scientific events: organization . . . 25
11.1.2 Scientific events: selection . . . 25
11.1.3 Journal . . . 26
11.1.4 Invited talks . . . 26
11.1.5 Contributed Talks . . . 28
11.1.6 Leadership within the scientific community . . . 28
11.1.7 Scientific expertise . . . 28
11.1.8 Research administration . . . 28
11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach . . . 29
11.2.1 Supervision . . . 29
11.2.2 Juries . . . 29
11.2.3 Teaching . . . 30
11.3 Popularization . . . 30
11.3.1 Productions (articles, videos, podcasts, serious games, ...) . . . 30
11.3.2 Participation in Live events . . . 30

12 Scientific production **31**
12.1 Major publications . . . 31
12.2 Publications of the year . . . 31
12.3 Cited publications . . . 37

1 Team members, visitors, external collaborators

Research Scientists

- Julie Josse [Team leader, INRIA, Senior Researcher, HDR]
- Aurélien Bellet [INRIA, Senior Researcher, HDR]
- Clement Berenfeld [INRIA, Advanced Research Position, from Oct 2025]
- Mathieu Even [INRIA, Researcher, from Oct 2025]
- Nicolas Papernot [INRIA, Chair, from Oct 2025, International Chair]

Faculty Members

- Pascal Demoly [UNIV MONTPELLIER, Professor]
- Nicolas Molinari [UNIV MONTPELLIER, Professor]

Post-Doctoral Fellows

- Clement Berenfeld [INRIA, Post-Doctoral Fellow, from Apr 2025 until Sep 2025]
- Linus Bleistein [UNIV PSL, from Mar 2025 until Sep 2025]
- Mathieu Dagreou [INRIA, Post-Doctoral Fellow]
- Mathieu Even [INRIA, Post-Doctoral Fellow, until Sep 2025]
- Jean-Baptiste Fermanian [INRIA, Post-Doctoral Fellow, from Oct 2025]
- Christian Janos Lebeda [INRIA, Post-Doctoral Fellow]
- Jeffrey Naef [INRIA, Post-Doctoral Fellow, until Jan 2025]

PhD Students

- Thomas Boudou [INRIA]
- Ahmed Boughdiri [INRIA]
- Tess Breton [UNIV PARIS - CITE, from Oct 2025]
- Ioan Tudor Cebere [INRIA]
- Agathe Chabassier [WITHINGS, CIFRE, from Sep 2025]
- Ghita Fassy El Fehri [INRIA]
- Maxime Fosset [UNIV MONTPELLIER]
- Laura Fuentes Vicente [UNIV MONTPELLIER]
- Remi Khellaf [UNIV MONTPELLIER]
- Charlotte Voinot [SANOFI, CIFRE]

Technical Staff

- Mariette Dupuy [INRIA, Engineer, from Nov 2025]
- Charif El Gataa [INRIA, Engineer, from May 2025 until Sep 2025]
- Dhia Eddine Merzougui [INRIA, Engineer, from Sep 2025]
- Christophe Muller [INRIA, Engineer, until Aug 2025]

Interns and Apprentices

- Devaganthan Sivakumar Srirangan [INRIA, Intern, from Oct 2025]

Administrative Assistant

- Claire-Marine Parodi [INRIA]

Visiting Scientists

- Charif El Gataa [Univ Torino, until Mar 2025]
- Krystyna Grzesiak [UNIV WROCLAW, until Apr 2025]
- Aishik Mandal [UNIV Darmstadt, from Oct 2025]
- Emma Torrini [UNIV FLORENCE, from Sep 2025]

External Collaborators

- Helene Bonneau-Chloup [ELIXIR HEALTH, until Mar 2025]
- Gaelle Dormion [ELIXIR HEALTH]
- Geneviève Robin [CNRS, from Oct 2025]

2 Overall objectives

The objective of the team (**Pre**cision **Me**dicine by **D**ata **I**ntegration and **Ca**usal **L**earning) is to develop the next generation of methods/algorithms to extract knowledge from health data and improve the care of patients. More specifically, the goal is to develop learning tools for **personalized treatment effect prediction** and for **predicting outcome**, while **integrating different data sources** to **guide decisions** made by clinicians and authorities. PreMeDICaL has three research axes:

1. *Personalized medicine through optimal treatment prescription.* The objective is to develop causal inference techniques for **(dynamic) policy learning**—allocating the best treatment to each individual at the right time—by leveraging both experimental data from Randomized Controlled Trials (RCTs) and non-experimental data (e.g., observational data from Electronic Health Records, cohorts, etc.). Combining these data sources will enable **better design of future RCTs** and, in the longer term, may **transform the standards of evidence required to bring treatments to market**, potentially allowing for the **launch of new drugs without traditional RCTs**, and doing so more efficiently.

2. *Personalized medicine through integration of diverse data sources.* Our purpose is to learn (e.g. build predictive models) from heterogeneous data, such as continuous time monitoring data and static clinical data, and from **decentralized data using federated learning** while **handling missing values** and **increasing the reliability and confidence of predictive model outputs**.

3. *Personalized medicine with privacy and fairness guarantees.* We seek to develop approaches to **ensure the confidentiality** of medical data and guarantee that models do not leak sensitive information. We additionally build methods to **handle fairness constraints** to ensure that models exhibit similar performance across different population groups.

Our ambition is to **bring methodological innovation directly to stakeholders**—including patients, clinicians, and regulators. Accordingly, beyond the development of novel methodologies, the project targets innovative solutions to major public health challenges across various application domains (e.g., respiratory allergies, traumatology, oncology, fertility, neurodegenerative diseases). In addition to **leveraging machine learning algorithms and relevant data**, it is essential to **integrate clinical expertise and existing guidelines** to ensure practical and effective outcomes. The long-term objective is to establish **clear, reproducible pipelines, methodologies, and software tools** (such as clinical decision support systems) that enable both significant scientific contributions and societal impact. These innovations aim to enhance the quality of patient care and create meaningful change in the medical profession by **facilitating earlier access to innovative solutions and more efficient treatments**. The team contributes to precision medicine (where the treatment/device is adapted on a patient basis) and to translational medicine, which aims at bridging the gap between fundamental research and its practical use.

3 Research program

3.1 Research Axis 1: Personalized medicine by optimal prescription of treatment

In machine learning (ML)/artificial intelligence (AI) progress has yielded powerful predictive models, yet they rely on correlations and lack an understanding of underlying mechanisms or intervention strategies. Causality is crucial for actionable insights, recommendations, and addressing "what if" scenarios, with applications in health, public policies, econometrics, and advertising. Causal inference gains prominence for addressing AI challenges like interpretability and robustness offering solutions akin to "AI-like human" approaches in novel settings. This axis aims to innovate causal machine learning at the AI-personalized medicine intersection, optimizing treatment allocation and enabling drug launches without randomized control trials (RCTs).

Randomized controlled trials are considered the gold standard approach for assessing the causal effect (i.e., the treatment effect) of an intervention or a treatment on an outcome of interest. Indeed, the allocation of the treatment is under control, which implies that there are no confounding factors (the distribution of covariates for treated and control patients is asymptotically balanced) that could interfere with the treatment and simple estimators (such as the difference in mean effect between the treated and controls) can be used to consistently estimate the average treatment effect (ATE). However, RCTs can come with drawbacks. They can be expensive, take a long time to set up, and be compromised by insufficient sample size due to either recruitment difficulties or restrictive inclusion/exclusion criteria. These criteria can lead to a narrowly defined trial sample that differs markedly from the population potentially eligible for the treatment (distributional shift). Therefore, the findings from RCTs can lack generalizability (or *external validity*). This has been largely published in the field of respiratory and allergic diseases, see for instance [81] which highlights that the population from RCTs represents less than 10% of the population that will receive treatments.

In contrast, there is an abundance of observational data, collected without systematically designed interventions. Such data can come from different sources: they can be collected from research sources (such as disease registries, cohorts, biobanks, epidemiological studies), or they can be routinely collected (through electronic health records, insurance claims, administrative databases, patients' App, etc). In that sense, observational data can be readily available, can include large samples representative of the target populations, and can be less costly than RCTs. To harness observational data for estimating treatment effects in health domains, regulatory frameworks—including those developed by the U.S. Food and Drug Administration (FDA)—promote the use of "real-world data" (RWD). RWD is defined as data derived from sources other than randomized clinical trials, and its use is encouraged for regulatory decision-making. Clinical evidence regarding the usage and potential benefits or risks of a medical product derived from the analysis of RWD is named Real World Evidence (RWE). The European Medicines Agency (EMA) is also a very active regulatory authority working with RWD to facilitate development and access to medicines. However, despite the large number of methods available to estimate the causal treatment effect from observational data such as matching, inverse probability weighting (IPW) or more recent doubly robust methods based on machine learning there are often concerns about the quality of these "big data" and causal claims. Indeed, building on observational data is still not consensual due to the lack of controlled experimental interventions, which opens the door to confounding biases (lack of *internal validity*).

Observational data and clinical trial data can provide different perspectives when evaluating an intervention or a medical treatment. Combining the information gathered from experimental and observational data is a promising avenue for medical research, because the knowledge acquired from integrative analyses could not be gathered from a single-source analysis alone. Three potential high impact applications of observational and clinical data are:

1. Predicting the effect of a treatment estimated on a RCT, on a new target population (generalization);
2. Comparing RCTs and RWE to validate observational methods;
3. Better estimation of heterogeneous treatment effects.

There is an abundant literature on bridging the findings from an RCT to a target population and combining both sources of information. Similar problems have been termed as *transportability*, and *data fusion* and have connections to the *covariate shift/domain generalization* problem in ML. [76] reviewed the methods to (a) generalize the treatment effect while integrating the distributional shift (IPSW, g-formula, AIPSW, calibration weighting, etc.), or (b) improve the estimate of the conditional average treatment effect (CATE, i.e. heterogeneous effect) while correcting for confounding factors not measured in the observational study. However, these methods have many shortcomings and there are still many challenges to address. We provide below examples of methodological locks we will overcome.

- Handling missing values and unmeasured covariates with multi-source data;
- Transfert Learning of optimal individualized treatment regimes with right-censored survival data;
- Policy learning and dynamic treatment policy with missing values;
- Generalization of different causal measures: Risk Ratio, Survival Ratio, etc;
- Providing finite sample guarantees;
- Study of causal effects in metric spaces
- Guide variable selection and provide variable-importance measures and tests in treatment effects setting

Such development will have significant societal impact in patient care and cost reduction, ultimately guiding future RCT designs.

3.2 Research axis 2: Personalized medicine by integration of different data sources

In this axis we focus both on integrating heterogeneous data/multiview/multimodal (time series, images, text, numerical or categorical data) potentially from different centers to establish predictive, as well as quantifying the uncertainty associated to predictive models. For the former, we will focus on handling missing values and on federated learning strategies, while for the latter we will consider uncertainty quantification approaches.

Federated learning [78] is a recent paradigm which enables model training across decentralized devices or servers holding local data samples, without exchanging them. Only the model updates, not the raw data, are sent to a central server, where they are aggregated to improve the global model. In the medical domain, federated learning helps to address privacy concerns by allowing models to be trained on data distributed across various healthcare institutions and/or companies without centrally aggregating sensitive patient information. This facilitates collaborative inference without compromising data security, making it particularly valuable for developing robust and generalizable medical AI models across diverse datasets while respecting privacy regulations.

Most statistical learning and artificial intelligence methodologies provide point predictions, without any indication of the degree of confidence that can be given to these predictions (i.e. without predictive intervals). This lack of uncertainty quantification of predictive models is a **major barrier to the adoption of**

powerful machine learning methods by society. Probabilistic forecasts, i.e. predicting the entire distribution probability and not only the conditional expectation, could partially tackle this issue but they are only valid asymptotically, require strong assumptions on the data (e.g. normality) or/and are model-dependent. The emergent field of conformal prediction (CP) [88, 83, 79] is a promising framework for **distribution-free uncertainty quantification**. It is a general procedure to build predictive intervals for any predictive model (including black-box methods such as deep learning), which are valid (i.e. achieve nominal marginal coverage), in finite sample, and without assumption on the data generation process except the exchangeability. This is extremely promising for decision support tools in critical applications: healthcare, autonomous driving, etc. An extension of CP (Conformalized Quantile Regression, [85]) was used to predict the U.S. presidential elections (2020) by the Washington Post.

We provide below examples of methodological challenges we will overcome.

- Relationship between the different sources;
- (Informative) missing values in time series and structured by blocks;
- Conformal prediction with missing values [90]; Relationship between predictive intervals and confidence intervals
- Federated learning with missing values;
- Federated causal inference.

3.3 Research Axis 3: Personalized medicine with privacy and fairness guarantees

In this axis, we aim to address privacy and fairness concerns in machine learning, with a focus on the challenges raised by medical applications. By integrating privacy and fairness into the design of the algorithms, we can enhance the trustworthiness of machine learning applications, promote ethical practices, and facilitate the responsible deployment of personalized medicine technologies for the benefit of diverse patient populations.

While training ML models on personal or otherwise confidential data can be beneficial in many applications such as healthcare, this can also lead to undesirable disclosure of sensitive information. Take for instance patient records, which often contain highly personal and identifiable information such as medical histories, diagnostic results, and genetic data. If a machine learning model trained on this data is not appropriately designed and secured, it may be possible for an attacker to deduce private information about individuals by analyzing the output of the model. Indeed, concrete attacks have been designed to predict whether a particular individual was part of the training set [87], and even to reconstruct some of the training data points [82]. Privacy-preserving machine learning aims to mitigate these concerns by incorporating techniques that safeguard sensitive information during the training and deployment of models. We focus on Differential Privacy (DP), a framework that provides a mathematical definition of privacy guarantees. In a nutshell, DP ensures that the inclusion or exclusion of any single data point does not significantly impact the output distribution of the training algorithm, thereby bounding the amount of information that can be inferred from the trained model about any individual in the dataset. DP requires to incorporate a certain amount of randomness into the algorithms, and thus yields a necessary trade-off between privacy and utility (e.g., accuracy of the resulting model). A key challenge is then to design methods that achieve the best possible trade-offs. We consider both centralized training by a trusted curator, and federated/decentralized training by participants who do not trust each other. We seek to characterize the achievable trade-offs, and to design algorithms with optimal privacy-utility trade-offs for a variety of machine learning and statistical inference tasks. Finally, we will also consider the relationship between missing values imputation methods and the generation of synthetic data which is often used to tackle privacy constraints.

Fairness considerations are also vital in machine learning to avoid bias in algorithms. Indeed, biased models could lead to unequal treatment of individuals based on factors like ethnicity or gender [86], potentially exacerbating healthcare disparities. For instance, if a machine learning model is trained predominantly on data from a specific demographic group, it may not generalize well to other groups, leading to inaccurate predictions for underrepresented populations. This can result in suboptimal healthcare outcomes, with certain individuals receiving inadequate attention or misdiagnoses. Additionally, historical biases present in healthcare data may be learned by machine learning models and perpetuated in their predictions. We aim to address these fairness challenges by incorporating fairness considerations into the machine learning pipeline, i.e., during data collection and preprocessing, model training and/or evaluation. An approach of particular interest is the introduction of group fairness constraints during the training phase [89]. Such constraints explicitly define the desired level of fairness and prevent the model from making predictions that disproportionately favor or disfavor specific population groups. As for privacy, we seek to study fairness in centralized training, but also in the context of federated learning which raises specific challenges as fairness on decentralized data becomes difficult to measure globally.

In addition to considering privacy and fairness in machine learning separately, we also aim to understand the interplay and potential tension between these two requirements, as well as to design algorithms that can provide optimal and tunable trade-offs.

4 Application domains

PreMeDICaL has a wide range of applications, including oncology, neurodegenerative diseases, fertility, and the use and evaluation of digital and medical devices. However, its main focus lies in trauma-care and respiratory diseases with a particular emphasis on asthma, as detailed below.

Traumatology: Trauma is the leading cause of death and disabilities among 16-45 year-olds, and a central challenge is to reduce both under- and over-triage to optimize resource allocation and patient outcomes. The Traumatrix project is a flagship collaboration between clinicians from the Traumabase network, Inria, CNRS, Ecole Polytechnique, EHESS, and the company Capgemini Invent (through skills sponsorship). Since 2019, the consortium has assembled a unique database of over 50,000 trauma cases from 40 centers, covering the entire care pathway from accident scene to hospital discharge. This resource has fueled the development of causal inference methods and predictive models capable of handling heterogeneous and incomplete data. The project also addresses critical methodological challenges: quantifying uncertainty so that algorithms can output "I don't know", ensuring fairness across patient subgroups, and delivering robust real-time decision support in high-stakes settings. Beyond methodological advances, Traumatrix has provided an exceptional training ground where doctoral and postdoctoral researchers can directly test innovations in clinical practice. Traumatrix has reached a decisive stage. Predictive models were validated for real-time deployment in collaboration with the SAMU, supported by a national PREPS grant (Programme de Recherche sur la Performance du Systeme de soins). In February 2026, a large-scale clinical trial begins across 16 emergency regulation centers, covering 22 million inhabitants in France, to evaluate the real-world impact of these models. The system predicts urgent needs—such as hemorrhage control or neurosurgery—at dispatch, with quantified confidence scores to guide decisions. Overall, this application domain fostered multiple methodological advancements in causal inference (e.g., generalization of treatment effects between populations, handling multiple outcomes), missing data imputation, synthetic data generation and federated learning.

Respiratory diseases: For more than 30 years, there has been an increase in the number of chronic non-communicable diseases (NCD), such as asthma and allergies. Allergies are the fourth most common chronic disease in the world. The World Health Organization (WHO) predicts that by 2050, one in two people in the world will suffer from allergies. In France, the number of people suffering from allergies has doubled in 20 years, particularly among children and young people. Although the expression of these diseases results from the interaction between the genetic background and the environment, especially through epigenetic mechanisms, their sudden increase is solely due to the environmental changes that occurred in the last decades because of the Western lifestyle, the genetic heritage requiring centuries to change. A full understanding of the complexity of chronic NCD prompts researchers to analyze large data utilizing proper markers and tools (e.g., biological, clinical, behavioral, economic, social, demographic, environmental data, patient experience, patient social networks) in an etiological and evaluative way to determine phenotypical patients' pathways, explain their impacts, their causes, their influences, prevent them and improve their prognosis. Integrating these different sources of information, collected by several actors (healthcare professionals, public authorities or patients themselves), thus offers new opportunities to design personalized solutions by adapting treatment to the patient and the organizational context, leading to improved patient care and prevention policies.

5 Social and environmental responsibility

5.1 Impact of research results

From a methodological point of view, the aim is to improve and develop new statistical and ML methods for establishing evidence on the efficiency of treatment by data enrichment (data fusion) and for predicting outcomes quantifying the uncertainty. An important output of this research is that these methodological works have a concrete impact on designing future clinical trials and that the new methodology will be supported by regulatory authorities. Indeed, exploiting both RCTs and observational data serves different purposes such as prediction of the treatment effect on new populations, increasing the generalization of clinical trials (so that they are more representative of the patient population who may benefit from the treatment) and also defining new inclusion criteria (because we identify subgroups who can benefit from treatment). This research is part of the PEPR project "Next methodological challenges in clinical trials in the era of digital health". Through axis 3 of our research program, we also aim to design methods that can effectively address and integrate societal requirements, with a particular focus on fairness and privacy. This involves developing algorithms that not only optimize performance but also ensure equitable treatment of diverse groups and protect sensitive data throughout the machine learning pipeline. By incorporating fairness, we strive to minimize biases and disparities in decision-making, ensuring that outcomes are inclusive and just. On the privacy front, our efforts include designing techniques that safeguard individuals' data, such as employing differential privacy, federated learning, or encryption mechanisms to prevent unauthorized access or misuse. Our overarching goal is to create systems that align with ethical principles and societal values, paving the way for responsible and trustworthy artificial intelligence applications.
From a technological point of view, the aim is to provide softwares (with open access softwares in a first place) for these methods to be applied in practice by studies stakeholders, clinicians and the clinical trial community.
From the clinical and patients point of view, the different projects aim to quantify the clinical benefit of intervention (over time), taking into account all patient characteristics, and to provide useful clinical prognosis tools allowing clinicians to optimally treat every patient, while also guaranteeing some level of fairness and privacy. The aim is to give patients better care and early access to innovation. In addition, these works can lead to a better adoption by the medical community of certain (advanced) techniques used to estimate the effects of treatment on patients (by comparing the results obtained in an RCT with the RWE).
From a public-health point of view, the aim is to guide decisions made by investigators, sponsors and authorities. Better trial designs may also have an important impact in terms of cost reduction. Finally, we aim at having a significant impact in the field of allergy treatments providing new knowledge that may change guidelines and practice.

6 Highlights of the year

6.1 Awards

- Tudor Cebere: recipient of the 2025 Google PhD Fellowship Program, recognizing outstanding graduate students who are conducting exceptional and innovative research in computer science and related fields. The program provides direct financial support for their PhD pursuits and connects each Fellow with a dedicated Google Research Mentor.
- Mathieu Even: Accessit for the Gilles Kahn PhD Award for the best thesis, French Computer Science Society (Société Informatique de France), sponsored by the Académie des Sciences.

6.2 Hackathon PREMEDICAL-CHU-CINES

CINES, Erios team of the University Hospital (CHU) Montpellier and the Inria PREMEDICAL team joined together in a seminar and hackathon from 3 to 5 November 2025 in Carcassone. The three teams collaborated on the Adastra supercomputer, working on large language model (LLM) evaluation from CHU medical texts. The topics spanned from LLM reliability with conformal prediction, to combining expert and LLM annotations with predictive powered inference, and finally assessing algorithm privacy through membership inference attacks. This hackathon serves as a basis and kickstart for longer term partnerships with the two other colaborators.

6.3 Inria International Chair

Nicolas Papernot has been awarded an Inria International Chair in the PREMEDICAL team. He is a professor at Toronto University, member of the Vector Institute and recipient of a CIFAR AI Chair. He is a leading international expert on privacy and federated learning. He will work together with the team, during his 2025/2026 sabbatical and the upcoming 3 years, to advance privacy and federated learning, with applications to precision medicine.

7 Latest software developments, platforms, open data

7.1 Latest software developments

7.1.1 declearn

Keyword: Federated learning

Scientific Description: declearn is a python package providing with a framework to perform federated learning, i.e. to train machine learning models by distributing computations across a set of data owners that, consequently, only have to share aggregated information (rather than individual data samples) with an orchestrating server (and, by extension, with each other).

The aim of declearn is to provide both real-world end-users and algorithm researchers with a modular and extensible framework that:

(1) builds on abstractions general enough to write backbone algorithmic code agnostic to the actual computation framework, statistical model details or network communications setup

(2) designs modular and combinable objects, so that algorithmic features, and more generally any specific implementation of a component (the model, network protocol, client or server optimizer...) may easily be plugged into the main federated learning process - enabling users to experiment with configurations that intersect unitary features

(3) provides with functioning tools that may be used out-of-the-box to set up federated learning tasks using some popular computation frameworks (scikit- learn, tensorflow, pytorch...) and federated learning algorithms (FedAvg, Scaffold, FedYogi...)

(4) provides with tools that enable extending the support of existing tools and APIs to custom functions and classes without having to hack into the source code, merely adding new features (tensor libraries, model classes, optimization plug-ins, orchestration algorithms, communication protocols...) to the party.

Parts of the declearn code (Optimizers,...) are included in the FedBioMed software.

At the moment, declearn has been focused on so-called "centralized" federated learning that implies a central server orchestrating computations, but it might become more oriented towards decentralized processes in the future, that remove the use of a central agent.

Functional Description: This library provides the two main components to perform federated learning:

(1) the client, to be run by each participant, performs the learning on local data et releases only the result of the computation

(2) the server orchestrates the process and aggregates the local models in a global model

News of the Year: Two major releases with key new functionalities including algorithms for group fairness and the ability to use secure aggregation.

URL: https://gitlab.inria.fr/magnet/declearn/declearn2

Contact: Aurélien Bellet

Participants: Paul Andrey, Aurélien Bellet, Nathan Bigaud, Marc Tommasi, Nathalie Vauquier

Partner: CHRU Lille

7.1.2 CaMeA

Name: Causal Meta-Analysis for Aggregated Data

Keywords: Causality, Randomised control trials

Functional Description: Based on results from multiple clinical trials (contingency tables cross-tabulating treatment and response), CaMeA measures the effect of a treatment or intervention using various metrics, such as Risk Ratio, Risk Difference, and others.

Publication: hal-05086779

Contact: Julie Josse

Participants: Julie Josse, Clement Berenfeld

7.1.3 missMDA

Keyword: Missing data

Functional Description: The missMDA package is dedicated to missing values in and with Multivariate Data Analysis. It allows one to apply PCA, MCA, FAMD and MFA on incomplete data. It performs single and multiple imputation for continuous, categorical and mixed data based on principal components methods

URL: `http://factominer.free.fr/missMDA/index.html`

Contact: Julie Josse

Partner: AGROCAMPUS

7.1.4 factominer

Keywords: Dimensionality reduction, PCA, Text mining, Clustering

Functional Description: The FactoMineR package is dedicated to performing principal components methods to explore, sum-up and visualize data. Dimensionality reduction methods include PCA, correspondence analysis (CA) for count data such as documents-words data, multiple correspondence analysis (MCA) for categorical data such as survey data, factorial analysis of mixed data (FAMD) for both types of variables as well as methods for groups of variables, of individuals (multiple factorial analysis, MFA), for hierarchy . . .

References: https://husson.github.io/MOOC_AnaDo/index.html https://husson.github.io/MOOC.html#PCAcourse

URL: `http://factominer.free.fr/index_fr.html`

Contact: Julie Josse

Partner: AGROCAMPUS

7.2 New platforms

Causal inference taskview: to list and organize all the R packages on causal inference

Participants: Julie Josse.

R-miss-tastic: a platform to gather and create resources on missing data, aimed at researchers and students who often do not have lecture on missing values. It includes bibliography, courses, tutorials, implementations, pipelines of analysis in R and Python, etc.

Participants: Julie Josse, Krystyna Grzesiak, Christophe Muller.

The Hitchhiker's Guide to Attacks on Output Privacy: Hosted by OpenDP, this website is a living database of research on attacks that infer sensitive information from statistical outputs. To help researchers and practitioners find relevant literature and understand privacy risks, it classifies papers by key dimensions. These include the attacker's objective (e.g., membership inference, reconstruction), the data modality (text, vision, tabular), and the type of statistical release.

Participants: Ioan Tudor Cebere.

8 New results

8.1 Treatment effect estimation

Results: Causal Meta Analysis: Rethinking the Foundations of Evidence-Based Medicine [59]

Participants: Julie Josse, Clement Berenfeld, Ahmed Boughdiri, Remi Khellaf, Aurélien Bellet.

Meta-analysis, by synthesizing effect estimates from multiple studies conducted in diverse settings, stands at the top of the evidence hierarchy in clinical research. Yet, conventional approaches based on fixed- or random-effects models lack a causal framework, which may limit their interpretability and utility for public policy. Incorporating causal inference reframes meta-analysis as the estimation of well-defined causal effects on clearly specified populations, enabling a principled approach to handling study heterogeneity. We show that classical meta-analysis estimators have a clear causal interpretation when effects are measured as risk differences. However, this breaks down for nonlinear measures like the risk ratio and odds ratio. To address this, we introduce novel causal aggregation formulas that remain compatible with standard meta-analysis practices and do not require access to individual-level data. To evaluate real-world impact, we apply both classical and causal meta-analysis methods to 500 published meta-analyses. While the conclusions often align, notable discrepancies emerge, revealing cases where conventional methods may suggest a treatment is beneficial when, under a causal lens, it is in fact harmful.

Results: A Unified Framework for the Transportability of Population-Level Causal Measures [62]

Participants: Julie Josse, Clement Berenfeld, Ahmed Boughdiri.

Generalization methods offer a powerful solution to one of the key drawbacks of randomized controlled trials (RCTs): their limited representativeness. By enabling the transport of treatment effect estimates to target populations subject to distributional shifts, these methods are increasingly recognized as the future of meta-analysis, the current gold standard in evidence-based medicine. Yet most existing approaches focus on the risk difference, overlooking the diverse range of causal measures routinely reported in clinical research. Reporting multiple effect measures-both absolute (e.g., risk difference, number needed to treat) and relative (e.g., risk ratio, odds ratio)-is essential to ensure clinical relevance, policy utility, and interpretability across contexts. To address this gap, we propose a unified framework for transporting a broad class of first-moment population causal effect measures under covariate shift. We provide identification results under two conditional exchangeability assumptions, derive both classical and semiparametric estimators, and evaluate their performance through theoretical analysis, simulations, and real-world applications. Our analysis shows the specificity of different causal measures and thus the interest of studying them all: for instance, two common approaches (one-step, estimating equation) lead to similar estimators for the risk difference but to two distinct estimators for the odds ratio.

Results: Rethinking the Win Ratio: A Causal Framework for Hierarchical Outcome Analysis. [63]

Participants: Julie Josse, Mathieu Even.

For hierarchical multivarariate outcomes, the FDA recommends the Win Ratio and Generalized Pairwise Comparisons approaches [84, 75]. However, as far as we know, these empirical methods lack causal or statistical foundations to justify their broader use in recent studies. To address this gap, we establish causal foundations for hierarchical comparison methods. We define related causal effect measures, and highlight that depending on the methodology used to compute Win Ratio, the causal estimand targeted can be different, as proved by our consistency results, which may then lead to reversed and incorrect treatment recommendations in heterogeneous populations, as we illustrate through striking examples. In order to compensate for this fallacy, we introduce a novel, individual-level yet identifiable causal effect measure that better approximates the ideal, non-identifiable individual-level estimand. We prove that computing Win Ratio or Net Benefits using a Nearest Neighbor pairing approach between treated and controlled patients, an approach that can be seen as an extreme form of stratification, leads to estimating this new causal estimand measure. We extend our methods to observational settings via propensity weighting, distributional regression to address the curse of dimensionality, and a doubly robust framework. We prove the consistency of our methods, and the double robustness of our augmented estimator. These methods are straightforward to implement, making them accessible to practitioners. Finally, we validate our approach using synthetic data and on CRASH-3 [CRASH et al., 2019], a major clinical trial focused on assessing the effects of tranexamic acid in patients with traumatic brain injury.

8.2 Federated Learning

Results: Federated Causal Inference from Multi-Site Observational Data via Propensity Score Aggregation. [66]

Participants: Remi Khellaf, Aurélien Bellet, Julie Josse.

Causal inference typically assumes centralized access to individual-level data. Yet, in practice, data are often decentralized across multiple sites, making centralization infeasible due to privacy, logistical, or legal constraints. We address this problem by estimating the Average Treatment Effect (ATE) from decentralized observational data via a Federated Learning (FL) approach, allowing inference through the exchange of aggregate statistics rather than individual-level data. We propose a novel method to estimate propensity scores by computing a federated weighted average of local scores with Membership Weights (MW)—probabilities of site membership conditional on covariates—which can be flexibly estimated using parametric or non-parametric classification models. Unlike density ratio weights (DW) from the transportability and generalization literature, which either rely on strong modeling assumptions or cannot be implemented in FL, MW can be estimated using standard FL algorithms and are more robust, as they support flexible, non-parametric models—making them the preferred choice in multi-site settings with strict data-sharing constraints. The resulting propensity scores are used to construct Federated Inverse Propensity Weighting (Fed-IPW) and Augmented IPW (Fed-AIPW) estimators. Unlike meta-analysis methods, which fail when any site violates positivity, our approach leverages heterogeneity in treatment assignment across sites to improve overlap. We show that Fed-IPW and Fed-AIPW perform well under site-level heterogeneity in sample sizes, treatment mechanisms, and covariate distributions. Both theoretical analysis and experiments on simulated and real-world data highlight their advantages over meta-analysis and related methods.

Results: Generalization under Byzantine & Poisoning Attacks: Tight Stability Bounds in Robust Distributed Learning [61]

Participants: Thomas Boudou, Aurélien Bellet.

Robust distributed learning algorithms aim to maintain good performance in distributed and federated settings, even in the presence of misbehaving workers. Two primary threat models have been studied: Byzantine attacks, where misbehaving workers can send arbitrarily corrupted updates, and data poisoning attacks, where misbehavior is limited to manipulation of local training data. While prior work has shown comparable optimization error under both threat models, a fundamental question remains open: How do these threat models impact generalization? Empirical evidence suggests a gap between the two threat models, yet it remains unclear whether it is fundamental or merely an artifact of suboptimal attacks. In this work, we present the first theoretical investigation into this problem, formally showing that Byzantine attacks are intrinsically more harmful to generalization than data poisoning.

8.3 Learning with Privacy Guarantees

Results: Model Agnostic Differentially Private Causal Inference. [67]

Participants: Aurélien Bellet, Julie Josse, Christian Janos Lebeda, Mathieu Even.

Estimating causal effects from observational data is essential in fields such as medicine, economics and social sciences, where privacy concerns are paramount. We propose a general, model-agnostic framework for differentially private estimation of average treatment effects (ATE) that avoids strong structural assumptions on the data-generating process or the models used to estimate propensity scores and conditional outcomes. In contrast to prior work, which enforces differential privacy by directly privatizing these nuisance components and results in a privacy cost that scales with model complexity, our approach decouples nuisance estimation from privacy protection. This separation allows the use of flexible, state-of-the-art black-box models, while differential privacy is achieved by perturbing only predictions and aggregation steps within a fold-splitting scheme with ensemble techniques. We instantiate the framework for three classical estimators – the G-formula, inverse propensity weighting (IPW), and augmented IPW (AIPW) – and provide formal utility and privacy guarantees. Empirical results show that our methods maintain competitive performance under realistic privacy budgets. We further extend our framework to support meta-analysis of multiple private ATE estimates. Our results bridge a critical gap between causal inference and privacy-preserving data analysis.

Results: Privacy Amplification Through Synthetic Data: Insights from Linear Regression [53]

Participants: Aurélien Bellet.

Synthetic data inherits the differential privacy guarantees of the model used to generate it. Additionally, synthetic data may benefit from privacy amplification when the generative model is kept hidden. While empirical studies suggest this phenomenon, a rigorous theoretical understanding is still lacking. In this paper, we investigate this question through the well-understood framework of linear regression. First, we establish negative results showing that if an adversary controls the seed of the generative model, a single synthetic data point can leak as much information as releasing the model itself. Conversely, we show that when synthetic data is generated from random inputs, releasing a limited number of synthetic data points amplifies privacy beyond the model's inherent guarantees. We believe our findings in linear regression can serve as a foundation for deriving more general bounds in the future.

Results: Tighter Privacy Auditing of DP-SGD in the Hidden State Threat Model [46]

Participants: Ioan Tudor Cebere, Aurélien Bellet, Nicolas Papernot.

Machine learning models can be trained with formal privacy guarantees via differentially private optimizers such as differentially private stochastic gradient descent DP-SGD. In this work, we focus on a threat model where the adversary has access only to the final model, with no visibility into intermediate updates. In the literature, this hidden state threat model exhibits a significant gap between the lower bound from empirical privacy auditing and the theoretical upper bound provided by privacy accounting. To challenge this gap, we propose to audit this threat model with adversaries that craft a gradient sequence designed to maximize the privacy loss of the final model without relying on intermediate updates. Our experiments show that this approach consistently outperforms previous attempts at auditing the hidden state model. Furthermore, our results advance the understanding of achievable privacy guarantees within this threat model. Specifically, when the crafted gradient is inserted at every optimization step, we show that concealing the intermediate model updates in DP-SGD does not enhance the privacy guarantees. The situation is more complex when the crafted gradient is not inserted at every step: our auditing lower bound matches the privacy upper bound only for an adversarially-chosen loss landscape and a sufficiently large batch size. This suggests that existing privacy upper bounds can be improved in certain regimes.

8.4 Handling missing data

Results: When Pattern-by-Pattern Works: Theoretical and Empirical Insights for Logistic Models with Missing Values [65]

Participants: Christophe Muller, Julie Josse.

Predicting a response with partially missing inputs remains a challenging task even in parametric models, since parameter estimation in itself is not sufficient to predict on partially observed inputs. Several works study prediction in linear models. In this paper, we focus on logistic models, which present their own difficulties. From a theoretical perspective, we prove that a Pattern-by-Pattern strategy (PbP), which learns one logistic model per missingness pattern, accurately approximates Bayes probabilities in various missing data scenarios such as missing completely at random, missing at random, and missing not at random (MCAR, MAR and MNAR). Empirically, we thoroughly compare various methods (constant and iterative imputations, complete case analysis, PbP, and an EM algorithm) across classification, probability estimation, calibration, and parameter inference. Our analysis provides a comprehensive view on the logistic regression with missing values. It reveals that mean imputation can be used as baseline for low sample sizes, and improved performance is obtained via nonlinear multiple iterative imputation techniques with the labels (MICE.RF.Y). For large sample sizes, PbP is the best method for Gaussian mixtures, and we recommend MICE.RF.Y in presence of nonlinear features.

Results: Do we Need Dozens of Methods for Real World Missing Value Imputation? [65]

Participants: Krystyna Grzesiak, Christophe Muller, Julie Josse, Jeffrey Naef.

Missing values pose a persistent challenge in modern data science. Consequently, there is an ever-growing number of publications introducing new imputation methods in various fields. While many studies compare imputation approaches, they often focus on a limited subset of algorithms and evaluate performance primarily through pointwise metrics such as RMSE, which are not suitable to measure the preservation of the true data distribution. In this work, we provide a systematic benchmarking method based on the idea of treating imputation as a distributional prediction task. We consider a large number of algorithms and, for the first time, evaluate them not only on synthetic missing mechanisms, but also on real-world missingness scenarios, using the concept of Imputation Scores. Finally, while the focus of previous benchmark has often been on numerical data, we also consider mixed data sets in our study. The analysis overwhelmingly confirms the superiority of iterative imputation algorithms, especially the methods implemented in the mice R package.

8.5 Application domain

Results: Sodium Bicarbonate for Severe Metabolic Acidemia and Acute Kidney Injury [77]

Participants: Maxime Fosset, Nicolas Molinari.

The effect of sodium bicarbonate infusion on outcome in patients with severe metabolic acidemia and moderate to severe acute kidney injury is unknown. The objective is to determine whether sodium bicarbonate infusion is associated with day 90 all-cause mortality in patients with severe metabolic acidemia and moderate to severe acute kidney injury. Randomized, open-label, clinical trial conducted with 640 patients in 43 French intensive care units from October 6, 2019, to December 19, 2023, with 90-day follow-up. The last date of follow-up was June 17, 2024. Adults with severe metabolic acidemia (defined as pH ≤7.20) and moderate to severe acute kidney injury were enrolled. Intervention Patients were randomized 1:1 to receive either intravenous sodium bicarbonate infusion or no sodium bicarbonate to target an arterial pH of 7.30 or higher.

Main Outcomes and Measures: The primary outcome was day 90 all-cause mortality. Secondary outcomes included day 28 and day 180 all-cause mortality; use of organ support therapy, vasopressors, or invasive mechanical ventilation; intensive care unit and hospital length of stay; intensive care unit-acquired infections; fluid balance; day-7 Sequential [Sepsis-related] Organ Failure Assessment score (6 organ systems' function is evaluated and scored from 0 [no dysfunction] to 4 [failure]; total score ranges from 0 [normal] to 24 [maximum failure]); and major adverse kidney events on day 90. Results Among 640 randomly assigned patients, 627 were analyzed (313 in the control group and 314 in the bicarbonate group). The median age was 67 years (IQR, 59-74 years); 194 of 314 patients (62%) in the bicarbonate group and 185 of 313 controls (59%) were male. In the primary analysis, day 90 all-cause mortality was 195 of 314 patients (62.1%) in the bicarbonate group and 193 of 313 (61.7%) in the control group (absolute difference, 0.4; 95% CI, -7.2 to 8.0; P =.91). There was no evidence of a group effect on day 28 or day 180 all-cause mortality. Among 18 secondary outcomes, kidney replacement therapy was used in 109 of 314 (35%) bicarbonate group patients and 157 of 313 (50%) controls (absolute difference, -15.5; 95% CI, -23.1 to -7.8). No evidence of a group effect was found on other secondary outcomes, including adverse events. Conclusions and Relevance: For patients with severe metabolic acidemia and moderate to severe acute kidney injury, intravenous sodium bicarbonate did not affect mortality. Trial Registration ClinicalTrials.gov Identifier: NCT04010630

Results: Allergen Chip Challenge: a nationwide open database supporting allergy prediction algorithms [80]

Participants: Pascal Demoly.

Background: Allergen chip technologies are a powerful tool for simultaneous analysis of hundreds of allergens, generating a comprehensive sensitization landscape for precision medicine in allergy. This considerable amount of data requires extensive knowledge for translation into clinically relevant conclusion. Objective: To harness Machine Learning (ML) for allergen chip interpretation in daily practice, we set out to establish a nationwide, open database of allergen chip, demographic and clinical information and to submit it to an international crowdsourced ML competition to generate a predictive allergy classification algorithm. Methods: The project consortium defined 20 clinical variables and 5 demographic factors for retrospective collection in conjunction with allergen chip IgE data (2014-2023) from 11 French University Hospitals. The dataset was processed to tag confirmed allergy, grade of severity, and culprit allergen identification associated with allergen chip data and submitted to the data challenge. Results: Data were collected for 4,271 patients, yielding a dataset with over 700,000 specific IgE data points. Sensitization was present in 3579 patients (84%). Allergy was confirmed in 2,236 patients (53%) and excluded in 1,076 patients, the remaining 959 being missing outcome data (allergy diagnosis labels). The competition attracted 292 data scientists who submitted 3,135 algorithms. The highest F-scores ranged from 0.780 to 0.786. The database was subsequently made available as an open source. Conclusions: We present a nationwide open allergy database designed to enable the development of predictive algorithms. This scalable framework, integrating clinical data with ML techniques paves the way for data-driven allergen chip use and interpretation by allergists.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Participants: Julie Josse, Gaelle Dormion.

- Title: Policy learning for personalized medicine. Finding the optimal dose of hormone for ovarian stimulation

  Infertility affects 1 in 5 couples of childbearing age. The most common solution is to resort to In Vitro Fertilization. However, the first challenge is to determine the initial dose and duration of gonadotropin hormone administration to maximize the number of oocytes retrieved at the end of stimulation, under the constraint that estradiol levels must not be too high to avoid hyperstimulation. The second challenge is to determine the ideal day for ovulation induction, to maximize the number of oocytes retrieved, and this is done by looking at the biological results of each monitoring. To tackle these two challenges, we will leverage rich observational multi-centric and longitudinal data as well as techniques of causal inference. More precisely, we will consider methods for learning optimal treatment policies and in particular for establishing the appropriate dose and duration of treatment for each patient. One of the challenges will be to propose methods to manage missing data in this framework. We will also consider techniques of dynamic treatment regimes to enrich the analysis with monitoring data, especially regarding hormone levels.
- Company: Elixir
- Duration: Feb 2023 -

Participants: Julie Josse, Mathieu Even.

- Title: (Longitudinal) Causal Machine Learning with Multiple Outcomes

  Context: The current healthcare system often employs a 'one size fits all' strategy, standardizing drug dosages, frequencies, and administration methods for all adults. However, this generalized approach fails to consider essential physio-pathological differences, such as sex, age, ethnicity, or disease progression, which significantly influence the efficacy and safety of medical treatments. This issue is particularly important in the fields of neurology and psychiatry, where interindividual patient characteristics play a crucial role in clinical symptoms, disease progression, and response to treatment.

Objective: Theremia aims to address these challenges by developing algorithms that analyze the response to central nervous system targeted drug treatments based on comprehensive patient characteristics (including sex, age, ethnic origin, disease progression, and genotype) and detailed drug properties (chemical and biological aspects).

By applying causal machine learning techniques to large observational clinical datasets, Theremia seeks to uncover the underlying factors that influence drug efficacy and the occurrence of side effects. This complex analysis often encounters methodological challenges, such as handling incomplete data and managing the intricacies of observational data, areas in which PreMeDICaL has considerable expertise.

Project Overview: This two-year collaborative research project will focus on methodological advancements in developing causal machine learning algorithms using clinical data related to Parkinson's disease. The primary objective is to analyze the effects of treatments and associated side effects in specific patient groups. The project is divided into two main phases, corresponding to the two years of research: 1) Static Causal Machine Learning (CML) with Multiple Outcomes, 2) Transition to Longitudinal Data Analysis

- Company: Theremia Health
- Duration: Dec 2024 -

Participants: Julie Josse, Agathe Chabassier.

- Title: Causal effects with digital devices

The overarching objective of this thesis is to investigate the effects of interventions on complex data, with a particular focus on highly granular time-series information. This endeavor raises profound conceptual challenges, particularly in defining what constitutes an average treatment effect on a dynamic trajectory—a task for which the inherent complexity of the data precludes any singular or straightforward solution. To address this, the work will establish a rigorous theoretical framework capable of identifying and characterizing both average and conditional effects, followed by the development of tailored estimators to quantify these effects. The proposed methodology will then be validated against real-world industrial data, which introduce additional layers of complexity, including recruitment biases, missing data mechanisms, and other practical constraints. By bridging theoretical precision with applied relevance, this research aims to advance robust analytical approaches that can inform decision-making in settings where data intricacy demands innovative solutions.

- Company: Withings
- Duration: Oct 2025 -

Participants: Pascal Demoly.

- Participation to the Fondation TEZOS (Vigicard digital health card project) with the startup CodInsight
- Co-creation of the startup AdviceMedica (collective intelligence for solving complex cases in medicine)

Participants: Aurélien Bellet, Ghita Fassy El Fehri.

- Title: Differentially private Federated learning in the framework of Bayesian Networks with application to cosmetic research

  The objective of this PhD is to develop a federated-learning type approach for Bayesian networks with additional privacy protection of model parameters by combining differential privacy with federated learning. The thesis will review the state of the art in this field, define the methodology and develop the associated algorithms in Python to learn the structure and estimate the parameters of the Bayesian networks in the context of federated learning with differential privacy guarantees.
- Company: L'Oréal
- Duration: December 2024 - December 2027

10 Partnerships and cooperations

10.1 International research visitors

10.1.1 Visits of international scientists

Inria International Chair

Nicolas Papernot, from October 2025

Status Professor

Institution of origin: Toronto

Country: Canada

Dates: 03/10/2025-

Context of the visit: Collaboration on Privacy

Mobility program/type of mobility: Research visit

Other international visits to the team

Uri Shalit

Status Professor

Institution of origin: Tel Aviv University

Country: Israel

Dates: 03/10/2025-07/10/2025

Context of the visit: Team hackathon and scientific collaboration on policy learning

Mobility program/type of mobility: Research visit

Ali Shojaie

Status Professor

Institution of origin: University of Washington

Country: USA

Dates: 09/12/2025-12/12/2025

Context of the visit: Scientific discussions and team seminar to prepare a potential sabbatical

Mobility program/type of mobility: Research visit

Krystyna Grzesiak

Status PhD student

Institution of origin: University of Wrocław

Country: Poland

Dates: 1/11/2024-01/04/2025

Context of the visit: Research work on missing data imputation

Mobility program/type of mobility: Research stay

Emma Torrini

Status PhD student

Institution of origin: Università degli Studi di Firenze

Country: Italy

Dates: 1/09/2025- 31/12/2025

Context of the visit: Research work on survival analysis

Mobility program/type of mobility: Research stay

10.1.2 Visits to international teams

Research stays abroad

Ahmed Boughdiri

Visited institution: UC Berkeley

Country: USA

Dates: 24/11/25 - 26/11/25

Context of the visit: invited talk to present work on generalization in causal inference before attending Neurips

Mobility program/type of mobility: Research visit

10.2 National initiatives

10.2.1 PEPR Digital Health

The "PEPR Santé Numérique", launched in June 2023 as part of the Plan Innovation Santé 2030, is a major initiative in the "Digital Health" acceleration strategy with a program dedicated to stimulating scientific research in this field.

PreMeDICaL is involved in three projects that have been launched:

- SMATCH "Statistical and AI Methods for the Challenges of Modern Clinical Trials in Digital Health" - Julie Josse, Pascal Demoly Mathieu Even
  - New clinical trial methods and designs based on animal-to-human, research-based disease models,
  - Enriching clinical trials with multi-source, multi-dimensional ancillary data,
  - Next-generation designs for clinical evaluation of digital medical devices based on AI algorithms,
  - Regulation, feasibility and dissemination of clinical trials
- Digital Pharmacological Twins "Multi-scale and longitudinal data modeling in pharmacology: toward digital pharmacological twins" - Julie Josse Jeffrey NaefClement Berenfeld
- Secure, safe and fair machine learning for healthcare - Aurélien Bellet

10.2.2 PEPR Cybersecurity

PreMeDICaL is involved in project IPoP (Interdisciplinary Project on Privacy) - Aurélien Bellet. The objectives of this project are to study the threats on privacy that have been introduced by these new services, and to conceive theoretical and technical privacy-preserving solutions that are compatible with French and European regulations, that preserve the quality of experience of the users. These solutions will be deployed and assessed, both on the technological and legal sides, and on their societal acceptability. In order to achieve these objectives, we adopt an interdisciplinary approach, bringing together many diverse fields: computer science, technology, engineering, social sciences, economy and law.

The project's scientific program focuses on new forms of personal information collection, on the learning of Artificial Intelligence (AI) models that preserve the confidentiality of personal information used, on data anonymization techniques, on securing personal data management systems, on differential privacy, on personal data legal protection and compliance, and all the associated societal and ethical considerations. This unifying interdisciplinary research program brings together internationally recognized research teams (from universities, engineering schools and institutions) working on privacy, and the French Data Protection Authority (CNIL).

This holistic vision of the issues linked to personal data protection will, on one hand, let us propose solutions to the scientific and technological challenges and, on the other hand, help us confront these solutions in many different ways in the context of interdisciplinary collaborations, thus leading to recommendations and proposals in the field of regulations or legal frameworks. This comprehensive consideration of all the issues aims at encouraging the adoption and acceptability of the solutions proposed by all stakeholders, legislators, data controllers, data processors, solution designers, developers all the way to end-users.

10.2.3 Inria Challenge FedMalin

Aurélien Bellet leads FedMalin. FedMalin is a research project that spans 11 Inria research teams and aims to push Federated Learning (FL) research and concrete use-cases through a multidisciplinary consortium involving expertise in ML, distributed systems, privacy and security, networks, and medicine. We propose to address a number of challenges that arise when FL is deployed over the Internet, including privacy & fairness, energy consumption, personalization, and location/time dependencies. FedMalin will also contribute to the development of open-source tools for FL experimentation and real-world deployments, and use them for concrete applications in medicine and crowdsensing.

The FedMalin Inria Challenge is supported by Groupe La Poste, sponsor of the Inria Foundation.

10.3 PANAME Project

The PANAME project, to audit the privacy of AI models, has been launched by the CNIL - Commission Nationale de l'Informatique et des Libertés, the Pôle d'Expertise de la Régulation Numérique (PEReN), the IPoP project (led by Inria) and the ANSSI - Agence nationale de la sécurité des systèmes d'information.

The aim of this project will be to develop a software library available in whole or in part as open source, enabling the implementation of data extraction and/or re-identification attacks on AI models. This will enable the confidentiality of AI models to be tested and audited.

10.3.1 ANR JCJC PRIDE

Aurélien Bellet leads PRIDE, a JCJC ANR project on privacy-preserving decentralized machine learning. The goal of PRIDE is to develop theoretical and algorithmic tools that enable differentially-private ML methods operating on decentralized datasets, through three complementary objectives:

- Prove that decentralized learning protocols naturally amplify DP guarantees;
- Propose algorithms at the intersection of decentralized ML and secure multi-party computation;
- Design data-adaptive communication schemes to speed up the convergence on heterogeneous datasets.

10.4 Regional initiatives

UM Envi-H: initiative by the University of Montpellier.

Participants: Pascal Demoly.

The University of Montpellier, with the support of the Regional Health Agency of Occitanie, is launching an innovative project in the field of environmental health education: the creation of a Small Private Online Course (SPOC) dedicated to environmental health (EH) for primary care. This project is part of Axis 1, "Inform, educate, and train in environmental health," of the Regional Environmental Health Plan for Occitanie (PRSE4 Occitanie 2023-2028), which "aims to provide professionals, local authorities, and citizens with the knowledge and skills needed to act on environmental and health issues."

In collaboration with the Hérault Primary Health Insurance Fund and the University Department of General Medicine, this SPOC will be a hybrid training program combining online modules with in-person sessions.

Available from early 2026, it aims to develop EH skills for learners in both continuing and initial education. It is primarily intended for coordinators of coordinated healthcare structures (Territorial Professional Health Communities - CPTS / Multidisciplinary Health Centers - MSP), as well as for students in related fields.

This program will focus on enhancing the EH competencies of participants through a hybrid format combining online and in-person learning.

ComexIA Health Occitanie:

Participants: Nicolas Molinari, Clement Berenfeld.

Members of the steering committee for the Occitanie region's key challenge "AI for health": preparation of the call for proposals (12 co-financed PhD positions), selection of applications, dossier follow-up, and management of a 1.2M Euros budget.

Ethical Committee University Hospital

Participants: Nicolas Molinari, , Julie Josse.

We are involved in the CSE (comité, scientifique et éthique) for the CHU Montpellier.

Other local Projects the team is part of: Muse, eDOL, expos-UM, viA-UM, Fondation One Science Montpellier.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organization

- Aurélien Bellet co-organizes the Federated Learning One World webinar (1100+ registered attendees) since May 2020.
- Mathieu Even: organization of NeurIPS In Paris, a 2-day event that took place at Sorbonne Université on the 25th and 26th of November 2025, providing a local alternative to the NeurIPS conference.
- Linus Bleistein: took part in the organization of Eurips in Copenhagen, a one-week event with over 3500 participants.
- Clement Berenfeld and Linus Bleistein: organization of RAHM 2026, a one-day workshop taking place in PariSanté Campus on the 27th of January 2026, gathering experts in health machine learning.

11.1.2 Scientific events: selection

Member of the conference program committees

- Aurélien Bellet: Senior Area Chair for Artificial Intelligence and Statistics, AISTATS 2026
- Aurélien Bellet: Area Chair for Neural Information Processing Systems, NeurIPS 2025
- Aurélien Bellet: Area Chair for International Conference on Machine Learning, ICML 2025
- Julie Josse: Member of the scientific committee IMS International Conference on Data Science, Seville, France, December 2025.
- Julie Josse: Steering committee Eurocim european conferences on causal inference, 2025 -

Reviewer

- Aurélien Bellet : Workshop on Collaborative and Federated Agentic Workflows at ICML 2025
- Aurélien Bellet : Workshop "Will Synthetic Data Finally Solve the Data Access Problem?" at ICLR 2025
- Aurélien Bellet : Workshop on Security, Privacy and Information Theory at AsiaCCS 2025
- Aurélien Bellet : CAp 2025
- Aurélien Bellet : APVP 2025
- Mathieu Even reviewed for the conferences NeurIPS conference 2025 and AISTATS 2026.
- Christian Janos Lebeda reviewed for the conferences SaTML 2026 and AISTATS 2026 and the workshops TPDP 2025 and PPML @ EurIPS 2025.

- Ioan Tudor Cebere: Reviewer for the conferences ICLR 2026, SaTML 2026, ICML 2025 and the TPDP 2025 workshop.
- Mathieu Dagréou reviewed for the conferences ICML 2025, NeurIPS 2025, AISTATS 2026, and ICLR2026.
- Linus Bleistein: reviewer for AISTATS 2025, NeurIPS 2025, AISTATS 2026
- Rémi Khellaf: reviewed for ICLR2026.
- Clement Berenfeld reviewed for the conferences NeurIPS conference 2025.
- Julie Josse reviewed for the conferences NeurIPS conference 2025.

11.1.3 Journal

Member of the editorial boards

- Aurélien Bellet is Action Editor for Transactions of Machine Learning Research (TMLR)
- Julie Josse: Associate Editor of Foundations and Trends in Machine Learning

Reviewer - reviewing activities

- Christian Janos Lebeda reviewed a paper for the Journal of Privacy and Confidentiality.
- Jean-Baptiste Fermanian reviewed a paper for the Electronic Journal of Statistics.
- Mathieu Dagréou reviewed 4 papers for the Transaction of Machine Learning Research journal and one for the Journal of Machine Learning Research
- Clement Berenfeld reviewed for Bernoulli journal, Statistics and Probability Letters, ESAIM Probability and Statistics.

11.1.4 Invited talks

- Julie Josse: AI, Science and Society, IPP Paris
- Julie Josse: Data Science, Statistics and Visualisation 2025, July 2025, South Africa. (Online).
- Julie Josse: Digicore meeting 2025, European research network on oncology
- Julie Josse: Online causal inference seminar
- Julie Josse: LMU AI Keynote Series
- Julie Josse: National Institute for Health and Care Excellence (NICEUK)
- Julie Josse: Toulouse School Economics
- Julie Josse: Academie of science, May 2025.
- Julie Josse: Bernouilli Lab, Paris
- Julie Josse: CAUSALab at Harvard T.H. Chan Methods series
- Aurélien Bellet : Keynote speaker at the Privacy-Preserving Machine Learning workshop at EurIPS 2025
- Aurélien Bellet : Invited talk at the Trustworthy AI Symposium (Paris AI Action Summit)
- Aurélien Bellet : Invited talk at Autumn School on Recent Advances in Machine Learning

- Aurélien Bellet : Invited talk at Dagstuhl Seminar "PETs and AI: Privacy Washing and the Need for a PETs Evaluation Framework"
- Aurélien Bellet : Talk in the BIPID Team, UMR INSERN IAME
- Aurélien Bellet : Invited talk in Bureau de Biostatistique et d'Epidémiologie, ONCOSTAT, Gustave Roussy
- Mathieu Even: invited talk at IBS in Liège (19/05/2025-21/05/2025)
- Mathieu Even: invited talk at ICSDS in Sevilla
- Mathieu Even: talk in the PEPR SMATCH days 2025 (07/10/2025)
- Mathieu Even: invited talk at the Biostatistic seminar of CIRC (OMS, Lyon, 04/09/2025)
- Mathieu Even: invited talk at the Journées de Société d'Informatique de France to present PhD works.
- Ahmed Boughdiri: invited talk at UC Berkeley to present generalization work (24/11/25)
- Ahmed Boughdiri: invited talk at Soda INRIA to present work on meta-analysis (21/10/25)
- Ahmed Boughdiri: invited talk in the PEPR SMATCH days 2025 (07/10/2025)
- Christian Janos Lebeda: invited talk at University of Toronto (10/01/2025).
- Christian Janos Lebeda: invited talk at BARC - University of Copenhagen (16/04/2025).
- Christian Janos Lebeda: Invited talk at IPoP réunion plénière 2025 (10/10/2025).
- Ioan Tudor Cebere: Invited talk at Inria Lille (16/10/2025).
- Clement Berenfeld : Invited talk at IDESP seminar (IDESP, Montpellier, 02/10/2025)
- Clement Berenfeld : Invited talk at the Biostatistic seminar of CIRC (OMS, Lyon, 12/10/2025)
- Jean-Baptiste Fermanian: talk at New challenges in high-dimensional statistics, Marseille.
- Jean-Baptiste Fermanian: talk at VITE 2025, Montpellier.
- Jean-Baptiste Fermanian: poster at Neurips 2025, San Diego.
- Jean-Baptiste Fermanian: talk at UC Berkeley.
- Jean-Baptiste Fermanian: talk at Seminaire Probabilités et Statistiques, Université d'Angers.
- Jean-Baptiste Fermanian: talk at Team seminar MAGNET team, Inria Lille.
- Jean-Baptiste Fermanian: talk at MADSTAT Seminar, Toulouse School of Economics, Toulouse.
- Laura Fuentes: talk at “Journées de l'IA pour la santé" organized by ANITI and the Occitanie region, Montpellier.
- Laura Fuentes: talk at Seminar#10 IDESP, Montpellier.
- Rémi Khellaf: talk at SODA team seminar, October 2025, Saclay
- Rémi Khellaf: talk at Causal Data Science meeting, October 2025, online
- Rémi Khellaf: talk at Pacific Causal Inference Conference, July 2025, Beijing (online)
- Ghita Fassy El Fehri: talk at ISI 2025, The Hague (05/10/2025-09/10/2025)
- Ghita Fassy El Fehri: poster at PPML @ Eurips, Copenhagen (02/12/2025-07/12/2025)

11.1.5 Contributed Talks

- Rémi Khellaf: talk at SMATCH PEPR annual meeting, November 2025, Paris
- Rémi Khellaf: talk at Cap25, July 2025, Dijon
- Laura Fuentes: talk at JDS 2025, Marseille.
- Rémi Khellaf: talk at Statistics and Biostatistics (SNB), November 2025, Paris
- Rémi Khellaf: talk at FedMalin third meeting, November 2025, Montpellier
- Rémi Khellaf: talk at Journées des Statistiques, May 2025, Marseille
- Laura Fuentes: poster at EUROCIM 2025 in Ghent (09/04/2025-11/04/2025).
- Ahmed Boughdiri: contributed talk at SNB 2025 in Paris (08/10/2025-10/10/2025)
- Mathieu Even: contributed talk at EUROCIM 2025 in Ghent (09/04/2025-11/04/2025)
- Christian Janos Lebeda: 2 contributed talks at SOSA 2025 in New Orleans, USA (13/01/2025).
- Christian Janos Lebeda: contributed talk at FORC 2025 at Stanford University, USA (05/06/2025).

11.1.6 Leadership within the scientific community

- Julie Josse is elected as a member of the R foundation and of the R Foundation Conference Committee. She is in the board of the French R committee (organization for coordinating R conferences "Les rencontres R") and involved in a task Forwards force on behalf of the R Foundation with the aim of increasing the participation of women and under-represented groups in the STEM community (founding member in 2015).
- Charlotte Voinot: Treasurer of "Groupe Jeunes Statisticien.ne.s"
- Ioan Tudor Cebere : Privacy Attacks Workgroup Leadership for OpenDP.

11.1.7 Scientific expertise

1. Aurélien Bellet : Member of the CNIL-Inria Privacy Award committee
2. Aurélien Bellet : Member of the OECD Expert Group on AI, Data, and Privacy
3. Aurélien Bellet : ethics advisor for the European Strategy Forum on Research Infrastructures (ESFRI) project SLICES-PP
4. Aurélien Bellet : Member of the scientific and ethics committee of AICET
5. Clement Berenfeld: member of the expert committee of ANITI IA for health program
6. Julie Josse: Scientific council Cluster IA PostGenAI@Paris Sorbonne Université
7. Julie Josse: Evaluation of ERC, and projects for tenured Professor positions (Harvard)

11.1.8 Research administration

- Aurélien Bellet : member of the Operational Committee for the assessment of Legal and Ethical risks (COERLE).

11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

11.2.1 Supervision

PhD students

- Julie Josse and Aurélien Bellet : Supervision of Rémi Khellaf (grant Montpellier), September 2023 -
- Aurélien Bellet : Supervision of Ioan Tudor Cebere, October 2022 -
- Aurélien Bellet : Supervision of Clément Pierquin with Marc Tommasi, June 2023 -
- Aurélien Bellet : Supervision of Brahim Erraji with Catuscia Palamidessi and Michael Perrot, September 2023 -
- Aurélien Bellet : Supervision of Thomas Boudou with Batiste Le Bars, October 2024 -
- Aurélien Bellet : Supervision of Ghita Fassy El Fehri, December 2024 -
- Julie Josse: Supervision of Laura Fuentes Vincente (grant Montpellier) with Antoine Chambaz, November 2024 -
- Julie Josse: Supervision of Ahmed Boughdiri (grant Inria), September 2023 -
- Julie Josse: Supervision of Charlotte Voinot with Bernard Sebastien (grant Phd thesis Cifre Sanofi), April 2023 -
- Julie Josse and Nicolas Molinari: Supervision of the MD Maxime Fosset (grant Montpellier University, MUSE) with Boris Jung (MD), May 2022 -
- Julie Josse: Supervision of Agathe Chabassier (PhD Cifre Withings) with Erwan Scornet, October 2025 -
- Julie Josse: Supervision of Tess Breton with Antoine Chambaz and Genevieve Robin, October 2025 -

Postdocs

- Aurélien Bellet : Jean-Baptiste Fermanian, October 2025 -
- Aurélien Bellet : Linus Bleistein, Mar - September 2025
- Aurélien Bellet : Mathieu Dagréou, December 2024 -
- Aurélien Bellet : Christian Janos Lebeda, October 2024 -
- Julie Josse : Mathieu Even, October 2024 - October 2025.
- Julie Josse : Clement Berenfeld, March 2025 - October 2025.

11.2.2 Juries

Member of PhD/HDR committees

- Aurélien Bellet : Reviewer for the habilitation thesis (HDR) of Cédric Gouy-Paillier. May 2025.
- Aurélien Bellet : Opponent for the PhD of Dominik Fay, KTH (Sweden), June 2025
- Aurélien Bellet : Reviewer for the PhD of Sadegh Farhadkhani, EPFL (Switzerland), August 2025
- Julie Josse and Aurélien Bellet : PhD defense committee of Linus Bleistein, Université Paris Saclay, June 2025
- Aurélien Bellet : PhD defense committee of Alexandre Rio, Université Grenoble Alpes, June 2025

- Aurélien Bellet : PhD defense committee of Romain Chor, Université Gustave Eiffel, September 2025
- Julie Josse: PhD defense committee of Stella Dimitsaki under the supervision of Marie-Christine Jaulent
- Julie Josse: PhD defense committee of Axel Roques under the supervision of Nicolas Vayatis
- Julie Josse: PhD defense committee of Antoine Pitoy under the supervision of Solène Desm'ee and Hoai Thu Thai (SANOFI)
- Julie Josse: PhD defense committee of Hava Chaptoukaev under the supervision of Maria A. Zuluaga
- Julie Josse: HDR defense committee of Myriam Tami

Member of hiring committees

- Aurélien Bellet : Member of full professor recruiting committee - Université de Saint-Etienne.
- Aurélien Bellet : Member of assistant professor recruiting committee - Université de Montpellier.

11.2.3 Teaching

- Aurélien Bellet gave a guest lecture for the *Responsible Machine Learning* course of Master MVA (2.5 hours).
- Clement Berenfeld has taught for the *Fondamentaux du machine learning* course of Polytech Montpellier (4th year course, 9h).
- Julie Josse: Winter School on Causality and Explainable AI (6 hours)
- Julie Josse: AI+Science Summer School. University of Chicago-Paris (3 hours).
- Julie Josse: Spring School in Statistics and Machine Learning: Oxford-Imperial-Bocconi, 1 week (35 hours).

11.3 Popularization

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

- Aurélien Bellet : Participation to the book "Tout comprendre (ou presque) sur l'intelligence artificielle" (Understanding (Almost) Everything About AI), CNRS Editions [link].
- Aurélien Bellet : Interview about AI in participatory democracy in the French media DémocratieS [link].
- Aurélien Bellet : Interview for L'ECO [link]
- Mathieu Even was invited by Serge Abiteboul to the podcast "Parlez-moi d'IA" on Cause Commune to talk about his PhD Thesis on federated model learning and causal inference applied to the medical field (link).
- Julie Josse: blog binaire la recherche: L'intelligence artificielle et la fabrique du savoir clinique.

11.3.2 Participation in Live events

- Julie Josse: Vivatech, round table Healthcare for All? Balancing Innovation and Equity Using AI. video.
- Julie Josse: Medtech IAvideo
- Julie Josse: Presentation de l'équipe Premedical, Mairie de Montpellier. Lancement de de l'association IA Montpellier Méditerrané.

12 Scientific production

12.1 Major publications

[1] C. Berenfeld, A. Boughdiri, B. Colnet, W. A. C. van Amsterdam, A. Bellet, R. Khellaf, E. Scornet and J. Josse. *Causal Meta-Analysis: Rethinking the Foundations of Evidence-Based Medicine*. 26th May 2025. URL: https://inria.hal.science/hal-05086779.

[2] T. Cebere, A. Bellet and N. Papernot. 'Tighter Privacy Auditing of DP-SGD in the Hidden State Threat Model'. In: ICLR 2025 - 13th International Conference on Learning Representations. Singapore, Singapore, 2025. URL: https://inria.hal.science/hal-04863204.

[3] B. Colnet, J. Josse, G. Varoquaux and E. Scornet. *Risk ratio, odds ratio, risk difference... Which causal measure is easier to generalize?* 2023. URL: https://hal.science/hal-04369607.

[4] B. Colnet, I. Mayer, G. Chen, A. Dieng, R. Li, G. Varoquaux, J.-P. Vert, J. Josse and S. Yang. 'Causal inference methods for combining randomized trials and observational studies: a review'. In: *Statistical Science* (2024). URL: https://hal.science/hal-03008276.

[5] T. Gauss, A. James, C. Colas, N. Delhaye, M. Holleville, B. Bijok, M. Werner, A. Meyer, V. Ramonda, E. Cesareo, H. de Cherisey, S. Medjkoune, S. Salah, J.-P. Nadal, J.-D. Moyer, A. Vilotitch, P. Bouzat and J. Josse. 'Comparison of machine learning and human prediction to identify trauma patients in need of hemorrhage control resuscitation (ShockMatrix study): a prospective observational study'. In: *The Lancet Regional Health - Europe* 55 (Aug. 2025), p. 101356. DOI: 10.1016/j.lanepe.2025.101356. URL: https://hal.science/hal-05113088.

[6] J. Josse, N. Prost, E. Scornet and G. Varoquaux. *On the consistency of supervised learning with missing values*. 29th June 2020. URL: https://hal.science/hal-02024202.

[7] R. Khellaf, A. Bellet and J. Josse. 'Federated Causal Inference: Multi-Study ATE Estimation beyond Meta-Analysis'. In: AISTATS 2025 - 28th International Conference on Artificial Intelligence and Statistics. Mai Khao, Thailand, 2025. URL: https://hal.science/hal-04747923.

[8] M. Le Morvan, J. Josse, E. Scornet and G. Varoquaux. 'What's a good imputation to predict with missing values?' In: NeurIPS 2021 - 35th Conference on Neural Information Processing Systems. Virtual, France, 6th Dec. 2021. URL: https://hal.science/hal-03243931.

[9] I. Mayer, A. Sportisse, N. Tierney, N. Vialaneix and J. Josse. 'R-miss-tastic: a unified platform for missing values methods and workflows'. In: *The R Journal* (30th July 2022). URL: https://hal.archives-ouvertes.fr/hal-02879337.

[10] I. Mayer, E. Sverdrup, T. Gauss, J.-D. Moyer, S. Wager and J. Josse. 'Doubly robust treatment effect estimation with missing attributes'. In: *Annals of Applied Statistics* 14.3 (Sept. 2020), pp. 1409–1431. DOI: 10.1214/20-AOAS1356. URL: https://hal.science/hal-02879332.

[11] M. Zaffran, O. Féron, Y. Goude, A. Dieuleveut and J. Josse. 'Adaptive Conformal Predictions for Time Series'. In: ICML 2022 - International Conference on Machine Learning. Baltimore, United States, 17th July 2022. URL: https://hal.archives-ouvertes.fr/hal-03573934.

12.2 Publications of the year

International journals

[12] M. Akkari, R. Lopez, I. Jaussent, C. Vidal, S. Skinner, D. Jaffuel and N. Molinari. 'Is a change in mouth opening associated with improvements in quality of life in children with type 1 obstructive sleep apnoea after adenotonsillectomy? Protocol for the JawChild prospective cohort study'. In: *BMJ Open* 15.12 (14th Dec. 2025), e105324. DOI: 10.1136/bmjopen-2025-105324. URL: https://hal.science/hal-05418376.

[13] A. Bourdin, S. Casanova, E. Ahmed, B. Brouard, A. Boher, I. Vachier, C. Suehs, F. Pujot, E. Nogue, D. Galeazzi, F. Cardon, M. Volpato, L. Halimi, S. Skinner and N. Molinari. 'Larger quality of life gains with asthma support app: a randomised controlled trial'. In: *ERJ Open Research* (30th Oct. 2025), pp. 00361–2025. DOI: 10.1183/23120541.00361-2025. URL: https://hal.science/hal-05343022.

[14] A. Bourdin, N. Malafaye, J. Pissarra, E. Ahmed, E. Nogué, C. M. Suehs, I. Vachier and N. Molinari. 'Healthcare Resource Utilisation and Costs of Mepolizumab Initiation: A 5-Year National Cohort Analysis'. In: *Allergy* (2025). DOI: 10.1111/all.16501. URL: https://hal.science/hal-04950053. In press.

[15] Y. Cabon, S. Bommart, G. Marin, E. Ahmed, C. Suehs, I. Vachier, A. S. Gamez, N. Molinari and A. Bourdin. 'Uniformity in air trapping distribution is associated with the severity of airway hyper-responsiveness in asthma'. In: *European Radiology* (28th June 2025). DOI: 10.1007/s00330-025-11751-6. URL: https://hal.science/hal-05138233.

[16] B. Canaud, A. Davenport, M. Morena, M. Amico, N. Molinari and J.-P. Cristol. 'Continuous beta-2 microglobulin–based clearance highlights superiority of high-Dose HDF over high-flux HD in predicting outcomes'. In: *Scientific Reports* 15.1 (1st July 2025), p. 20421. DOI: 10.1038/s41598-025-07497-2. URL: https://hal.science/hal-05139263.

[17] M. Capdevila, A. de Jong, F. Belafia, A. Vonarb, J. Carr, N. Molinari, O. Choquet, X. Capdevila and S. Jaber. 'Ultrasound guided transcutaneous phrenic nerve stimulation in critically ill patients: a new method to evaluate diaphragmatic function'. In: *Anesthesiology* 142.3 (2025), pp. 522–531. DOI: 10.1097/ALN.0000000000005267. URL: https://hal.science/hal-04751240.

[18] M. Capdevila, J. Pensier, A. de Jong, B. Jung, J. Beghin, T. Laumon, Y. Aarab, L. Deffontis, T. Sfara, A. Cuny, J. Carr, N. Molinari, J.-Y. Le Guennec, F. Raynaud, S. Matecki, L. Brochard, A. Lacampagne and S. Jaber. 'Impact of under-assisted ventilation on diaphragm function and structure in a porcine model'. In: *Anesthesiology* (24th Jan. 2025). DOI: 10.1097/ALN.0000000000005390. URL: https://hal.science/hal-04911765.

[19] M. Conseil, S. Jaber, F. Galia, N. Molinari, G. Chanques, A. de Jong and M. Capdevila. 'Neurally adjusted ventilatory assist in critical care patients with and without obesity: a prospective randomized crossover study'. In: *Annals of Intensive Care* 15.1 (29th Aug. 2025), p. 128. DOI: 10.1186/s13613-025-01552-x. URL: https://hal.science/hal-05229333.

[20] C. Desbordes, V. Szabo, F. Greco, K. Chalard, C. Dargazanli, N. Molinari, E. Matzner, V. Macioce, J. Pissarra, G. Chanques and P. F. Perrigault. 'Influence of meteorological changes on the occurrence of cerebral aneurysm rupture in the Montpellier region: A retrospective study'. In: *Neurochirurgie* 71.2 (Mar. 2025), p. 101630. DOI: 10.1016/j.neuchi.2025.101630. URL: https://hal.science/hal-04882378.

[21] T. Gauss, P. Bouzat and J. Josse. 'From bedside to data and back to bedside: why do we need better data to perform trauma phenotyping and improve trauma care?' In: *Intensive Care Medicine* 51 (4th June 2025), pp. 1161–1163. DOI: 10.1007/s00134-025-07946-6. URL: https://hal.science/hal-05098316.

[22] T. Gauss, A. James, C. Colas, N. Delhaye, M. Holleville, B. Bijok, M. Werner, A. Meyer, V. Ramonda, E. Cesareo, H. de Cherisey, S. Medjkoune, S. Salah, J.-P. Nadal, J.-D. Moyer, A. Vilotitch, P. Bouzat and J. Josse. 'Comparison of machine learning and human prediction to identify trauma patients in need of hemorrhage control resuscitation (ShockMatrix study): a prospective observational study'. In: *The Lancet Regional Health - Europe* 55 (Aug. 2025), p. 101356. DOI: 10.1016/j.lanepe.2025.101356. URL: https://hal.science/hal-05113088.

[23] I.-M. Ghiordanescu, N. Molinari, A.-M. Forsea, R. Schrijvers, C. Hoarau and P. Demoly. 'Penicillin allergy management strategies relevant for clinical practice - a narrative review'. In: *Romanian Journal of Internal Medicine* (1st Jan. 2025). DOI: 10.2478/rjim-2024-0035. URL: https://hal.science/hal-04921907.

[24] T. Hannequin, A. Ughetto, J. Eliet, C. Blin, A. Canon, M. Saour, P. Gaudard, C. Vidal, N. Molinari, P. Colson and M. Mourad. 'End-tidal carbon dioxide changes induced by passive leg raising can predict fluid responsiveness in patients on veno-arterial extracorporeal membrane oxygenation: a prospective, interventional study'. In: *Annals of Intensive Care* 15.1 (20th Nov. 2025), p. 185. DOI: 10.1186/s13613-025-01604-2. URL: https://hal.science/hal-05374118.

[25] A. de Jong, H. Huguet, A. Bignon, F. Stephan, T. Godet, L. Collet, K. Asehnoune, A. Sylvestre, J. Sautillet, R. Blondonnet, M. Ferrandiere, Y. Launey, S. Lasocki, A. Rollé, P.-M. Fayolle, L. Muller, E. Pardo, N. Terzi, S. Ramin, B. Jung, E. Weiss, P.-S. Abback, P. Guerci, B. Sarton, H. Rozé, C. Dupuis, J. Cousson, M. Faucher, V. Lemiale, B. Cholley, G. Chanques, F. Belafia, C. Chauveton, E. Futier, E. Azoulay, N. Molinari and S. Jaber. 'Noninvasive ventilation in postoperative critically ill patients with morbid obesity: secondary analysis of the EXTUBOBESE multicentre randomised clinical trial'. In: *British Journal of Anaesthesia* (Oct. 2025). DOI: 10.1016/j.bja.2025.09.002. URL: https://hal.science/hal-05298562.

[26] L. Knabe, C. Tanaka, S. Tebeka, S. Neilson, C. Cauvin, N. Mercier, D. Cayron, M. Savelli, P. Deruelle, F. Fuchs, J. Pissarra, N. Molinari and R. Belzeaux. 'Prevalence of psychiatric disorders during pregnancy—a feasibility study at second trimester ultrasound in the general population (GROUP study): study protocol'. In: *BMJ Open* 15.2 (2025), e091923. DOI: 10.1136/bmjopen-2024-091923. URL: https://hal.science/hal-04978711.

[27] C. J. Lebeda and J. Tetek. 'Better Differentially Private Approximate Histograms and Heavy Hitters using the Misra-Gries Sketch'. In: *ACM Transactions on Database Systems* 50.3 (14th May 2025), pp. 1–26. DOI: 10.1145/3716375. URL: https://hal.science/hal-05240815.

[28] A. Marouf, S. Chaubard, R. Liévin, J.-M. Michot, N. Molinari, J. Rossignol, D. Cavalieri, C. Golfier, O. Allangba, L. Philippe, B. Tessoulin, A. Chauchet, B. Deau, L. Obéric, J. Vargaftig, A. Moignet, A. Clavert, R. Duléry, G. Brisou, S. Tardy, V. Fataccioli, R. Houot, R.-O. Casasnovas, C. Thieblemont, H. Ghesquieres, S. Carras, S. Le Gouill, G. Cartron, V. Ribrag, M. Cheminant, A. Marçais, F. Suarez, A. Marabelle, O. Tournilhac, G. Damaj, P. Gaulard, L. de Leval, F. Lemonnier, E. Bachy, S. Chevret, O. Hermine, L. Couronné and A. Jaccard. 'Efficacy of anti-PD1 therapy in extranodal NK/T cell lymphoma: A matched cohort analysis from the LYSA'. In: *HemaSphere* 9.1 (20th Jan. 2025), e70081. DOI: 10.1002/hem3.70081. URL: https://hal.science/hal-04926143.

[29] N. Molinari, N. Roche, A.-L. Vataire, S. Perrier, N. Pagès, A. Panes, A. Schmidt, A. Bourdin and L. Watier. 'Economic Burden of Uncontrolled COPD Under Triple Therapy: A Population-Based Study'. In: *International Journal of Chronic Obstructive Pulmonary Disease* 20 (Sept. 2025), pp. 3295–3304. DOI: 10.2147/COPD.S532553. URL: https://hal.science/hal-05286462.

[30] C. Monet, L. Piron, M. Pressac, N. Molinari, A. de Jong, B. Guiu and S. Jaber. 'Study protocol for the HONIVAH trial: a single-centre randomised study assessing high-flow oxygen therapy versus non-invasive ventilation on lung volumes and the upper airway in hypoxemic critically ill patients'. In: *BMJ Open* 15.9 (16th Sept. 2025), e106340. DOI: 10.1136/bmjopen-2025-106340. URL: https://hal.science/hal-05266532.

[31] L. Monnier, C. Colette, E. Renard, P.-Y. Benhamou, S. Aouinti, N. Molinari and D. Owens. 'Prevent hypoglycaemia when using automated insulin delivery systems in type 1 diabetes requires near normal glycaemic variability'. In: *Diabetes & Metabolism* 51.1 (Jan. 2025), p. 101589. DOI: 10.1016/j.diabet.2024.101589. URL: https://hal.science/hal-04926784.

[32] A. Naudet-Lasserre, J. Pensier, A. de Jong, M. Capdevila, S. Jaber and N. Molinari. 'Effect of noninvasive ventilation on tracheal reintubation among patients with hypoxemic respiratory failure following abdominal surgery: a bayesian post-hoc analysis of the NIVAS trial'. In: *Critical Care* (10th Dec. 2025). DOI: 10.1186/s13054-025-05795-9. URL: https://hal.science/hal-05411292.

[33] S. Ng, I. Derraz, G. Marnat, V. Jecko, J. Papaxanthos, G. Bellanger, C. Cognard, J.-C. Sol, T. Le Van, M. Berhouma, P. Thouant, T. Jacquesson, O. Eker, H. El Fertit, D. Liebskind, N. Molinari, N. Lonjon and V. Costalat. 'Embolization of the middle meningeal artery for chronic subdural hematoma: The OTEMACS multicenter, randomized, clinical trial protocol'. In: *Interventional Neuroradiology* (July 2025), p. 15910199251356098. DOI: 10.1177/15910199251356098. URL: https://hal.science/hal-05154946.

[34] J. Pensier, A. Naudet-Lasserre, C. Monet, M. Capdevila, Y. Aarab, I. Lakbar, G. Chanques, N. Molinari, A. de Jong and S. Jaber. 'Noninvasive respiratory support following extubation in critically ill adults with obesity: a systematic review and network meta-analysis'. In: *EClinicalMedicine* 79 (Jan. 2025), p. 103002. DOI: 10.1016/j.eclinm.2024.103002. URL: https://hal.science/hal-04847797.

[35] J. Pensier, M. Touaibia, M. A. Meerun, P. Hefteh, N. Bloncourt, A. Vonarb, A. Prades, B. Al Taweel, A. Debourdeau, L. Monino, G. Chanques, N. Molinari, A. de Jong, B. Guiu and S. Jaber. 'Morphological subphenotypes of acute pancreatitis-related acute respiratory distress syndrome'. In: *Critical Care* (16th Jan. 2026), Online ahead of print. DOI: 10.1186/s13054-026-05837-w. URL: https://hal.science/hal-05463296.

[36] V. Puel, F. Gagnadoux, I. Godard, G. Papaioannou, N. Zarqane, S. Ibrahim, N. Molinari, J.-L. Pepin and A. Sabil. 'Obstructive sleep apnoea syndrome and recurrent vasovagal syncope: Insights from a multicentre observational study'. In: *Archives of cardiovascular diseases* (Nov. 2025). DOI: 10.1016/j.acvd.2025.09.008. URL: https://hal.science/hal-05410043. In press.

[37] H. Susmann, A. Chambaz, J. Josse, P. Aegerter, M. Wargon and E. Bacry. 'Probabilistic Prediction of Arrivals and Hospitalizations in Emergency Departments in Île-de-France'. In: *International Journal of Medical Informatics* 195 (9th Mar. 2025), p. 105728. DOI: 10.1016/j.ijmedinf.2024.105728. URL: https://hal.science/hal-04539380.

[38] H. Susmann, A. Chambaz, J. Josse, P. Aegerter, M. Wargon and E. Bacry. 'Probabilistic prediction of arrivals and hospitalizations in emergency departments in Île-de-France'. In: *International Journal of Medical Informatics* 195 (Mar. 2025), p. 105728. DOI: 10.1016/j.ijmedinf.2024.105728. URL: https://hal.science/hal-05458755.

[39] C. Vidal, F. Bertelli, J.-P. Mallet, R. Gilson, J.-C. Borel, F. Gagnadoux, S. Skinner, A. Bourdin, N. Molinari and D. Jaffuel. 'Added value of technician intervention to improve mask management for apneic patients treated with long-term CPAP'. In: *Sleep and Breathing* 29.1 (28th Feb. 2025), p. 117. DOI: 10.1007/s11325-025-03279-2. URL: https://hal.science/hal-04972489.

[40] C. Vidal, J.-P. Mallet, S. Skinner, R. Gilson, O. Gaubert, A. Prigent, F. Gagnadoux, J.-C. Borel, A. Bourdin, N. Molinari and D. Jaffuel. 'Positive Airway Pressure-Related Aerophagia in Obstructive Sleep Apnea: Results from the InterfaceVent Real-Life Study'. In: *Journal of Clinical Medicine* 14.18 (11th Sept. 2025), p. 6424. DOI: 10.3390/jcm14186424. URL: https://hal.science/hal-05287086.

[41] C. Vidal, J.-P. Mallet, S. Skinner, R. Gilson, O. Gaubert, A. Prigent, F. Gagnadoux, A. Bourdin, N. Molinari and D. Jaffuel. 'Concerns arising from the calculation of the apnea-hypopnea index during CPAP-telemonitoring of patients with obstructive sleep apnea'. In: *Respiratory Research* 26.1 (12th July 2025), p. 244. DOI: 10.1186/s12931-025-03324-4. URL: https://hal.science/hal-05160145.

[42] C. Vidal, J.-P. Mallet, S. Skinner, R. Gilson, O. Gaubert, A. Prigent, F. Gagnadoux, A. Bourdin, N. Molinari and D. Jaffuel. 'The influence of atmospheric temperature on long-term CPAP-usage in patients with obstructive sleep apnea.' In: *Sleep Medicine* 128 (Jan. 2025), pp. 22–28. DOI: 10.1016/j.sleep.2025.01.023. URL: https://hal.science/hal-04912274.

[43] J.-B. Woillard, S. Benzekry, J. Josse, M. White-Koning, E. Chatelut, E. Comets, F. Lemaitre, B. Franck, M. Grégoire, F. Stanke-Labesque, S. Zohar, M. Ursino and C. Battail. 'Digital Pharmacological Twins: Bridging Multi-scale Modelling and Artificial Intelligence for Precision Medicine : the DIGPHAT consortium'. In: *Therapies* (2025). DOI: 10.1016/j.therap.2025.09.006. URL: https://hal.science/hal-05288607.

[44] P. Zhao, J. Josse and S. Yang. 'Efficient and Robust Transfer Learning of Optimal Individualized Treatment Regimes with Right-Censored Survival Data'. In: *Journal of Machine Learning Research* 26.48 (3rd Feb. 2025), pp. 1–54. URL: https://hal.science/hal-05244892.

International peer-reviewed conferences

[45] J. Abécassis, H. Zenati, S. Boumaïza, J. Josse and B. Thirion. 'CO11.2 - Explorer les fonctions cognitives dans UK Biobank avec une analyse de médiation causale'. In: EPICLIN 2025 - Conférence francophone d'EPIdémiologie CLINique. Vol. 73. Bordeaux, France, May 2025, p. 203025. DOI: 10.1016/j.jeph.2025.203025. URL: https://hal.science/hal-05241488.

[46] T. Cebere, A. Bellet and N. Papernot. 'Tighter Privacy Auditing of DP-SGD in the Hidden State Threat Model'. In: ICLR 2025 - 13th International Conference on Learning Representations. Singapore, Singapore, 2025. URL: https://inria.hal.science/hal-04863204 (cit. on p. 17).

[47] A. Filiot, N. Dop, O. Tchita, A. Riou, R. Dubois, T. Peeters, D. Valter, M. Scalbert, C. Saillard, G. Robin and A. Olivier. 'Distilling Foundation Models for Robust and Efficient Models in Digital Pathology'. In: *MICCAI 2025 - Open Access / Springerlink*. MICCAI 2025 - International Conference on Medical Image Computing and Computer Assisted Intervention. Vol. 15966. Lecture Notes in Computer Science. Daejeon, South Korea: Springer, 20th Sept. 2025, pp. 162–172. DOI: 10.1007/978-3-032-04981-0_16. URL: https://hal.science/hal-05404020.

[48] R. Khellaf, A. Bellet and J. Josse. 'Federated Causal Inference: Multi-Study ATE Estimation beyond Meta-Analysis'. In: AISTATS 2025 - 28th International Conference on Artificial Intelligence and Statistics. Mai Khao, Thailand, 2025. URL: https://hal.science/hal-04747923.

[49] C. J. Lebeda. 'Better Gaussian Mechanism using Correlated Noise'. In: *2025 Symposium on Simplicity in Algorithms (SOSA)*. SIAM Symposium on Simplicity in Algorithms (SOSA25). New Orleans (Louisiana), United States: Society for Industrial and Applied Mathematics, 15th Jan. 2025. DOI: 10.1137/1.9781611978315.9. URL: https://inria.hal.science/hal-05238002.

[50] C. J. Lebeda, M. Regehr, G. Kamath and T. Steinke. 'Avoiding Pitfalls for Privacy Accounting of Subsampled Mechanisms Under Composition'. In: *IEEE Xplore*. SaTML 2025 - IEEE Conference on Secure and Trustworthy Machine Learning. 2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML). Copenhagen, Denmark: IEEE, 22nd May 2025, pp. 996–1006. DOI: 10.1109/SaTML64287.2025.00060. URL: https://hal.science/hal-05240803.

[51] C. J. Lebeda and L. Retschmeier. 'The Correlated Gaussian Sparse Histogram Mechanism'. In: *6th Symposium on Foundations of Responsible Computing (FORC 2025)*. Symposium on Foundations of Responsible Computing (FORC 2025). Standford, United States, 3rd June 2025. DOI: 10.4230/LIPIcs.FORC.2025.23. URL: https://inria.hal.science/hal-04895301.

[52] C. J. Lebeda and J. Tětek. 'Testing Identity of Distributions under Kolmogorov Distance in Poly-logarithmic Space'. In: *2025 Symposium on Simplicity in Algorithms (SOSA)*. SIAM Symposium on Simplicity in Algorithms (SOSA25). New orleans, USA, United States, 15th Jan. 2025. DOI: 10.1137/1.9781611978315.4. URL: https://inria.hal.science/hal-04895297.

[53] C. Pierquin, A. Bellet, M. Tommasi and M. Boussard. 'Privacy Amplification Through Synthetic Data: Insights from Linear Regression'. In: ICML 2025 - 42nd International Conference on Machine Learning. Vancouver, Canada, 2025. DOI: 10.48550/arXiv.2506.05101. URL: https://hal.science/hal-05237878 (cit. on p. 16).

[54] A. S. Shamsabadi, P. Snyder, R. Giles, A. Bellet and H. Haddadi. 'Nebula: Efficient, Private and Accurate Histogram Estimation'. In: CCS 2025 - 32nd ACM Conference on Computer and Communications Security. Tapei (Taiwan), Taiwan, 2025. URL: https://inria.hal.science/hal-04863194.

[55] H. Zenati, J. Abécassis, J. Josse and B. Thirion. 'Double Debiased Machine Learning for Mediation Analysis with Continuous Treatments'. In: *Proceedings of Machine Learning Research*. AISTATS - 28th International Conference on Artificial Intelligence and Statistics. Vol. PMLR-. Mai Khao, Thailand, 3rd May 2025. URL: https://hal.science/hal-04982992.

Reports & preprints

[56] J. Abécassis, H. Zenati, S. Boumaïza, J. Josse and B. Thirion. *Causal mediation analysis with one or multiple mediators: a comparative study*. 7th May 2025. URL: https://hal.science/hal-05060162.

[57] M.-F. Beclin, G. Marchello, J.-B. Woillard and J. Josse. *Super learner for tabular synthetic data generation*. 26th Sept. 2025. URL: https://hal.science/hal-05285855.

[58] A. Bellet, E. Cyffers, D. Frey, R. Gaudel, D. Lerévérend and F. Taïani. *Unified Privacy Guarantees for Decentralized Learning via Matrix Factorization*. 20th Oct. 2025. URL: https://hal.science/hal-05462250.

[59] C. Berenfeld, A. Boughdiri, B. Colnet, W. A. C. van Amsterdam, A. Bellet, R. Khellaf, E. Scornet and J. Josse. *Causal Meta-Analysis: Rethinking the Foundations of Evidence-Based Medicine*. 26th May 2025. URL: https://inria.hal.science/hal-05086779 (cit. on p. 14).

[60] L. Bleistein, A. Bellet and J. Josse. *Optimal Transport with Heterogeneously Missing Data*. 22nd May 2025. URL: https://hal.science/hal-05245177.

[61] T. Boudou, B. L. Bars, N. Gupta and A. Bellet. *Generalization under Byzantine & Poisoning Attacks: Tight Stability Bounds in Robust Distributed Learning*. 2025. URL: https://hal.science/hal-05245060 (cit. on p. 16).

[62] A. Boughdiri, C. Berenfeld, J. Josse and E. Scornet. *A Unified Framework for the Transportability of Population-Level Causal Measures*. 16th May 2025. URL: https://hal.science/hal-05070952 (cit. on p. 14).

[63] M. Even and J. Josse. *Rethinking the Win Ratio: A Causal Framework for Hierarchical Outcome Analysis*. 24th Apr. 2025. URL: https://hal.science/hal-05244324 (cit. on p. 15).

[64] G. Fassy El Fehri, A. Bellet and P. Bastien. *Differentially Private and Federated Structure Learning in Bayesian Networks*. 1st Dec. 2025. URL: https://inria.hal.science/hal-05462153.

[65] K. Grzesiak, C. Muller, J. Josse and J. Näf. *Do we Need Dozens of Methods for Real World Missing Value Imputation?* 6th Nov. 2025. URL: https://hal.science/hal-05366545 (cit. on p. 17).

[66] R. Khellaf, A. Bellet and J. Josse. *Federated Causal Inference from Multi-Site Observational Data via Propensity Score Aggregation*. 30th Sept. 2025. URL: https://hal.science/hal-05081887 (cit. on p. 15).

[67] C. J. Lebeda, M. Even, A. Bellet and J. Josse. *Model Agnostic Differentially Private Causal Inference*. 27th May 2025. URL: https://hal.science/hal-05241348 (cit. on p. 16).

[68] C. J. Lebeda, A. Nikolov and H. Tang. *Weighted Fourier Factorizations: Optimal Gaussian Noise for Differentially Private Marginal and Product Queries*. 25th Dec. 2025. URL: https://inria.hal.science/hal-05436327.

[69] C. Muller, E. Scornet and J. Josse. *When Pattern-by-Pattern Works: Theoretical and Empirical Insights for Logistic Models with Missing Values*. 11th Sept. 2025. URL: https://hal.science/hal-05150753.

[70] J. Näf, E. Scornet and J. Josse. *What Is a Good Imputation Under MAR Missingness?* 20th Mar. 2025. URL: https://hal.science/hal-04521894.

[71] G. Ninot, E. Descamps, G. Achalid, S. Abad, F. Berna, C. Belhomme, F. Carbonnel, P. M. Carrieri, P. Dargent-Molina, F. Fiteni, A.-M. Foucaut, A. Guyon, A. Legout, B. Lognos, N. Molinari, J. Nizard, M. Nogues, F. Paille, P. Poisbeau, L. Rochaix and B. Falissard. *The NPIS Model: A Standardized, Consensus-Based Framework for Evaluating Non-Pharmacological Interventions*. 5th Apr. 2025. DOI: 10.1101/2025.04.04.25325250. URL: https://hal.science/hal-05023209.

[72] V. Roca, M. Tommasi, P. Andrey, A. Bellet, M. Schirmer, H. Henon, L. Puy, J. Ramon, G. Kuchcinski, M. Bretzner and R. Lopes. *Federated Learning for MRI-based BrainAGE: a multicenter study on post-stroke functional outcome prediction*. 2025. DOI: 10.48550/arXiv.2506.15626. URL: https://inria.hal.science/hal-05244644.

[73] M. Yaghini, T. Cebere, M. Menart, A. Bellet and N. Papernot. *Private Rate-Constrained Optimization with Applications to Fair Learning*. 13th June 2025. URL: https://hal.science/hal-05237961.

Other scientific publications

[74] T. Fehr Delude, T. Gauss, A. Kalimouttou, J. Josse, J. Grèze, P. Bouzat and B. Lemasson. 'FASTDIAG-TC: Outcome Prediction for Traumatic Brain Injuries: Multimodal AI Combining Clinical Data and CT Scans'. In: IABM 2025 - 3e édition du Colloque Français d'Intelligence Artificielle en Imagerie Biomédicale. Nice, France, 17th Mar. 2025. URL: https://hal.science/hal-05391250.

12.3 Cited publications

[75] M. Buyse. 'Generalized pairwise comparisons of prioritized outcomes in the two-sample problem'. In: *Statistics in Medicine* 29.30 (Dec. 2010), pp. 3245–3257 (cit. on p. 15).

[76] B. Colnet, I. Mayer, G. Chen, A. Dieng, R. Li, G. Varoquaux, J.-P. Vert, J. Josse and S. Yang. 'Causal inference methods for combining randomized trials and observational studies: a review'. In: *Statistical Science* (2024). URL: https://hal.science/hal-03008276 (cit. on p. 8).

[77] B. Jung, M. Jabaudon, A. de Jong, L. Bitker, J. Audard, K. Klouche, B. Sarton, C. Guitton, S. Lasocki, B. Rieu et al. 'Sodium Bicarbonate for Severe Metabolic Acidemia and Acute Kidney Injury'. In: *JAMA Cardiology* (Oct. 2025). DOI: 10.1001/jama.2025.20231. URL: https://hal.science/hal-05338215 (cit. on p. 18).

[78] P. Kairouz, H. B. McMahan, B. Avent, **Bellet, A.**, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings et al. 'Advances and Open Problems in Federated Learning'. In: *Foundations and Trends® in Machine Learning* 14.1–2 (2021), pp. 1–210 (cit. on p. 8).

[79] J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani and L. Wasserman. 'Distribution-Free Predictive Inference for Regression'. In: *Journal of the American Statistical Association* 113.523 (2018), pp. 1094–1111 (cit. on p. 9).

[80] G. Martinroche, A. Guemari, P. A. Apoil, I. Annesi-Maesano, E. Fromentin, L. Guilleminault, D. Caimmi, C. Klingebiel, C. Beauvillain, A. Didier, J. Corriger, P. Demoly, J. Vitte and J. Goret. 'Allergen Chip Challenge: a nationwide open database supporting allergy prediction algorithms'. In: *Journal of Allergy and Clinical Immunology* (Sept. 2025). DOI: 10.1016/j.jaci.2025.09.017. URL: https://univ-reims.hal.science/hal-05291392 (cit. on p. 18).

[81] L. Pahus, D. Jaffuel, I. Vachier, A. Bourdin, C. M. Suehs, N. Molinari and P. Chanez. 'Randomised controlled trials in severe asthma: selection by phenotype or stereotype'. In: *European Respiratory Journal* 53.2 (2019) (cit. on p. 7).

[82] B. Paige, J. Bell, **Bellet, A.**, A. Gascón and D. Ezer. 'Reconstructing Genotypes in Private Genomic Databases from Genetic Risk Scores'. In: *Journal of Computational Biology* 28.5 (2021), pp. 435–451 (cit. on p. 9).

[83] H. Papadopoulos, K. Proedrou, V. Vovk and A. Gammerman. 'Inductive Confidence Machines for Regression'. In: *Machine Learning: ECML 2002*. Springer, 2002, pp. 345–356 (cit. on p. 9).

[84] S. J. Pocock, C. A. Ariti, T. J. Collier and D. Wang. 'The win ratio: a new approach to the analysis of composite endpoints in clinical trials based on clinical priorities'. In: *European Heart Journal* 33.2 (Sept. 2011), pp. 176–182 (cit. on p. 15).

[85] Y. Romano, E. Patterson and E. Candès. 'Conformalized Quantile Regression'. en. In: *Advances in Neural Information Processing Systems*. Vol. 32. 2019. URL: https://papers.nips.cc/paper/2019/hash/5103c3584b063c431bd1268e9b5e76fb-Abstract.html (cit. on p. 9).

[86] A. D. Selbst, D. Boyd, S. A. Friedler, S. Venkatasubramanian and J. Vertesi. 'Fairness and Abstraction in Sociotechnical Systems'. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 2019, pp. 59–68 (cit. on p. 9).

[87] R. Shokri, M. Stronati, C. Song and V. Shmatikov. 'Membership Inference Attacks Against Machine Learning Models'. In: *IEEE Symposium on Security and Privacy*. 2017 (cit. on p. 9).

[88] V. Vovk, A. Gammerman and G. Shafer. *Algorithmic Learning in a Random World*. en. Springer US, 2005 (cit. on p. 9).

[89] M. B. Zafar, I. Valera, M. Gomez-Rodriguez and K. P. Gummadi. ‘Fairness Constraints: A Flexible Approach for Fair Classification’. In: *Journal of Machine Learning Research* 20.75 (2019), pp. 1–42 (cit. on p. 10).

[90] M. Zaffran, A. Dieuleveut, J. Josse and Y. Romano. ‘Conformal Prediction with Missing Values’. In: *Proceedings of Machine Learning Research*. Vol. 202. ICML 2023 - 40th International Conference on Machine Learning. Honolulu (Hawai), United States, July 2023, p. 40578 (cit. on p. 9).