

RESEARCH CENTRE

**Inria Branch at the University of  
Montpellier**

IN PARTNERSHIP WITH:

INSERM, Université de Montpellier

2024

ACTIVITY REPORT

Project-Team

PREMEDICAL

**Precision Medicine by Data Integration  
and Causal Learning**

IN COLLABORATION WITH: Institut Desbrest d'Épidémiologie et de Santé  
Publique (IDESP)

**DOMAIN**

**Digital Health, Biology and Earth**

**THEME**

**Computational Neuroscience and  
Medicine**

*Inria*

# Contents

<b>Project-Team PREMEDICAL</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>2</b>
<b>2 Overall objectives</b>	<b>3</b>
<b>3 Research program</b>	<b>4</b>
3.1 Research Axis 1: Personalized medicine by optimal prescription of treatment . . . . .	4
3.2 Research axis 2: Personalized medicine by integration of different data sources . . . . .	5
3.3 Research Axis 3: Personalized medicine with privacy and fairness guarantees . . . . .	6
<b>4 Application domains</b>	<b>7</b>
<b>5 Social and environmental responsibility</b>	<b>7</b>
5.1 Impact of research results . . . . .	7
<b>6 Highlights of the year</b>	<b>8</b>
6.1 Awards . . . . .	8
6.2 PhD defenses . . . . .	8
6.3 Other . . . . .	8
<b>7 New software, platforms, open data</b>	<b>9</b>
7.1 New software . . . . .	9
7.1.1 declearn . . . . .	9
7.2 New platforms . . . . .	9
<b>8 New results</b>	<b>10</b>
8.1 Treatment effect estimation . . . . .	10
8.2 Handling missing values . . . . .	11
8.3 Learning with privacy guarantees . . . . .	12
8.4 Federated learning . . . . .	15
8.5 Fair machine learning . . . . .	16
8.6 Uncertainty quantification . . . . .	16
8.7 Application domain: allergies, ICU care . . . . .	17
<b>9 Bilateral contracts and grants with industry</b>	<b>18</b>
9.1 Bilateral contracts with industry . . . . .	18
<b>10 Partnerships and cooperations</b>	<b>20</b>
10.1 International research visitors . . . . .	20
10.1.1 Visits of international scientists . . . . .	20
10.1.2 Visits to international teams . . . . .	20
10.2 National initiatives . . . . .	21
10.2.1 PEPR Digital Health . . . . .	21
10.2.2 PEPR Cybersecurity . . . . .	21
10.2.3 Inria Challenge FedMalin . . . . .	21
10.2.4 ANR JCJC PRIDE . . . . .	22
10.2.5 Allergen-Chip-Challenge . . . . .	22
10.2.6 Grant from the National Interministerial Road Safety Observatory . . . . .	22
10.2.7 Grant from PHRC . . . . .	22
10.2.8 Grant from Institut Exposum Doctoral Nexus . . . . .	22
10.2.9 Grant from Directorate General for Healthcare Services (DGOS) . . . . .	22
10.3 Regional initiatives . . . . .	22

<b>11 Dissemination</b>	<b>23</b>
11.1 Promoting scientific activities	23
11.1.1 Scientific events: organisation	23
11.1.2 Scientific events: selection	23
11.1.3 Journal	24
11.1.4 Invited talks	24
11.1.5 Leadership within the scientific community	25
11.1.6 Scientific expertise	26
11.1.7 Research administration	26
11.2 Teaching - Supervision - Juries	26
11.2.1 Teaching	26
11.2.2 Supervision	26
11.2.3 Juries	28
11.3 Popularization	29
11.3.1 Specific official responsibilities in science outreach structures	29
11.3.2 Productions (articles, videos, podcasts, serious games, ...)	29
11.3.3 Participation in Live events	29
<b>12 Scientific production</b>	<b>29</b>
12.1 Major publications	29
12.2 Publications of the year	30
12.3 Cited publications	33

## **Project-Team PREMEDICAL**

*Creation of the Project-Team: 2022 June 01*

### **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. – Health
- B2.2. – Physiology and diseases
- B2.3. – Epidemiology

# 1 Team members, visitors, external collaborators

## Research Scientists

- Julie Josse [Team leader, INRIA, Senior Researcher, from Mar 2024, HDR]
- Aurélien Bellet [INRIA, Senior Researcher, HDR]

## Faculty Members

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

## Post-Doctoral Fellows

- Mathieu Dagereu [INRIA, Post-Doctoral Fellow, from Dec 2024]
- Mathieu Even [INRIA, Post-Doctoral Fellow, from Oct 2024]
- Christian Janos Lebeda [INRIA, Post-Doctoral Fellow, from Oct 2024]
- Giulia Marchello [UNIV MONTPELLIER, Post-Doctoral Fellow, from Feb 2024 until Sep 2024]
- Jeffrey Naef [INRIA, Post-Doctoral Fellow, from Feb 2024]
- Jeffrey Naef [UNIV MONTPELLIER, until Jan 2024]

## PhD Students

- Marie Felicia Beclin [UNIV MONTPELLIER, ATER, from Oct 2024]
- Marie Felicia Beclin [UNIV MONTPELLIER, until Sep 2024]
- Thomas Boudou [INRIA, from Oct 2024]
- Ahmed Boughdiri [INRIA]
- Ioan Tudor Cebere [INRIA, from Sep 2024]
- Ghita Fassy El Fehri [INRIA, from Dec 2024]
- Maxime Fosset [UNIV MONTPELLIER]
- Laura Fuentes Vicente [UNIV MONTPELLIER, from Oct 2024]
- Remi Khellaf [UNIV MONTPELLIER]
- Charlotte Voinot [SANOFI, CIFRE]
- Margaux Zaffran [INRIA, until Jun 2024]
- Pan Zhao [UNIV MONTPELLIER, until Sep 2024]

## Technical Staff

- Christophe Muller [INRIA, Engineer, from Oct 2024]

## Interns and Apprentices

- Pauline Bian [ENSAE, Intern, until Mar 2024]
- Helene Bonneau-Chloup [UNIV CAMBRIDGE, Intern, until Mar 2024]
- Laura Fuentes Vicente [INRIA, Intern, from Apr 2024 until Sep 2024]

## Administrative Assistant

- Claire-Marine Parodi [INRIA]

## Visiting Scientists

- Clement Berenfeld [UNIV POTSDAM, from Sep 2024 until Oct 2024]
- Charif El Gataa [Univ Torino, from Oct 2024]
- Krystyna Grzesiak [UNIV WROCLAW, from Nov 2024]

## External Collaborators

- Helene Bonneau-Chloup [ELIXIR HEALTH, from Apr 2024]
- Gaelle Dormion [ELIXIR HEALTH, from Sep 2024]
- Pierre Lafaye De Micheaux [Univ New Souht Wales, until Jan 2024]

## 2 Overall objectives

The objective of the PreMeDICAL team (**P**recision **M**edicine by **D**ata Integration and **C**ausal Learning) is to develop the next generation of methods/algorithms to extract knowledge from health data and improve patient care. More specifically, the aim 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 by optimal prescription of treatment. We will develop causal inference techniques for **(dynamic) policy learning** (allocating the best treatment for each person at the right time), that handle missing values and leverage both RCTs and observational data. Using both data sources allow to **better design future RCTs** or to **launch a drug without running RCTs** and **in the longer term to rethink the evidence needed to bring treatments to the market** and to do so more quickly.
2. Personalized medicine by integration of different data sources. We will build predictive models for heterogeneous data: for instance given monitoring data in continuous time, images and clinical data, what is the risk for an event to occur? Is it useful to have all the sources or do they provide the same information? We will additionally develop solutions to **learn from decentralized data (federated learning)**, to **handle missing values** in a supervised learning setting and to **improve the confidence of the outputs of the predictive models**.
3. Personalized medicine with privacy and fairness guarantees. We 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.

The aim is to **push methodological innovation up to the stakeholders** (patients, clinicians, regulators, etc.). Consequently, beyond these methodological developments, innovative responses to the public health challenge posed by respiratory allergies are targeted. In addition to **leveraging machine learning algorithms and leveraging appropriate data, combining them with clinical expertise and existing recommendations** is necessary. Long-term aims are to have both a strong scientific and societal impact with a substantial impact on the quality of care for patients and major consequences for the medical profession by **providing a much earlier access to innovative solutions and more efficient treatment and care**. With a successful proof of concept in the domain of allergies, by having clear **reproducible pipelines, methodologies, software** (by providing clinical decision making system tools) we could thereafter consider other pathologies (such as traumatology and oncology studied at IDESP). Hence, a joint team between Inria and Inserm provides a unique opportunity for trans-disciplinary research and collaboration bringing together mathematical, methodological, technological and medical expertise. The PreMeDiCaL 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 is 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 [50] 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 leverage observational data for treatment effect estimation in health domains, several laws built on studies by the USA Food and Drug Administration (FDA) encourage the use of "real world data" (RWD), defined as data "derived from sources other than randomized clinical trials", 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. [13] 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 importance variables 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 [48] 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) [56, 52, 49] 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, [53]) 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 [9]; 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 [55], and even to reconstruct some of the training data points [51]. 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 [54], 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 health-care 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 [57]. 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

The first application domain of PreMeDICAL is **respiratory diseases** and in particular **Asthma**. For more than 30 years, there has been an increase in a 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 offer new opportunities to design personalized solutions by adapting treatment to the patient and the organizational context, leading to improved patient care and prevention policies.

With a successful proof of concept in the domain of allergies, by having clear **reproducible pipelines, methodologies, software**, we will thereafter consider other pathologies (such as traumatology and oncology studied at IDESP).

## 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 serve 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 software (starting with open access) 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 trials' 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

- Julie Josse won the **Inria - French Academy of Sciences Young Researchers Prize**. This prize is awarded to a scientist under forty years of age, working in a French institution, who has made a major contribution to the field of computer and mathematical sciences through his or her research, transfer or innovation activities.
- Maxime Fosset got a **fulbright French-USA PhD grant** and a mobility grant from Societe Française de Réanimation en langues Française. He is spending 6 months at Harvard Medical School (Nov. 2024- ).
- Pan Zhao received the Institute of Mathematical Statistics (IMS) **Hannan Graduate Student Travel Awards**. The award recipients, who are IMS members, can use the funds to attend any IMS-sponsored or co-sponsored meeting.

### 6.2 PhD defenses

- Margaux Zaffran defended her Phd "Post-hoc predictive uncertainty quantification: methods with applications to electricity price forecasting" on June 25, 2024.
- Pan Zhao defended his Phd "Topics in Causal Inference and Policy Learning with Applications to Precision Medicine" on Wednesday September 4, 2024.
- Marie Felicia Beclin defended her Phd "Development of intelligent models from CT scan data of patients treated with Benralizumab," on December 5, 2024.

### 6.3 Other

Following the **health data hub challenge allergen-chip**, the Premedical team and clinical collaborators specialized in allergies have started a collaboration on data from the Société Française d'Allergies to determine molecular allergen profiles and their links to clinical symptoms. The stakes are high: the WHO estimates that by 2050, one in two people will suffer from respiratory diseases, like allergies and asthma.

## 7 New software, platforms, open data

### 7.1 New software

#### 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

### 7.2 New platforms

- **Causal inference taskview:** to list and organize all the R packages on causal inference
- **R-miss-tastic a platform to gather and create resources on missing data, aimed at** researchers and students who often don't have lecture on missing values. It includes bibliography, courses, tutorials, implementations, pipelines of analysis in R and Python, etc.

**Participants:** Julie Josse, Pan Zhao.

## 8 New results

### 8.1 Treatment effect estimation

**Results: Choice of the causal measure [2]**

**Participants:** Julie Josse.

There are many measures to report so-called treatment or causal effect: absolute difference, ratio, odds ratio, number needed to treat, and so on. The choice of a measure, e.g. absolute versus relative, is often debated because it leads to different appreciations of the same phenomenon; but it also implies different heterogeneity of treatment effect. In addition some measures – but not all – have appealing properties such as collapsibility, matching the intuition of a population summary. We review common measures and their pros and cons typically brought forward. Doing so, we clarify notions of collapsibility and treatment effect heterogeneity, unifying different existing definitions. Our main contribution is to propose to reverse the thinking: rather than starting from the measure, we start from a non-parametric generative model of the outcome. Depending on the nature of the outcome, some causal measures disentangle treatment modulations from baseline risk. Therefore, our analysis outlines an understanding of what heterogeneity and homogeneity of treatment effect mean, not through the lens of the measure, but through the lens of the covariates. Our goal is the generalization of causal measures. We show that different sets of covariates are needed to generalize an effect to a different target population depending on (i) the causal measure of interest, (ii) the nature of the outcome, and (iii) the generalization's method itself (generalizing either conditional outcome or local effects).

**Results: Federated Causal Inference [36]**

**Participants:** Remi Khellaf, Aurélien Bellet, Julie Josse.

Randomized Controlled Trials (RCTs) are the gold standard for estimating the Average Treatment Effect (ATE) in evidence-based medicine, but their limitations—such as stringent eligibility criteria and small sample sizes—have led to the prominence of meta-analyses, the pinnacle of evidence in clinical research, which aggregate evidence from multiple studies to enhance statistical power and precision.

Despite extensive guidelines on conducting meta-analyses, multi-centric approaches still face significant challenges. These primarily arise from heterogeneity caused by imbalances in datasets, variations in populations across studies, and center effects due to differing practices across institutions. Moreover, simply aggregating local estimates is not the only approach to conducting meta-analyses. However, implementing “one-stage” meta-analyses that pool individual patient data from all centers is practically challenging due to data silos and personal data regulations.

Federated causal inference offers a promising alternative by allowing decentralized data sources to collaborate without sharing raw data, thus maintaining privacy and compliance with regulations. This work investigates three federated ATE estimation approaches—meta-analysis estimators, one-shot federated estimators, and gradient-based federated estimators—comparing their trade-offs in statistical efficiency, communication costs, and robustness to heterogeneity. The study demonstrates that meta-analysis estimators can achieve statistical efficiency comparable to pooled data analysis when sufficient data is available at each center, while naturally accommodating center effects. In contrast, while gradient-based approaches excel in low-data scenarios, one-shot estimators can be robust to distributional shifts but suffer from increased variance when center effects are present.

Guidelines and a decision diagram are provided to help practitioners choose the most appropriate approach based on data and heterogeneity conditions.

### **Results: Distribution on Distribution Regression to model Treatment Response Assessment in Asthma Patients**

**Participants:** Marie Felicia Beclin, Nicolas Molinari.

Medical imaging plays a crucial role in evaluating treatment efficacy. While practitioners traditionally rely on specific biomarkers and clinical data, incorporating informative features derived from medical imaging can enhance treatment response prediction. This research focuses on thoracic scans taken in expiration and inspiration before and after one year of Benralizumab treatment for asthma patients.

Following image segmentation, histograms are calculated to represent the distribution of voxel intensities. The underlying hypothesis posits that patients with improved conditions will exhibit enhanced expiration scans after treatment, evident in the histograms through a rightward shift, indicating higher Hounsfield Unit (HU) values. To predict treatment's response, we develop an histogram on histogram regression. Unlike existing methods, our proposed model goes beyond point-wise estimation of coefficient, offering an inferential framework to obtain p-values and confidence intervals for assessing treatment effects.

## **8.2 Handling missing values**

### **Results: Missing values imputation [41]**

**Participants:** 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. The present paper attempts to take a step back and provide a more systematic analysis. Starting from an in-depth discussion of the Missing at Random (MAR) condition for nonparametric imputation, we first develop an identification result, showing that the widely used Multiple Imputation by Chained Equations (MICE) approach indeed identifies the right conditional distributions. Building on this analysis, we propose three essential properties a successful imputation method should meet, thus enabling a more principled evaluation of existing methods and more targeted development of new methods. In particular, we introduce a new imputation method, denoted mice-DRE, that meets two out of the three criteria. We then discuss and refine ways to rank imputation methods, developing a powerful, easy-to-use scoring algorithm to rank missing value imputations.

### **Results: Conformal prediction with missing values [47]**

**Participants:** Margaux Zaffran, Julie Josse.

By leveraging increasingly large data sets, statistical algorithms and machine learning methods can be used to support, high-stakes decision-making problems such as autonomous driving, medical or civic applications, and more. To ensure the safe deployment of predictive models, it is crucial to quantify the uncertainty of the resulting predictions, communicating the limits of predictive performance. Uncertainty quantification attracts a lot of attention in recent years, particularly methods that are based on Conformal Prediction.

We investigate how to adequately quantify predictive uncertainty with missing covariates. A bottleneck is that missing values induce heteroskedasticity on the response's predictive distribution given

the observed covariates. Thus, we focus on building predictive sets for the response that are valid conditionally to the missing values pattern. We show that this goal is impossible to achieve informatively in a distribution-free fashion, and we propose useful restrictions on the distribution class. Motivated by these hardness results, we characterize how missing values and predictive uncertainty intertwine. Particularly, we rigorously formalize the idea that the more missing values, the higher the predictive uncertainty. Then, we introduce a generalized framework, coined CP-MDA-Nested, outputting predictive sets in both regression and classification. Under independence between the missing value pattern and both the features and the response (an assumption justified by our hardness results), these predictive sets are valid conditionally to any pattern of missing values. Moreover, it provides great flexibility in the trade-off between statistical variability and efficiency. Finally, we experimentally assess the performances of CP-MDA-Nested beyond its scope of theoretical validity, demonstrating promising outcomes in more challenging configurations than independence.

### 8.3 Learning with privacy guarantees

#### Results: Rényi Pufferfish Privacy [27]

**Participants:** Aurélien Bellet.

Pufferfish privacy is a flexible generalization of differential privacy that allows to model arbitrary secrets and adversary’s prior knowledge about the data (e.g., correlation across individuals). Unfortunately, designing general and tractable Pufferfish mechanisms that do not compromise utility is challenging. Furthermore, this framework does not provide the composition guarantees needed for a direct use in iterative machine learning algorithms. To mitigate these issues, we introduce a Rényi divergence-based variant of Pufferfish and show that it allows us to extend the applicability of the Pufferfish framework. We first generalize the Wasserstein mechanism to cover a wide range of noise distributions and introduce several ways to improve its utility. We also derive stronger guarantees against out-of-distribution adversaries. Finally, as an alternative to composition, we prove privacy amplification results for contractive noisy iterations and showcase the first use of Pufferfish in private convex optimization. A common ingredient underlying our results is the use and extension of shift reduction lemmas.

#### Results: Relative Gaussian Mechanism [24]

**Participants:** Aurélien Bellet.

The Gaussian Mechanism (GM), which consists in adding Gaussian noise to a vector-valued query before releasing it, is a standard privacy protection mechanism. In particular, given that the query respects some L2 sensitivity property (the L2 distance between outputs on any two neighboring inputs is bounded), GM guarantees Rényi Differential Privacy (RDP). Unfortunately, precisely bounding the L2 sensitivity can be hard, thus leading to loose privacy bounds. In this work, we consider a Relative L2 sensitivity assumption, in which the bound on the distance between two query outputs may also depend on their norm. Leveraging this assumption, we introduce the Relative Gaussian Mechanism (RGM), in which the variance of the noise depends on the norm of the output. We prove tight bounds on the RDP parameters under relative L2 sensitivity, and characterize the privacy loss incurred by using output-dependent noise. In particular, we show that RGM naturally adapts to a latent variable that would control the norm of the output. Finally, we instantiate our framework to show tight guarantees for Private Gradient Descent, a problem that naturally fits our relative L2 sensitivity assumption.

#### Results: Confidential Proof of Differentially Private Training [29]

**Participants:** Ioan Tudor Cebere, Aurélien Bellet.

Post hoc privacy auditing techniques can be used to test the privacy guarantees of a model, but come with several limitations: (i) they can only establish lower bounds on the privacy loss, (ii) the intermediate model updates and some data must be shared with the auditor to get a better approximation of the privacy loss, and (iii) the auditor typically faces a steep computational cost to run a large number of attacks. In this paper, we propose to proactively generate a cryptographic certificate of privacy during training to forego such auditing limitations. We introduce Confidential-DPproof, a framework for Confidential Proof of Differentially Private Training, which enhances training with a certificate of the  $(\epsilon, \delta)$ -DP guarantee achieved. To obtain this certificate without revealing information about the training data or model, we design a customized zero-knowledge proof protocol tailored to the requirements introduced by differentially private training, including random noise addition and privacy amplification by subsampling. In experiments on CIFAR-10, Confidential-DPproof trains a model achieving state-of-the-art 91% test accuracy with a certified privacy guarantee of  $(\epsilon = 0.55, \delta = 10^{-5})$ -DP in approximately 100 hours.

#### **Results: Private Training of Lipschitz Neural Networks [21]**

**Participants:** Aurélien Bellet.

State-of-the-art approaches for training Differentially Private (DP) Deep Neural Networks (DNN) face difficulties to estimate tight bounds on the sensitivity of the network's layers, and instead rely on a process of per-sample gradient clipping. This clipping process not only biases the direction of gradients but also proves costly both in memory consumption and in computation. To provide sensitivity bounds and bypass the drawbacks of the clipping process, we propose to rely on Lipschitz constrained networks. Our theoretical analysis reveals an unexplored link between the Lipschitz constant with respect to their input and the one with respect to their parameters. By bounding the Lipschitz constant of each layer with respect to its parameters, we prove that we can train these networks with privacy guarantees. Our analysis not only allows the computation of the aforementioned sensitivities at scale, but also provides guidance on how to maximize the gradient-to-noise ratio for fixed privacy guarantees. The code has been released as a [Python package](#).

#### **Results: Private Decentralized Learning with Random Walks [22]**

**Participants:** Aurélien Bellet.

The popularity of federated learning comes from the possibility of better scalability and the ability for participants to keep control of their data, improving data security and sovereignty. Unfortunately, sharing model updates also creates a new privacy attack surface. In this work, we characterize the privacy guarantees of decentralized learning with random walk algorithms, where a model is updated by traveling from one node to another along the edges of a communication graph. Using a recent variant of differential privacy tailored to the study of decentralized algorithms, namely Pairwise Network Differential Privacy, we derive closed-form expressions for the privacy loss between each pair of nodes where the impact of the communication topology is captured by graph theoretic quantities. Our results further reveal that random walk algorithms tends to yield better privacy guarantees than gossip algorithms for nodes close from each other. We supplement our theoretical results with empirical evaluation on synthetic and real-world graphs and datasets.

#### **Results: Privacy Attacks in Decentralized Learning [26]**

**Participants:** Aurélien Bellet.

Decentralized Gradient Descent (D-GD) allows a set of users to perform collaborative learning without sharing their data by iteratively averaging local model updates with their neighbors in a network graph. The absence of direct communication between non-neighbor nodes might lead to the belief that users cannot infer precise information about the data of others. In this work, we demonstrate the opposite, by proposing the first attack against D-GD that enables a user (or set of users) to reconstruct the private data of other users outside their immediate neighborhood. Our approach is based on a reconstruction attack against the gossip averaging protocol, which we then extend to handle the additional challenges raised by D-GD. We validate the effectiveness of our attack on real graphs and datasets, showing that the number of users compromised by a single or a handful of attackers is often surprisingly large. We empirically investigate some of the factors that affect the performance of the attack, namely the graph topology, the number of attackers, and their position in the graph.

### **Results: Privacy Auditing of Machine Learning [33]**

**Participants:** Ioan Tudor Cebere, Aurélien Bellet.

Machine learning models can be trained with formal privacy guarantees via differentially private optimizers such as Differential Privacy 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 amplify privacy. 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.

### **Results: Private Histogram Estimation [43]**

**Participants:** Aurélien Bellet.

We present Nebula, a system for differential private histogram estimation of data distributed among clients. Nebula enables clients to locally subsample and encode their data such that an untrusted server learns only data values that meet an aggregation threshold to satisfy differential privacy guarantees. Compared with other private histogram estimation systems, Nebula uniquely achieves all of the following: i) a strict upper bound on privacy leakage; ii) client privacy under realistic trust assumptions; iii) significantly better utility compared to standard local differential privacy systems; and iv) avoiding trusted third-parties, multi-party computation, or trusted hardware. We provide both a formal evaluation of Nebula's privacy, utility and efficiency guarantees, along with an empirical evaluation on three real-world datasets. We demonstrate that clients can encode and upload their data efficiently (only 0.0058 seconds running time and 0.0027 MB data communication) and privately (strong differential privacy guarantees  $\epsilon = 1$ ). On the United States Census dataset, the Nebula's untrusted aggregation server estimates histograms with above 88% better utility than the existing local deployment of differential privacy. Additionally, we

describe a variant that allows clients to submit multi-dimensional data, with similar privacy, utility, and performance. Finally, we provide an open source implementation of Nebula.

#### **Results: Correlated Gaussian Mechanism [37]**

**Participants:** Christian Janos Lebeda.

We consider the problem of releasing a sparse histogram under  $(\epsilon, \delta)$ -differential privacy. The stability histogram independently adds noise from a Laplace or Gaussian distribution to the non-zero entries and removes those noisy counts below a threshold. Thereby, the introduction of new non-zero values between neighboring histograms is only revealed with probability at most  $\delta$ , and typically, the value of the threshold dominates the error of the mechanism. We consider the variant of the stability histogram with Gaussian noise. Recent works reduced the error for private histograms using correlated Gaussian noise. However, these techniques can not be directly applied in the very sparse setting. Instead, we adopt Lebeda's technique and show that adding correlated noise to the non-zero counts only allows us to reduce the magnitude of noise when we have a sparsity bound. This, in turn, allows us to use a lower threshold by up to a factor of 1/2 compared to the non-correlated noise mechanism. We then extend our mechanism to a setting without a known bound on sparsity. Additionally, we show that correlated noise can give a similar improvement for the more practical discrete Gaussian mechanism.

## **8.4 Federated learning**

#### **Results: Generalization Guarantees for Decentralized SGD [25]**

**Participants:** Aurélien Bellet.

This work presents a new generalization error analysis for Decentralized Stochastic Gradient Descent (D-SGD) based on algorithmic stability. The obtained results overhaul a series of recent works that suggested an increased instability due to decentralization and a detrimental impact of poorly-connected communication graphs on generalization. On the contrary, we show, for convex, strongly convex and non-convex functions, that D-SGD can always recover generalization bounds analogous to those of classical SGD, suggesting that the choice of graph does not matter. We then argue that this result is coming from a worst-case analysis, and we provide a refined optimization-dependent generalization bound for general convex functions. This new bound reveals that the choice of graph can in fact improve the worst-case bound in certain regimes, and that surprisingly, a poorly-connected graph can even be beneficial for generalization.

#### **Results: Federated Conformal Prediction [35]**

**Participants:** Aurélien Bellet.

We study conformal prediction in the one-shot federated learning setting. The main goal is to compute marginally and training-conditionally valid prediction sets, at the server-level, in only one round of communication between the agents and the server. Using the quantile-of-quantiles family of estimators and split conformal prediction, we introduce a collection of computationally-efficient and distribution-free algorithms that satisfy the aforementioned requirements. Our approaches come from theoretical results related to order statistics and the analysis of the Beta-Beta distribution. We also prove upper bounds on the coverage of all proposed algorithms when the nonconformity scores are almost surely distinct. For algorithms with training-conditional guarantees, these bounds are of the same order of magnitude as those of the centralized case. Remarkably, this implies that the one-shot

federated learning setting entails no significant loss compared to the centralized case. Our experiments confirm that our algorithms return prediction sets with coverage and length similar to those obtained in a centralized setting.

## 8.5 Fair machine learning

**Results: Synthetic Data Generation for Intersectional Fairness [39]**

**Participants:** Aurélien Bellet.

In this work, we introduce a data augmentation approach specifically tailored to enhance intersectional fairness in classification tasks. Our method capitalizes on the hierarchical structure inherent to intersectionality, by viewing groups as intersections of their parent categories. This perspective allows us to augment data for smaller groups by learning a transformation function that combines data from these parent groups. Our empirical analysis, conducted on four diverse datasets including both text and images, reveals that classifiers trained with this data augmentation approach achieve superior intersectional fairness and are more robust to "leveling down" when compared to methods optimizing traditional group fairness metrics.

## 8.6 Uncertainty quantification

**Participants:** Julie Josse.

**Results: Probabilistic Prediction of Arrivals and Hospitalizations in Emergency Departments in Île-de-France [45]**

Adaptive probabilistic forecasting of French electricity spot prices

**Background:** Forecasts of future demand is foundational for effective resource allocation in emergency departments (EDs). As ED demand is inherently variable, it is important for forecasts to characterize the range of possible future demand. However, extant research focuses primarily on producing point forecasts using a wide variety of prediction algorithms. In this study, our objective is to generate point and interval predictions that accurately characterize the variability in ED demand using ensemble methods that combine predictions from multiple base algorithms based on their empirical performance.

**Methods:** Data consisted in daily arrivals and subsequent hospitalizations at 72 emergency departments in Ile-de-France from 2014-2018. Additional explanatory variables were collected including public and school holidays, meteorological variables, and public health trends. One-day ahead point and 80% interval pre- dictions of arrivals and hospitalizations were produced by predicting the 10%, 50%, and 90% quantiles of the forecast distribution. Quantile prediction algorithms included methods such as ARIMAX, variations of random forests, and generalized additive models. Ensemble predictions were then formed using Exponentially Weighted Averaging, Bernstein Online Aggregation, and Super Learning. Prediction intervals were post-processed using Adaptive Conformal Inference techniques. Point predictions were evaluated by their Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE), and 80% interval predictions by their empirical coverage and mean interval width.

**Results:** For point forecasts, ensemble methods achieved lower average MAE and MAPE than any of the base algorithms. All of the base algorithms and ensemble methods yielded prediction intervals with near optimal empirical coverage after conformalization. For hospitalizations, the shortest mean interval widths were achieved by the ensemble methods.

**Conclusions:** Ensemble methods yield joint point and prediction intervals that adapt to individual EDs and achieve better performance than individual algorithms. Conformal inference techniques improves the performance of the prediction intervals.

**Keywords** Emergency department, Time series forecasting, Machine learning, Ensemble learning, Confor- mal inference

**Participants:** Margaux Zaffran.

### **Results: Adaptive probabilistic forecasting of French electricity spot prices [34]**

Electricity price forecasting (EPF) plays a major role for electricity companies as a fundamental entry for trading decisions or energy management operations. As electricity can not be stored, electricity prices are highly volatile which make EPF a particularly difficult task. This is all the more true when dramatic fortuitous events disrupt the markets. Trading and more generally energy management decisions require risk management tools which are based on probabilistic EPF (PEPF). In this challenging context, we argue in favor of the deployment of highly adaptive black-boxes strategies allowing to turn any forecasts into a robust adaptive predictive interval, such as conformal prediction and online aggregation, as a fundamental last layer of any operational pipeline. We propose to investigate a novel data set containing the French electricity spot prices during the turbulent 2020-2021 years, and build a new explanatory feature revealing high predictive power, namely the nuclear availability. Benchmarking state-of-the-art PEPF on this data set highlights the difficulty of choosing a given model, as they all behave very differently in practice, and none of them is reliable. However, we propose an adequate conformalization, coined Online Sequential Split Conformal Prediction (OSSCP-horizon), that improves the performances of PEPF methods, even in the most hazardous period of late 2021. Finally, we emphasize that combining it with online aggregation significantly outperforms any other approaches, and should be the preferred pipeline, as it provides trustworthy probabilistic forecasts.

## **8.7 Application domain: allergies, ICU care**

**Participants:** Pascal Demoly.

### **Results: Impact of liquid sublingual immunotherapy on asthma onset and progression in patients with allergic rhinitis: a nationwide population-based study (EfficAPSI study)**

**Background:** The only disease-modifying treatment currently available for allergic rhinitis (AR) is allergen immunotherapy (AIT). The main objective of the EfficAPSI real-world study (RWS) was to evaluate the impact of liquid sublingual immunotherapy (SLIT-liquid) on asthma onset and evolution in AR patients.

**Methods:**

An analysis with propensity score weighting was performed using the EfficAPSI cohort, comparing patients dispensed SLIT-liquid with patients dispensed AR symptomatic medication with no history of AIT (controls). Index date corresponded to the first dispensation of either treatment. The sensitive definition of asthma event considered the first asthma drug dispensation, hospitalization or long-term disease (LTD) for asthma, the specific one omitted drug dispensation and the combined one considered omalizumab or three ICS ± LABA dispensation, hospitalization or LTD. In patients with pre-existing asthma, the GINA treatment step-up evolution was analyzed.

**Findings:** In this cohort including 112,492 SLIT-liquid and 333,082 controls, SLIT-liquid exposure was associated with a significant lower risk of asthma onset. Exposure to SLIT was associated with a one-third reduction in GINA step-up, irrespective of baseline treatment steps

**Interpretation:** In this national RWS with the largest number of person-years of follow-up to date in the field of AIT, SLIT-liquid was associated with a significant reduction in the risk of asthma onset or worsening. The use of three definitions (sensitive or specific) and GINA step-up reinforced the rigorous methodology, substantiating SLIT-liquid evidence as a causal treatment option for patients with respiratory allergies.

**Participants:** Julie Josse.

### **Results: Pilot deployment of a machine-learning enhanced prediction of need for hemorrhage resuscitation after trauma - the ShockMatrix pilot study [14]**

Importance: Decision-making in trauma patients remains challenging and often results in deviation from guidelines. Machine-Learning (ML) enhanced decision-support could improve hemorrhage resuscitation. Aim: To develop a ML enhanced decision support tool to predict Need for Hemorrhage Resuscitation (NHR) (part I) and test the collection of the predictor variables in real time in a smartphone app (part II). Design, setting, and participants: Development of a ML model from a registry to predict NHR relying exclusively on prehospital predictors. Several models and imputation techniques were tested. We also assess the feasibility to collect the predictors of the model in a customized smartphone app during prealert and generate a prediction in four level-1 trauma centers to compare the predictions to the gestalt of the trauma leader.

Main outcomes and measures: Part 1: Model output was NHR defined by 1) at least one RBC transfusion in resuscitation, 2) transfusion  $\geq 4$  RBC within 6 h, 3) any hemorrhage control procedure within 6 h or 4) death from hemorrhage within 24 h. The performance metric was the F4-score and compared to reference scores (RED FLAG, ABC). In part 2, the model and clinician prediction were compared with Likelihood Ratios (LR).

Results: From 36,325 eligible patients in the registry (Nov 2010—May 2022), 28,614 were included in the model development (Part 1). Median age was 36 [25-52], median ISS 13 [5-22], 3249/28614 (11%) corresponded to the definition of NHR. A XGBoost model with nine prehospital variables generated the best predictive performance for NHR according to the F4-score with a score of 0.76 [0.73-0.78]. Over a 3-month period (Aug-Oct 2022), 139 of 391 eligible patients were included in part II (38.5%), 22/139 with NHR. Clinician satisfaction was high, no workflow disruption observed and LRs comparable between the model and the clinicians.

Conclusions and relevance: The ShockMatrix pilot study developed a simple ML-enhanced NHR prediction tool demonstrating a comparable performance to clinical reference scores and clinicians. Collecting the predictor variables in real-time on prealert was feasible and caused no workflow disruption.

## **9 Bilateral contracts and grants with industry**

### **9.1 Bilateral contracts with industry**

**Participants:** Julie Josse, Helene Bonneau-Chloup, 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 obtained 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:** 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 objectives 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 provide the state of the art in this scientific 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

**Participants:** Julie Josse, Nicolas Molinari, Aurélien Bellet, Pascal Demoly.

## 10 Partnerships and cooperations

### 10.1 International research visitors

#### 10.1.1 Visits of international scientists

##### **Shu Yang**

**Status** Assistant Professor

**Institution of origin:** North Carolina University

**Country:** USA

**Dates:** May, 17 to 23

**Context of the visit:** Research work on causal measures and transportability of treatment effects

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

#### Other international visits to the team

##### **Lena Stempfle**

**Status** PhD

**Institution of origin:** Chalmers University of Technology

**Country:** Sweden

**Dates:** April 22 to May 22

**Context of the visit:** **Research collaboration on Interpretable Machine Learning with missing values and medical applications**

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

#### 10.1.2 Visits to international teams

##### Research stays abroad

##### **Maxime Fosset**

**Visited institution:** Harvard Medical School

**Country:** USA

**Dates:** November 2024 - April 2025

**Context of the visit:** FullBright grant

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

## 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
  - 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 modelling in pharmacology: toward digital pharmacological twins"** - Julie Josse
- **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.

#### 10.2.4 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.

#### 10.2.5 Allergen-Chip-Challenge

The challenge **L'allergen-chip-challenge** aimed at creating a national dataset for artificial intelligence-assisted allergy diagnosis using semantic attributes and allergen multiplex technology. The challenge was supported by the Health Data Hub in collaboration with the company Trustee - Pascal Demoly

Three follow-up projects:

- grant PNRIA 2023 with Olivier Saut
- AAP MESSIDORE 2024 submitted, Pascal Demoly and Julie Josse lead one research axis
- **Team retreat** with Pascal Demoly and Julie Josse Julien Goret on Determination of molecular allergen profiles and links with respiratory and food allergies

#### 10.2.6 Grant from the National Interministerial Road Safety Observatory

Julie Josse - In collaboration with **Traumabase**. Grant for the SPOTE project (Specificities of Populations and Impact of Territories) aimed at studying the intra-hospital outcome of victims of road accidents treated, in critical care, in France, between 2013 and 2027.

#### 10.2.7 Grant from PHRC

Nicolas Molinari leads 3 work packages

- Evaluation of early venous stenting treatment of patients with newly diagnosed idiopathic intracranial hypertension
- Evaluation of venous stenting treatment of patients with idiopathic intracranial hypertension to pursue acetazolamide withdrawal
- REVERT - Reversing airway remodeling with Tezepelumab

#### 10.2.8 Grant from Institut Exposum Doctoral Nexus

Nicolas Molinari obtained a grant from **ExposUM Nexus 2024 Doctoral Nexus** for Phd students on "Modeling suicide risk," principal investigator of the axis (196,000 Euros).

#### 10.2.9 Grant from Directorate General for Healthcare Services (DGOS)

Nicolas Molinari obtained a grant from **the Health Data and Applications (DA<sub>t</sub>AE)** call for projects launched by the Directorate General for Healthcare Services (DGOS) and operated by the Health Data Hub for the APPCMMAF study to improve the care of patients on continuous positive airway pressure (CPAP), principal investigator (269,648 Euros).

### 10.3 Regional initiatives

Pascal Demoly

**UM Envi-H** Initiative by the University of Montpellier.

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.

**Participants:** Pascal Demoly, Nicolas Molinari, Julie Josse.

**ComexIA Health Occitanie** 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.

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: organisation

- Aurélien Bellet co-organizes the **Federated Learning One World webinar** (1100+ registered attendees) since May 2020.
- Aurélien Bellet: membre du comité scientifique des **55ièmes Journées de Statistique (JDS 2024)**
- Margaux Zaffran: Journée Young Statisticians and Probabilists **12th YSP**, Institut Henri Poincaré, Paris, Jan 2024.
- Margaux Zaffran, Charlotte Voinot: Recontres avec les conférencier.e.s invité.e.s et acteur.rice.s de la SFdS **Rencontres JdS**, Bordeaux, France, May 2024.
- Margaux Zaffran, Charlotte Voinot: **Déjeuners scientifiques JdS**, Bordeaux, France, May 2024.
- Margaux Zaffran: Mathematical Statistics Day, Paris, France
- Nicolas Molinari: Chair of the session "Explicability and causal inference: new ways of using data" at the 1st Biotherapies & AI in Occitanie workshop, October 2024.

#### 11.1.2 Scientific events: selection

##### Member of the conference program committees

- Julie Josse: IMS **International Conference on Data Science**, Nice, France, December 2024.

- Julie Josse: [Methodological and Computational Advances in Survival Analysis](#) Workshop, Nov 2024.
- Julie Josse: [useR!2024](#), Salzburg, July 2024.
- Aurélien Bellet: Area Chair for Neural Information Processing Systems, NeurIPS 2024
- Aurélien Bellet: Area Chair for International Conference on Machine Learning, ICML 2024
- Aurélien Bellet : Area Chair for Artificial Intelligence and Statistics, AISTATS 2025

### Reviewer

- Aurélien Bellet : [Workshop on Privacy and Security in Augmented, Virtual, and eXtended Realities](#) at WoWMoM 2024
- Aurélien Bellet : [Workshop on Privacy Regulation and Protection in Machine Learning](#) at ICLR 2024
- Aurélien Bellet : [Workshop on Security, Privacy and Information Theory](#) at CSF 2024
- Aurélien Bellet : [Workshop on Privacy-Preserving Artificial Intelligence](#) at AAAI 2025
- Aurélien Bellet : [CAp 2024](#)
- Aurélien Bellet : [APVP 2024](#)
- Ioan Tudor Cebere: ICML 2024
- Ioan Tudor Cebere: ICLR 2025
- Christian Janos Lebeda: AISTATS 2025
- Christian Janos Lebeda : OpenDP Privacy Proof Review Board member

### 11.1.3 Journal

#### Member of the editorial boards

- Julie Josse: 2024 - . Associate editor of [Foundations and Trends in Machine Learning](#)
- Aurélien Bellet is Action Editor for [Transactions of Machine Learning Research \(TMLR\)](#)
- Nicolas Molinari Statistics Editor for the journal [European Respiratory Journal](#)

#### Reviewer - reviewing activities

- Jeffrey Naf: Reviews for Transactions on Machine Learning Research (TMLR), 2024
- Jeffrey Naf: Reviews for Conference on Causal Learning and Reasoning (CLear), 2024

### 11.1.4 Invited talks

- Pascal Demoly: [Futurapolis Santé](#), "Exposome : la chasse aux ennemis de nos poumons est ouverte", Oct. 2024.
- Pascal Demoly: [Congrès Francophone d'Allergologie](#), Apr. 2024.
- Pascal Demoly: [Journée annuelle de l'Institut ExposUM](#), Nov. 2024.
- Julie Josse : [Bernoulli-IMS 11th World Congress in Probability and Statistics 2024](#), Talk in session on missing values.
- Julie Josse : [Symposium on Causality](#) (Panel), Sept. 2024, Florence.

- Julie Josse: **50 ans du CMAP**, Centre de Mathématiques Appliquées de l'Ecole Polytechnique, Sept. 2024, Paris.
- Julie Josse: 2nd Global Symposium of Research Methodology Innovation in Trauma and Emergency Care, May 2024, Columbus Ohio.
- Julie Josse: **European Conference of Causal Inference** (Eurocim), Apr. 2024., Copenhagen.
- Julie Josse: NIH Biostatistics Branch, Division of Cancer Epidemiology and Genetics, National Cancer Institute, February 2024
- Aurélien Bellet: **IMS International Conference on Statistics and Data Science** (ICSDS), Nice.
- Aurélien Bellet: **3rd Workshop on Principles of Distributed Learning**, Paris.
- Aurélien Bellet: **Workshop on AI Auditing**, Paris.
- Aurélien Bellet: **Learning and Optimization in Côte d'Azur Workshop**, Sophia-Antipolis.
- Aurélien Bellet: Inserm Workshop on Massive Genomic Data: Statistical and Bioinformatic Advances, Bordeaux.
- Aurélien Bellet: Privacy Alpine Seminar (Privaski), Corrençon en Vercors.
- Aurélien Bellet: **L'Oréal**, virtual.
- Aurélien Bellet: **Owkin**, virtual.
- Charlotte Voinot: Talk on introduction to causal inference, **Stat4Plant Seminar INRAE**, Jouy en josas, March 2024
- Charlotte Voinot: Talk on Causal survival analysis, **MECOSA (Methodological and Computational Advances in Survival Analysis)**, Inria Paris, November 2024
- Remi Khellaf: Talk on introduction to causal inference, **Stat4Plant Seminar INRAE**, Jouy en josas, March 2024
- Jeffrey Näf: Talk on Missing Values **icsds2024**, Nice, December 2024
- Pan Zhao: **IMS Asia-Pacific Rim Meeting**. Invited Talk. January 4 - 7, 2024, Melbourne, Australia.
- Ioan Tudor Ceberé: Talks on privacy auditing at Microsoft Research, Google Deepmind, Harvard and University of Toronto.
- Nicolas Molinari Invited speaker at the workshop "AI for bronchial diseases," December 13, 2024.

#### 11.1.5 Leadership within the scientific community

- Pascal Demoly : **full member of the Academy of Medicine**, 1st division
- Pascal Demoly : Animation of the network **e-allergies**
- Pascal Demoly : president of the "French Society of Allergology"
- Pascal Demoly : WHO Collaborating Center for "Scientific Support for Classifications".
- 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).
- Margaux Zaffran: President of "**Groupe Jeunes Statisticien.ne.s**"
- Charlotte Voinot: Treasurer of "**Groupe Jeunes Statisticien.ne.s**"
- Ioan Tudor Ceberé : Privacy Attacks Workgroup Leadership for OpenDP.

### 11.1.6 Scientific expertise

- Julie Josse: Member of the Searching Committee for ENDOMIC, Inria. 2024.
- Julie Josse: Advisory Board of HORIZON EUROPE (HORIZON-HLTH-2022-TOOL-11-02), more-europa 2023-.
- Julie Josse and Nicolas Molinari: Comité scientifique et éthique du CHU de Montpellier. Dec 2023 -
- Julie Josse: Evaluation of research projects for funding agency or promotions for tenured Professor positions. Washington University; John Hopkins University; ANRT (PhD Cifre);
- Aurélien Bellet: Member of the CNIL-Inria Privacy Award committee
- Aurélien Bellet: ethics advisor for the European Strategy Forum on Research Infrastructures (ESFRI) project [SLICES-PP](#)
- Nicolas Molinari: president of the Institutional Review Board (IRB) of the Adène group
- Nicolas Molinari: Expert for DGOS, ANR, and several GIRCI (research project evaluations).
- Nicolas Molinari: [Scientific Advisory Board of Nomics, "Make sleep medicine accessible"](#).

### 11.1.7 Research administration

- Aurélien Bellet: member of the Operational Committee for the assessment of Legal and Ethical risks (COERLE).
- Julie Josse member of CSD ("Comité Suivi Doctoral") Inria
- Nicolas Molinari: elected member of "Commissions scientifique spécialisées"(CSS) 6 of INSERM
- Margaux Zaffran Elected member, parity and diversity committee, CMAP, École polytechnique.

## 11.2 Teaching - Supervision - Juries

### 11.2.1 Teaching

- Engineering School: M2 students, 40heqTD, Introduction to Probabilistic Graphical Models and Deep Generative Models, Master recherche spécialité "Mathématiques Appliquées", M2 Mathématiques, Vision et Apprentissage (ENS Paris-Saclay), 1er semestre, 2024/2025 Rémi Khellaf
- Master: Institut de formation en masso-kinésithérapie, 9heqTD, statistics, Montpellier - Nicolas Molinari
- Master: Institut de formation en masso-kinésithérapie, head of the program, statistics, Montpellier - Nicolas Molinari
- Ecoles d'étiopathie, head of the program, statistics, Montpellier - Nicolas Molinari
- Master: EDSB « Epidémiologie, Données de Santé, Biostatistique », head of « Grands enjeux en santé », Université de Montpellier - Pascal Demoly

### 11.2.2 Supervision

#### PhD students:

- Julie Josse: Supervision of [Laura Fuentes Vincente](#) (grant Montpellier) with Antoine Chambaz, Nov 2024 -
- Julie Josse: Supervision of [Ahmed Boughdiri](#) (grant Inria), Sep 2023 -

- Julie Josse and Aurélien Bellet: Supervision of Rémi Khellaf (grant Montpellier) with Erwan Scornet, Sep 2023 -
- Julie Josse: Supervision of Charlotte Voinot with Bernard Sebastien (grant Phd thesis Cifre Sanofi), Apr. 2023 -
- Julie Josse: Supervision of the medical doctor (MD) Tobias Gauss with Pierre Bouzat (MD), Feb. 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 Margaux Zaffran (Cifre EDF) with Aymeric Dieuleveut, Yannig Goude and Olivier Ferron, Defended June 2024.
- Julie Josse: Supervision of Pan Zhao (grant MUSE) with Antoine Chambaz, Defended September 2024.
- Aurélien Bellet: Supervision of Jean-Rémy Conti with Stéphan Cléménçon, October 2021 -
- Aurélien Bellet: Supervision of Edwige Cyffers (defended in December 2024)
- 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 -
- Nicolas Molinari: Supervision of Coutureau J., December 2024 -
- Nicolas Molinari: Supervision of Ibrahim S., October 2024.
- Nicolas Molinari: Supervision of Marie Felicia Beclin, defended December 2024.
- Pascal Demoly Supervision of of Ileana Ghiordanescu, defended on December 3, 2024, entitled "Mathematical Modeling of Drug Hypersensitivity Reactions - From Phenotyping to Endotyping."

**Postdocs:**

- Julie Josse: Mathieu Even, Oct. 2024 - .
- Julie Josse: Houssam Zenati, Dec. 2023 - Dec. 2024. Joint supervision with Bertrand Thirion and Judith Abecassis.
- Julie Josse: Herb Susmann, Sept. 2023 - 2024. Joint supervision with Antoine Chambaz. Current position: postdoc NYU Grossman School of Medicine
- Julie Josse: Jeffrey Naf, Feb. 2023 -
- Aurélien Bellet: Batiste Le Bars, until July 2024
- Aurélien Bellet: Mathieu Dagneou, Dec 2024 -

**Masters:**

- Nicolas Molinari: Supervisor of the Master 2 internship (EDSB) of J. Coutureau (100%), "Score to differentiate malignant non-mass lesions and benign breast cancer," defended in June 2024.
- Nicolas Molinari: Co-supervisor of the Master 2 internship (EDSB) of M. Meerun (50%), "Prediction of mortality in severe acute pancreatitis," defended in June 2024.
- Nicolas Molinari: Supervisor of the Master 2 internship (EDSB) of F. Kucharczak (100%), "Contribution of statistical variability quantification in the diagnosis of Parkinson's disease," defended in June 2024.

### 11.2.3 Juries

#### Member of PhD/HDR committees:

- Julie Josse: CSI Eugène Berta, under the supervision of Francis Bach and Michael Jordan. 2024 -
- Julie Josse: PhD defense committee of Alexis Ayme under the supervision of Erwan Scornet, Claire Boyer and Aymeric Dieuleveut. Oct 2024.
- Julie Josse: PhD defense committee of Noemie Simon Tillaux, under the supervision of Florence tubach. Nov 2024.
- Julie Josse: HDR defense committee of Emilie Devijver. Nov 2024.
- Julie Josse: PhD defense committee Floriane Jochum, under the supervision of Anne Sophie Hamy. Dec. 2024.
- Julie Josse: PhD defense committee (reviewer) of Sophia Yazzourh under the supervision of Nicolas Savy and Philippe Saint Pierre.
- Julie Josse: Habilitation of Boris Hejblum, May 2024
- Julie Josse: CSI Rémy Chapelle, supervised by Bruno Falissard, Mohammed Sedki and Nicolas Vayatis. 2024 -
- Julie Josse: PhD defense committee of Armand Lacombe under the supervision of Michelle Sebag. Jan. 2024.
- Aurélien Bellet: Reviewer for the habilitation thesis (HDR) of Antoine Boutet. Dec. 2024.
- Aurélien Bellet: Reviewer for the PhD of Louis Leconte under the supervision of Eric Moulines, Lionel Trojman and Van Minh Nguyen. June 2024.
- Aurélien Bellet: Reviewer for the PhD of Mathieu Dagréou under the supervision of Samuel Vaiter and Thomas Moreau. Oct. 2024.
- Aurélien Bellet: PhD defense committee of Marie Garin under the supervision of Nicolas Vayatis. June 2024.
- Aurélien Bellet: PhD defense committee of Tanguy Lefort under the supervision of Joseph Salmon and Alexis Joly. Sep. 2024.
- Aurélien Bellet: PhD defense committee of Tuan-Anh Nguyen under the supervision of Denis Trystram and Kim Thang Nguyen. Oct. 2024.

#### Member of hiring committees:

- Julie Josse: Member of the committee Chaire de Professeur Junior, CBIO "Artificial Intelligence for Digital Health". Sep. 2024.
- Julie Josse: Member of the committee Chaire de Professeur Junior, ENS Lyon. June 2024.
- Julie Josse: Member of the committee Chaire de Professeur Junior, Statistics and Public Health - Inria Rennes. May 2024.
- Aurélien Bellet: Member of assistant professor recruiting committee - Université de Montpellier.

## 11.3 Popularization

### 11.3.1 Specific official responsibilities in science outreach structures

- Julie Josse: Committee on [Nomination for the Institute of Mathematical Statistics \(IMS\)](#) to select one candidate for IMS President. 2024 -

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

- Aurélien Bellet : article on federated learning for healthcare in Télécom Paris Alumni [\[link\]](#).
- Aurélien Bellet : Interview for LUM Magazine [\[link\]](#)
- Ioan Tudor Ceberu : hosting OpenMined's Privacy Tech Talk Series on Youtube, see [\[link\]](#) and [\[link\]](#)

### 11.3.3 Participation in Live events

- Julie Josse: Table ronde sur l'IA - Infravia.
- Margaux Zaffran, Charlotte Voinot: Session groupe jeune JdS [Sexisme ordinaire, violences sexistes et sexuelles, biais de genre. Quel est le constat aujourd'hui dans la recherche académique en France ?](#), Bordeaux, France, May 2024.
- Margaux Zaffran Volunteer, Séphora Berrebi Association. Participation to various masterclasses for high school girls.

## 12 Scientific production

### 12.1 Major publications

- [1] A. Bourdin, S. Bommart, G. Marin, I. Vachier, A. S. Gamez, E. Ahmed, C. Suehs and N. Molinari. 'Obesity in women with asthma: baseline disadvantage plus greater small-airway responsiveness'. In: *Allergy* (2022). DOI: [10.1111/all.15509](#). URL: [https://hal.science/hal-03772214](#).
- [2] 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](#) (cit. on p. 10).
- [3] 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](#).
- [4] 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](#).
- [5] 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](#).
- [6] 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](#).
- [7] 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](#).
- [8] F. Roubille, E. Matzner-Lober, S. Aguilhon, M. Rene, L. Lecourt, M. Galinier, J.-e. Ricci and N. Molinari. 'Impact of global warming on weight in patients with heart failure during the 2019 heatwave in France'. In: *ESC Heart Failure* (2022). DOI: [10.1002/ehf2.14129](#). URL: [https://hal.science/hal-03783874](#).

- [9] M. Zaffran, A. Dieuleveut, J. Josse and Y. Romano. ‘Conformal Prediction with Missing Values’. In: *Proceedings of Machine Learning Research*. ICML 2023 - 40 th International Conference on Machine Learning. Vol. PMLR202. Honolulu (Hawaii), United States, 23rd July 2023, p. 40578. URL: <https://hal.science/hal-03896384> (cit. on p. 6).
- [10] 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

- [11] L. S. Cardelli, M. Magaldi, A. Agullo, G. Richard, E. Nogue, P. Berdague, M. Galiner, F. Georger, F. Picard, E. Prunet, N. Molinari, A. Bourdin, D. Jaffuel and F. Roubille. ‘Sacubitril/valsartan has an underestimated impact on the right ventricle in patients with sleep-disordered breathing, especially central sleep apnoea syndrome’. In: *Archives of cardiovascular diseases* (2024), Online ahead of print. DOI: [10.1016/j.acvd.2024.04.003](https://doi.org/10.1016/j.acvd.2024.04.003). URL: <https://hal.science/hal-04580913>. In press.
- [12] B. Colnet, J. Josse, G. Varoquaux and E. Scornet. ‘Reweighting the RCT for generalization: finite sample error and variable selection’. In: *Journal of the Royal Statistical Society: Series A Statistics in Society* (24th May 2024). DOI: [10.1093/jrssa/qnae043](https://doi.org/10.1093/jrssa/qnae043). URL: <https://hal.science/hal-03822662>.
- [13] 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>. In press (cit. on p. 5).
- [14] T. Gauss, J.-D. Moyer, C. Colas, M. Pichon, N. Delhay, M. Werner, V. Ramonda, T. Sempe, S. Medjkoune, J. Josse et al. ‘Pilot deployment of a machine-learning enhanced prediction of need for hemorrhage resuscitation after trauma – the ShockMatrix pilot study’. In: *BMC Medical Informatics and Decision Making* 24.1 (28th Oct. 2024), p. 315. DOI: [10.1186/s12911-024-02723-9](https://doi.org/10.1186/s12911-024-02723-9). URL: <https://hal.science/hal-04832640> (cit. on p. 18).
- [15] D. Jaffuel, E. Serrano, C. Leroyer, A. Chartier and P. Demoly. ‘SQ HDM sublingual immunotherapy tablet for the treatment of HDM allergic rhinitis and asthma improves subjective sleepiness and insomnia: an exploratory analysis of the real-life CARIOCA study’. In: *Journal of Investigational Allergology and Clinical Immunology* 34.5 (2024). DOI: [10.18176/jiaci.0934](https://doi.org/10.18176/jiaci.0934). URL: <https://hal.science/hal-04232004>.
- [16] J. Josse, J. M. Chen, N. Prost, G. Varoquaux and E. Scornet. ‘On the consistency of supervised learning with missing values’. In: *Statistical Papers* 65.9 (4th Mar. 2024), pp. 5447–5479. DOI: [10.1007/s00362-024-01550-4](https://doi.org/10.1007/s00362-024-01550-4). URL: <https://hal.science/hal-02024202>.
- [17] H. Pan, D. Jarvis, J. Potts, L. Casas, D. Nowak, J. Heinrich, J. G. Aymerich, I. Urrutia, J. Martinez-Moratalla, J.-A. Gullon, A. Pereira-Vega, C. Raheison, S. Chanoine, P. Demoly, B. Leynaert, T. Gislason, N. Probst, M. J. Abramson, R. Jogi, D. Norback, T. Sigsgaard, M. Olivieri, C. Svanes and E. Fuertes. ‘Gas cooking indoors and respiratory symptoms in the ECRHS cohort’. In: *International Journal of Hygiene and Environmental Health* 256 (1st Mar. 2024), p. 114310. DOI: [10.1016/j.ijheh.2023.114310](https://doi.org/10.1016/j.ijheh.2023.114310). URL: <https://hal.science/hal-04467856>.
- [18] A. Sportisse, M. Marbac, F. Laporte, G. Celeux, C. Boyer, J. Josse and C. Biernacki. ‘Model-based Clustering with Missing Not At Random Data’. In: *Statistics and Computing* (18th June 2024). DOI: [10.1007/s11222-024-10444-2](https://doi.org/10.1007/s11222-024-10444-2). URL: <https://hal.science/hal-03494674>.
- [19] J.-b. Woillard, C. Benoist, A. Destere, M. Labriffe, G. Marchello, J. Josse and P. Marquet. ‘To be or not to be, when synthetic data meet clinical pharmacology: A focused study on pharmacogenetics’. In: *CPT: Pharmacometrics and Systems Pharmacology* (1st Sept. 2024), Online ahead of print. DOI: [10.1002/psp4.13240](https://doi.org/10.1002/psp4.13240). URL: <https://hal.science/hal-04747078>.

**International peer-reviewed conferences**

- [20] C. Bénard, J. Naf and J. Josse. ‘MMD-based Variable Importance for Distributional Random Forest’. In: *Proceedings of Machine Learning Research*. AISTATS 2024 - The 27th International Conference on Artificial Intelligence and Statistics. Vol. PMLR-238. Volume 238: International Conference on Artificial Intelligence and Statistics, 2-4 May 2024, Palau de Congressos, Valencia, Spain. Valence, Spain, 2024, pp. 1324–1332. URL: <https://hal.science/hal-04251588>.
- [21] L. Béthune, T. Massena, T. Boissin, Y. Prudent, C. Friedrich, F. Mamalet, A. Bellet, M. Serrurier, D. Vigouroux and C. Friedrich. ‘DP-SGD Without Clipping: The Lipschitz Neural Network Way’. In: ICLR 2024 - 12th International Conference on Learning Representations. Vienna (Austria), Austria, 2024. URL: <https://hal.science/hal-04610647> (cit. on p. 13).
- [22] E. Cyffers, A. Bellet and J. Upadhyay. ‘Differentially Private Decentralized Learning with Random Walks’. In: ICML 2024 - Forty-first International Conference on Machine Learning. Vienne (Autriche), Austria: arXiv, 2024. DOI: [10.48550/arXiv.2402.07471](https://doi.org/10.48550/arXiv.2402.07471). URL: <https://hal.science/hal-04610660> (cit. on p. 13).
- [23] M. Even, L. Ganassali, J. Maier and L. Massoulié. ‘Aligning Embeddings and Geometric Random Graphs: Informational Results and Computational Approaches for the Procrustes-Wasserstein Problem’. In: NeurIPS 2024 - 38th Conference on Neural Information Processing Systems. Vancouver (BC), Canada, 9th Dec. 2024. URL: <https://hal.science/hal-04895718>.
- [24] H. Hendrikx, P. Mangold and A. Bellet. ‘The Relative Gaussian Mechanism and its Application to Private Gradient Descent’. In: *PMLR*. AISTATS 2024 - 27th International Conference on Artificial Intelligence and Statistics. 238. Valencia, Spain, 29th Aug. 2024, pp. 3079–3087. URL: <https://hal.science/hal-04370596> (cit. on p. 12).
- [25] B. Le Bars, A. Bellet, M. Tommasi, K. Scaman and G. Neglia. ‘Improved Stability and Generalization Guarantees of the Decentralized SGD Algorithm’. In: ICML 2024 - The Forty-first International Conference on Machine Learning. Vienne, Austria, 21st July 2024. URL: <https://hal.science/hal-04611418> (cit. on p. 15).
- [26] A. E. Mrini, E. Cyffers and A. Bellet. ‘Privacy Attacks in Decentralized Learning’. In: ICML 2024 - Forty-first International Conference on Machine Learning. Vienne (Austria), Austria: arXiv, 2024. DOI: [10.48550/arXiv.2402.10001](https://doi.org/10.48550/arXiv.2402.10001). URL: <https://hal.science/hal-04610652> (cit. on p. 13).
- [27] C. Pierquin, A. Bellet, M. Tommasi and M. Boussard. ‘Rényi Pufferfish Privacy: General Additive Noise Mechanisms and Privacy Amplification by Iteration via Shift Reduction Lemmas’. In: International Conference on Machine Learning (ICML 2024). Vienna (Austria), Austria, 2024. URL: <https://inria.hal.science/hal-04363020> (cit. on p. 12).
- [28] T. Seoudi, D. Ayache, O. BENABBAD, F. PAGES, J. Charensol, M. Bahriz, N. Molinari, F. Gouzi and A. Vicet. ‘Breath analysis by quartz enhanced photoacoustic spectroscopy: A clinical study’. In: FLAIR 2024 - Field Laser Applications in Industry and Research 2024. Assise, Italy, 16th Sept. 2024. URL: <https://hal.science/hal-04770606>.
- [29] A. S. Shamsabadi, G. Tan, T. I. Cebere, A. Bellet, H. Haddadi, N. Papernot, X. Wang and A. Weller. ‘Confidential-DPproof: Confidential Proof of Differentially Private Training’. In: ICLR 2024 - 12th International Conference on Learning Representations. Vienna (Austria), Austria, 2024. URL: <https://hal.science/hal-04610635> (cit. on p. 12).
- [30] P. Zhao, A. Chambaz, J. Josse and S. Yang. ‘Positivity-free Policy Learning with Observational Data’. In: *Proceedings of Machine Learning Research*. AISTATS 2024 - The 27th International Conference on Artificial Intelligence and Statistics. Vol. PMLR-238. Volume 238: International Conference on Artificial Intelligence and Statistics, 2-4 May 2024, Palau de Congressos, Valencia, Spain. Valence, Spain, 2024, pp. 1918–1926. URL: <https://hal.science/hal-04236430>.

### Conferences without proceedings

- [31] T. Seoudi, D. Ayache, O. Benabbad, J. Charenzol, F. Pagès, E. Rosenkrantz, N. Molinari, M. Bahriz, F. Gouzi and A. Vicet. 'Photoacoustic sensing based on resonant mechanical transducers: application to diagnosis in breath'. In: ICPPP 2024 - 22st International Conference on Photoacoustic and Photothermal Phenomena. Coimbra (Portugal), Portugal, 8th July 2024. URL: <https://hal.umontpellier.fr/hal-04751842>.

### Reports & preprints

- [32] A. Boughdiri, J. Josse and E. Scornet. *Quantifying Treatment Effects: Estimating Risk Ratios in Causal Inference*. 15th Oct. 2024. URL: <https://hal.science/hal-04588018>.
- [33] T. Cebere, A. Bellet and N. Papernot. *Tighter Privacy Auditing of DP-SGD in the Hidden State Threat Model*. 14th Oct. 2024. URL: <https://inria.hal.science/hal-04863204> (cit. on p. 14).
- [34] G. Dutot, M. Zaffran, O. Féron and Y. Goude. *Adaptive probabilistic forecasting of French electricity spot prices*. 24th May 2024. URL: <https://hal.science/hal-04586509> (cit. on p. 17).
- [35] P. Humbert, B. Le Bars, A. Bellet and S. Arlot. *Marginal and training-conditional guarantees in one-shot federated conformal prediction*. 18th May 2024. URL: <https://hal.science/hal-04579882> (cit. on p. 15).
- [36] R. Khellaf, A. Bellet and J. Josse. *Federated Causal Inference: Multi-Studies ATE Estimation beyond Meta-Analysis*. 14th Oct. 2024. URL: <https://hal.science/hal-04747923> (cit. on p. 10).
- [37] C. J. Lebeda and L. Retschmeier. *The Correlated Gaussian Sparse Histogram Mechanism*. 13th Dec. 2024. URL: <https://inria.hal.science/hal-04895301> (cit. on p. 15).
- [38] C. J. Lebeda and J. Tëtek. *Testing Identity of Distributions under Kolmogorov Distance in Polylogarithmic Space*. 29th Oct. 2024. URL: <https://inria.hal.science/hal-04895297>.
- [39] G. Maheshwari, A. Bellet, P. Denis and M. Keller. *Synthetic Data Generation for Intersectional Fairness by Leveraging Hierarchical Group Structure*. 23rd May 2024. URL: <https://inria.hal.science/hal-04863199> (cit. on p. 16).
- [40] J. Näf, P. Bachmann and M. Meierer. *Customer Base Analysis in Non-Contractual Settings: A Model of Customer Attrition, Transactions, and Spending*. 9th Sept. 2024. URL: <https://hal.science/hal-04692045>.
- [41] J. Näf, E. Scornet and J. Josse. *What Is a Good Imputation Under MAR Missingness?* 14th Jan. 2025. URL: <https://hal.science/hal-04521894> (cit. on p. 11).
- [42] J. Näf and H. Susmann. *Causal-DRF: Conditional Kernel Treatment Effect Estimation using Distributional Random Forest*. 13th Nov. 2024. URL: <https://hal.science/hal-04854037>.
- [43] A. S. Shamsabadi, P. Snyder, R. Giles, A. Bellet and H. Haddadi. *Nebula: Efficient, Private and Accurate Histogram Estimation*. 17th Sept. 2024. URL: <https://inria.hal.science/hal-04863194> (cit. on p. 14).
- [44] L. Stempfle, A. James, J. Josse, T. Gauss and F. Johansson. *Expert Study on Interpretable Machine Learning Models with Missing Data*. 2024. DOI: [10.48550/arXiv.2411.09591](https://doi.org/10.48550/arXiv.2411.09591). URL: <https://hal.science/hal-04894332>.
- [45] H. Susmann, A. Chambaz, J. Josse, M. Wargon, P. Aegerter and E. Bacry. *Probabilistic Prediction of Arrivals and Hospitalizations in Emergency Departments in Île-de-France*. 9th Apr. 2024. URL: <https://hal.science/hal-04539380> (cit. on p. 16).
- [46] C. Voinot, C. Berenfeld, I. Mayer, B. Sebastien and J. Josse. *Causal survival analysis, Estimation of the Average Treatment Effect (ATE): Practical Recommendations*. 20th Dec. 2024. URL: <https://hal.science/hal-04852134>.
- [47] M. Zaffran, J. Josse, Y. Romano and A. Dieuleveut. *Predictive Uncertainty Quantification with Missing Covariates*. 24th May 2024. URL: <https://hal.science/hal-04587674> (cit. on p. 11).

### 12.3 Cited publications

- [48] 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. 5).
- [49] 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. 6).
- [50] L. Pahuš, 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. 4).
- [51] 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. 6).
- [52] 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. 6).
- [53] 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. 6).
- [54] 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. 6).
- [55] 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. 6).
- [56] V. Vovk, A. Gammerman and G. Shafer. *Algorithmic Learning in a Random World*. en. Springer US, 2005 (cit. on p. 6).
- [57] 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. 7).