Offer #2025-09177

Post-Doctoral Research Visit F/M Differential Privacy and Fairness-Aware AI

Contract type : Fixed-term contract

Level of qualifications required : PhD or equivalent

Fonction : Post-Doctoral Research Visit

About the research centre or Inria department

The Inria Grenoble research center groups together almost 600 people in 27 research teams and 8 research support departments.

Staff is present on three campuses in Grenoble, in close collaboration with other research and higher education institutions (University Grenoble Alpes, CNRS, CEA, INRAE, …), but also with key economic players in the area.

Inria Grenoble is active in the fields of high-performance computing, verification and embedded systems, modeling of the environment at multiple levels, and data science and artificial intelligence. The center is a top-level scientific institute with an extensive network of international collaborations in Europe and the rest of the world.

Context

Context. This postdoctoral position is part of the ANR JCJC project AI-PULSE (Aligning Privacy, Utility, and Fairness for Responsible AI), led by Héber H. Arcolezi within the Privatics team at Inria Grenoble.
AI-PULSE addresses a key open question in responsible AI: can we design practical machine learning systems that satisfy strong privacy guarantees [1] and fairness [2] constraints simultaneously, without sacrificing utility?

The postdoctoral researcher will play a central role in the project's core scientific tasks. The work will build on preliminary findings that local differential privacy (LDP) can help decrease biases [3, 4] instead of always worsening them, a challenging conventional assumption in the field [5].
The position will be embedded in a network of national and international collaborations, especially with ÉTS Montréal and UQAM, and includes opportunities for short- and medium-term research visits, collaborative development, and joint publications.

We are looking for a candidate who can bring new expertise to the team and/or help strengthen collaborations within the project consortium. The postdoc will work closely with project members to deliver high-impact theoretical and practical contributions to responsible AI.

Starting Date. Flexible start in late 2025 or early 2026.

Selected references.

- [1] Dwork, Cynthia, and Aaron Roth. "The algorithmic foundations of differential privacy." Foundations and Trends® in Theoretical Computer Science 9.3–4 (2014): 211-407.
- [2] Barocas, Solon, Moritz Hardt, and Arvind Narayanan. Fairness and machine learning: Limitations and opportunities. MIT press, 2023.
- [3] Makhlouf, Karima, et al. "A systematic and formal study of the impact of local differential privacy on fairness: Preliminary results." 2024 IEEE 37th computer security foundations symposium (CSF). IEEE, 2024.
- [4] Makhlouf, Karima, et al. "On the impact of multi-dimensional local differential privacy on fairness." Data Mining and Knowledge Discovery 38.4 (2024): 2252-2275.
- [5] Bagdasaryan, Eugene, Omid Poursaeed, and Vitaly Shmatikov. "Differential privacy has disparate impact on model accuracy." NeurIPS 2019.

Assignment

Assignment. The recruited postdoc will collaborate with the project team to:
- Develop and rigorously evaluate new methodological frameworks that combine differential privacy and fairness-aware learning for real-world ML applications.
- Explore decentralized learning scenarios, including federated learning, to address privacy and fairness simultaneously when data must remain local.
- Theorize and quantify privacy–fairness–utility trade-offs under local privacy constraints, and contribute new LDP mechanisms that satisfy fairness guarantees.
- Translate research outputs into open-source building blocks, contributing to a broader DP-Fairness Toolkit for the research and practitioner community.

Research Topics. Possible directions include (but are not limited to):
- Design of new locally private mechanisms for high-dimensional or structured data that account for fairness constraints.
- Theoretical and empirical analysis of fairness mitigation steps in decentralized learning pipelines.
- Development of scalable prototypes and benchmarks for practical deployment in privacy-sensitive settings.

These outcomes will directly support the milestones of AI-PULSE and contribute broadly to the responsible AI research landscape.

Main activities

Main activities.
- Design and analyze new locally private and fairness-aware ML algorithms.
- Collaborate with Privatics members, and contribute to research visits to

international partners.
- Prototype methods and contribute to open-source code repositories.
- Publish and present results at top-tier international conferences.

Skills

Profile. We are looking for a candidate with:
- PhD in Computer Science, Statistics, or related fields.
- Strong background in one or more of the following: differential privacy, algorithmic fairness, or federated learning.
- Strong Python development skills; experience with open-source projects is valued.
- Strong analytical, communication, and collaboration skills.
- Proficiency in English.

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking (90 days / year) and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Complementary health insurance under conditions

Remuneration

2788€ gross salary / month

General Information

- **Theme/Domain :** Security and Confidentiality
  Statistics (Big data) (BAP E)
- **Town/city :** Montbonnot
- **Inria Center :** Centre Inria de l'Université Grenoble Alpes
- **Starting date :** 2025-11-01
- **Duration of contract :** 2 years
- **Deadline to apply :** 2025-08-31

Contacts

- **Inria Team :** PRIVATICS

- **Recruiter :**
  Hwang Arcolezi Heber / heber.hwang-arcolezi@inria.fr

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

Instruction to apply

Applications must be submitted online on the Inria website.

Processing of applications sent by other channels is not guaranteed.

Defence Security :
This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST).Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :
As part of its diversity policy, all Inria positions are accessible to people with disabilities.