PostDoc position - Inria n°2026-10132: Structured matrices for geometric computations

Summary

- **Start date:** September to December 2026
- **Duration:** 2-year full-time position (37.5 hours per week)
- **Gross salary:** 3,362€ gross / month (around 2,600€ / month after tax)
- **Location:** PariSanté Campus, 2-10 rue d'Oradour-sur-Glane, 75015 Paris, France
- **Remote Work:** Flexible arrangements available
- **Keywords:** geometric data, sparsity, low-rank approximation, SIMT programming

How to apply

First, **write us an email**: jean.feydy@inria.fr, alexis.glaunes@parisdescartes.fr and benjamin.charlier@inrae.fr. We'll be happy to discuss the project further and see how we could accomodate your constraints, career path and scientific interests.

Then, go through the **official recruitment process** at:
https://recrutement.inria.fr/public/classic/en/offres/2026-10132.
The start date is negotiable.

About the team

This postdoctoral position will be hosted within the **HeKA team at PariSanté Campus** and supervised by the **KeOps development team**: Jean Feydy (Inria, HeKA), Joan Glaunès (Université Paris Cité, MAP5) and Benjamin Charlier (INRAE, MIAT).

Based at PariSanté Campus, the HeKA team is a **multidisciplinary group** specializing in applied mathematics for clinical decision support. The team brings together researchers, clinician-scientists, and faculty members from Inria, Inserm, Université Paris Cité, and AP-HP.

Context

Massively parallel accelerators such as **Graphics Processing Units (GPUs)** now provide significant computational power at a fraction of the cost of a high-performance cluster. Providing **user-friendly libraries** that leverage these capabilities while remaining compatible with high-level development environments is essential for developping new methodological approach to analyze real-world datasets.

The **KeOps library (https://kernel-operations.io/) (1M+ downloads)** follows this approach and focuses on geometric computations based on the manipulation of distance and kernel matrices. These are widely used to compute interactions between large collections of samples, with applications that range from 3D shape processing to machine learning and computational physics.

KeOps introduces a high-level abstraction based on symbolic matrices (LazyTensors), offering a memory- and compute-efficient, transparent framework that is fully compatible with Python (NumPy, PyTorch) and R. We refer to this discussion for more details.

Assignment

The current implementation of KeOps provides efficient acceleration for dense, tensor-like operators. However, recent advances in biological and medical imaging, such as spatial omics, are producing a new class of **high-dimensional datasets with strong underlying structures**, including sparsity patterns. The goal of this postdoctoral project is therefore to move **beyond bruteforce dense computations** and leverage these structures to enable efficient processing of high-dimensional data.

Several research directions are possible, each tied to a specific application context. **Depending on your scientific profile and interests**, we will be able to focus on one or more of the following:

Strategy 1: sparse neighborhoods and spatial structures: generalize block-wise reduction schemes to handle **sparse matrices** which have few non-zero coefficients per row. Combined with the symbolic engine of the KeOps library, this could lead to very efficient routines for geometry processing.

Strategy 2: sparse, high-dimensional features: Address the computational limitations associated to high intrinsic dimensionality by developing sparse variants of the core operations used to process such datasets. A particular emphasis will be placed on efficient GPU implementations within the KeOps framework. Applications to spatial omics data can be found in [1].

Strategy 3: low rank kernel approximation: Investigate scalable approaches to kernel matrix-vector multiplication that rely on low-rank approximations, such as the *Fast and Free Memory method* explored in [2].

[1] xIV-LDDMM Toolkit: A Suite of Image-Varifold Based Technologies for Representing and Mapping 3D Imaging and Spatial-omics Data Simultaneously Across Scales. K. M. Stouffer et al. Nature Commun Biol

[2] Giga-scale Kernel Matrix-Vector Multiplication on GPU. R. Hu et al. NeurIps 2022

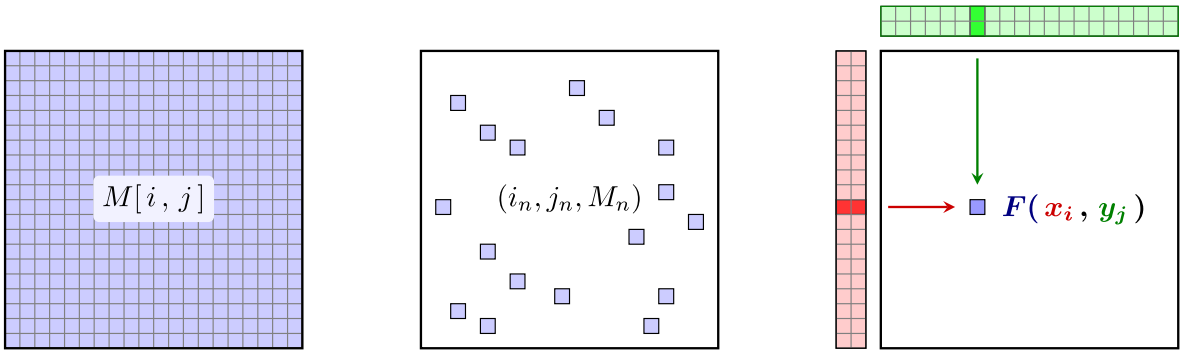

(a) Dense matrix (b) Sparse matrix (c) Symbolic matrix

Figure: Computing libraries represent most object as tensors. (a) Dense arrays are versatile, but have a large memory footprint and are costly to transfer between processing cores. (b) Sparse matrices alleviate this problem but do not stream well on GPUs as they rely on « random », non-contiguous memory transfers. (c) KeOps provides a convenient interface to manipulate **symbolic matrices**, whose coefficients are given by a formula that is evaluated lazily. It is now a popular Python extension to accelerate computations with distance and kernel matrices. The aim of this post-doc is to study **mathematical acceleration techniques** that could be applied to extend the KeOps framework, without compromising on ease-of-use and versatility.

Equipe HeKA
Health data- and model- driven approaches for Knowledge Acquisition
PariSanté Campus
10 Rue d'Oradour-sur-Glane, 75015 Paris, France

Requirements

- PhD in applied mathematics, computer science, physics, or a related field, with a strong interest in computational methods and real-world applications.
- Proficiency in Python and/or a compiled language (e.g., C++, CUDA).
- Strong background in at least one of the following research areas:
    - Computer graphics
    - Computational geometry or geometric data processing
    - Scientific computing and numerical methods
    - High-performance computing (HPC), GPU programming, or parallel computing
    - Machine learning and kernel methods
    - Optimization and large-scale linear algebra
    - Sparse and/or low-rank methods for large-scale data
    - Computational imaging or shape analysis
- Some experience with the open-source development ecosystem (version control, collaborative workflows, software design) would be a plus.

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Approximately 9 weeks of paid time off per year: 7 weeks of annual leave + 10 extra days off thanks to to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking 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
- Contribution to mutual insurance (subject to conditions)
- Gross Salary : 3,362 € per month

Funding

This position is funded by a 125k€ grant from the PR[AI]RIE-PSAI Institute (ANR-23-IACL-0008). It is hosted by the HeKA team (Inria Paris, Inserm, Université Paris-Cité).

Equipe HeKA
Health data- and model- driven approaches for Knowledge Acquisition
PariSanté Campus
10 Rue d'Oradour-sur-Glane, 75015 Paris, France