Reconceptualizing Provenance for Reproducibility in Agentic AI Scientific Workflows

Location	PSL Research University - LAMSADE Université Paris-Dauphine, France
Supervisor	Khalid Belhajjame
Contact	khalid.belhajjame@dauphine.fr
Candidate profile	Master in Computer Science and/or Bioinformatics.
	Proficiency in algorithmic design and programming is essential.

Context and Motivation

Scientific workflows have become a cornerstone of computational science, enabling analysts to express complex pipelines as structured compositions of steps, tools, and data transformations. Provenance models, most notably those grounded in the W3C PROV standard, support repeatability by recording workflow structure, execution traces, data lineage, and aspects of the computational environment.

However, the emergence of agent-augmented workflows, where LLM-based or human agents influence execution through context-sensitive choices, changes what must be captured to sustain reproducibility. In hybrid workflows, some steps are not mere deterministic transformations: they interpret intermediate results, select parameters, decide whether to branch or filter data, and sometimes synthesize conclusions. These agent-mediated decisions can alter downstream computations and, crucially, the scientific claims derived from outputs. Classical execution provenance can explain what happened, but it often fails to explain why it happened and under which decision conditions the results would be stable.

This thesis addresses that gap by developing a provenance-centered approach to reproducibility that explicitly represents decisions, reasoning, and transferable methodological knowledge in hybrid scientific workflows.

Research Topic and Core Questions

The thesis will investigate how to extend provenance models for hybrid workflows so that they support not only replaying computations but also comparing decisions, validating reasoning, and reusing accumulated analytical knowledge. The central question is how to design a provenance model that remains PROV-compliant while making agentic control explicit and analyzable across runs.

This involves representing decision-making events and their contextual evidence, capturing reasoning traces that relate evidence to scientific claims, and extracting higher-level knowledge—such as reusable decision patterns. A key aim is to align these provenance layers with the classical reproducibility spectrum (repeatability, replicability, reproducibility, reuse) without redefining that spectrum, but by making its provenance requirements explicit in hybrid settings.

Objectives

The PhD will pursue the following objectives:

1. Formalize a stratified provenance model for hybrid workflows, structured into cumulative layers capturing execution traces, decision traces, reasoning traces, and cross-run knowledge traces.
2. Operationalize reproducibility through layered provenance by defining how each provenance layer supports a level of the reproducibility spectrum and how traces can be compared across independent runs.
3. Design mechanisms for extracting reusable methodological knowledge, such as decision patterns and applicability conditions derived from accumulated workflow executions.
4. Develop and validate a prototype implementation, applying the framework to realistic hybrid scientific workflows (e.g., in computational biology or other data-intensive domains).

Required Skills and Profile

Candidates should have a strong background in computer science and/or bioinformatics, with solid programming skills and interest in research at the intersection of data management and AI.

The topic involves both conceptual modeling and system prototyping; candidates should be comfortable with abstraction, formal reasoning, and implementation work. Prior exposure to scientific workflows or reproducibility research is a plus but not required.

How to Apply

Interested candidates should contact Khalid Belhajjame at:
`khalid.belhajjame@lamsade.dauphine.fr`
The application should include:

- A detailed CV
- Academic transcriptsof the last three years.
- A short statement of interest (approximately one page) describing motivation and relevant experience
- A copy of the Master's thesis (if available)
- Contact information for two referees
- Optionally, links to publications, technical reports, or code repositories

Candidates are encouraged to briefly highlight any prior experience related to provenance, workflows, reproducibility, or AI-based systems.