

2025 Activity Report

RESEARCH CENTRE: Inria Centre at Université Grenoble Alpes
IN PARTNERSHIP WITH: CNRS, Université de Grenoble Alpes



Project-Team

TYREX

Types and Reasoning for the Web



In collaboration with Laboratoire d'Informatique de Grenoble (LIG)



Project-Team TYREX

Creation of the Project-Team: 2014 July 01

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

Keywords

Computer sciences and digital sciences

- A2.1.1. – Semantics of programming languages
- A2.1.4. – Functional programming
- A2.1.10. – Domain-specific languages
- A2.2.1. – Static analysis
- A2.2.8. – Code generation
- A2.5.1. – Software Architecture & Design
- A3.1.1. – Modeling, representation
- A3.1.2. – Data management, quering and storage
- A3.1.4. – Uncertain data
- A3.1.6. – Query optimization
- A3.1.9. – Database
- A3.1.11. – Structured data
- A3.2.1. – Knowledge bases
- A3.2.2. – Knowledge extraction, cleaning
- A3.2.3. – Inference
- A3.2.5. – Ontologies
- A3.2.6. – Linked data
- A3.3.3. – Big data analysis
- A3.4. – Machine learning and statistics
- A3.5.1. – Analysis of large graphs
- A4.5. – Formal method for verification, reliability, certification
- A6.3.3. – Data processing
- A7. – Theory of computation
- A7.1. – Algorithms
- A7.2. – Logic in Computer Science
- A9.1. – Knowledge
- A9.2. – Machine learning
- A9.2.1. – Supervised learning
- A9.7. – AI algorithmics
- A9.8. – Reasoning
- A9.10. – Hybrid approaches for AI
- A9.11. – Generative AI
- A9.14. – Evaluation of AI models
- A9.15. – Symbolic AI

Other research topics and application domains

- B2. – Digital health
- B6.1. – Software industry
- B6.5. – Information systems
- B8. – Smart Cities and Territories
- B9.5.1. – Computer science
- B9.5.6. – Data science
- B9.7.2. – Open data
- B9.8. – Reproducibility
- B9.11.2. – Financial risks

Contents

Project-Team TYREX	1
1 Team members, visitors, external collaborators	5
2 Overall objectives	5
2.1 Objectives	5
3 Research program	6
3.1 Algebraic Foundations for Robust Expressive and Efficient Information Extraction	6
3.2 Neuro-Symbolic Programming	6
4 Application domains	6
4.1 Querying Large Graphs	6
4.2 Predictive Analytics for Healthcare	6
5 Social and environmental responsibility	6
5.1 Impact of research results	6
6 Latest software developments, platforms, open data	7
6.1 Latest software developments	7
6.1.1 MuIR	7
6.1.2 KeGNN	7
6.1.3 Reproducibility-aaai24	7
6.1.4 MedAnalytics	7
7 New results	8
7.1 Foundations of next-generation data management systems	8
7.2 Neurosymbolic AI	9
8 Bilateral contracts and grants with industry	10
8.1 Bilateral contracts with industry	10
9 Partnerships and cooperations	11
9.1 International initiatives	11
9.1.1 Visits to international teams	11
9.2 National initiatives	11
9.2.1 ANR	11
10 Dissemination	11
10.1 Promoting scientific activities	11
10.1.1 Scientific events: selection	11
10.1.2 Scientific expertise	11
10.1.3 Research administration	11
10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach	12
10.2.1 Supervision	12
10.2.2 Juries	12
11 Scientific production	12
11.1 Major publications	12
11.2 Publications of the year	13

1 Team members, visitors, external collaborators

Research Scientists

- Pierre Genevès [Team leader, CNRS, Senior Researcher, HDR]
- Nabil Layaïda [INRIA, Senior Researcher, HDR]
- Chandan Sharma [INRIA, Starting Research Position]

Faculty Members

- Ugo Comignani [GRENOBLE INP, Associate Professor]
- Nils Gesbert [GRENOBLE INP, Associate Professor]

Post-Doctoral Fellow

- Luisa Werner [INRIA, Post-Doctoral Fellow]

PhD Students

- Richard Casetta [BNP PARIBAS , CIFRE]
- Guillaume Delplanque [UGA]
- Maroua Zeblah [OPENSEE SAS, CIFRE]

Technical Staff

- Sarah Chlyah [INRIA, Engineer]

Administrative Assistant

- Helen Pouchot-Rouge-Blanc [INRIA]

External Collaborator

- Laurent Carcone [W3C - ERCIM]

2 Overall objectives

2.1 Objectives

We develop the foundations for the next generation of information extraction, data analysis and neuro-symbolic programming systems. Our research extends ideas from data management, artificial intelligence, programming languages and logic.

Extracting value from data increasingly requires sophisticated algorithms to represent, query, process, analyze and interpret data. We develop the foundations of data processing systems and neuro-symbolic programming, with a focus on extracting information from graph structures. These graph structures are obtained from raw data that may be more or less structured, noisy, uncertain or incomplete. Challenges include robust, efficient and scalable processing of large graphs obtained from such data. We study and build new information extraction methods, as well as new robust and scalable programming methods for rich graph data structures.

3 Research program

3.1 Algebraic Foundations for Robust Expressive and Efficient Information Extraction

We investigate intermediate languages based on algebraic foundations for the representation, characterization, transformations and compilation of queries. We develop the algebraic and logical foundations of advanced data programming languages (extended relational algebras, algorithms, compilers) for more expressive and efficient query languages, in particular through aspects such as recursion, types, analytics, and provenance.

3.2 Neuro-Symbolic Programming

We investigate neuro-symbolic programming methods with graphs. This includes studying the integration between neural networks and symbolic logic and/or algebra. Challenges include bridging the gap between neural networks and symbolic logic, injecting knowledge in learning processes, supporting rich knowledge and property graphs, and dealing with scalability issues for large graphs.

4 Application domains

4.1 Querying Large Graphs

The increasing availability of large-scale graph-structured data presents both opportunities and challenges. Our research focuses on efficient methods for evaluating graph queries at scale, particularly in knowledge graphs structured in the Resource Description Framework (RDF) and property graphs.

We design advanced query languages to extract insights from these graphs and compile queries into algebraic representations. These representations are then translated into executable low-level code, optimized for various backends, including relational database management systems like PostgreSQL, and big data frameworks like Apache Spark.

Graph querying has applications across diverse domains, including large knowledge bases, social networks, road networks, trust and fraud detection in cryptocurrencies, citation and web graphs, and recommendation systems.

4.2 Predictive Analytics for Healthcare

A major expectation of data science in healthcare is the ability to leverage digitized health information and computer systems to better apprehend and improve care. The availability of clinical data and in particular electronic health records opens the way to the development of models for patients that can be used to predict health status, as well as to help prevent disease and adverse effects.

In collaboration with the Grenoble University Hospital (CHUGA), we explore solutions to the problem of predicting important clinical outcomes such as risks of adverse effects, nosocomial infections or inpatient mortality, based on large amounts of clinical data.

5 Social and environmental responsibility

5.1 Impact of research results

Our work on graph query optimization helps in reducing resource consumption in information extraction. Our work in neuro-symbolic programming helps in reducing the amount of data required when training accurate artificial intelligence models, thanks to the integration of symbolic concepts and reasoning rules.

6 Latest software developments, platforms, open data

6.1 Latest software developments

6.1.1 MuIR

Name: Mu Intermediate Representation System

Keywords: Optimizing compiler, Querying

Functional Description: This is a prototype of an intermediate language representation, i.e. an implementation of algebraic terms, rewrite rules, query plans, cost model, query optimizer, and query evaluators. This includes query evaluators for a variety of RDBMS backends including PostgreSQL as well a distributed evaluator of algebraic terms using Apache Spark. This also includes an implementation of an efficient enumerator for recursive query plans, cost estimations, and compilers for recursive graph queries. The overall system is described in the CIKM 2023 demonstration paper.

Publications: [hal-01673025](#), [hal-03295445](#), [hal-03004218](#), [hal-03517826](#)

Contact: Pierre Genevès

6.1.2 KeGNN

Name: Knowledge Enhanced Graph Neural Networks

Functional Description: We propose KeGNN, a neuro-symbolic framework for learning on graph data that combines both paradigms and allows for the integration of prior knowledge into a graph neural network model. In essence, KeGNN consists of a graph neural network as a base on which knowledge enhancement layers are stacked with the objective of refining predictions with respect to prior knowledge. We instantiate KeGNN in conjunction with two standard graph neural networks: Graph Convolutional Networks and Graph Attention Networks, and evaluate KeGNN on multiple benchmark datasets for node classification.

URL: <https://gitlab.inria.fr/tyrex-public/kegnn>

Publication: [hal-04041691](#)

Contact: Pierre Genevès

6.1.3 Reproducibility-aaai24

Functional Description: This is a re-implementation of the experiments conducted with Knowledge Enhanced Neural Networks (KENN) on the Citeseer Dataset, including the re-implementation of the Experiments in PyTorch and PyTorch Geometric. We also extended the experiments to the datasets Cora and PubMed.

URL: <https://gitlab.inria.fr/tyrex-public/reproducibility-aaai24>

Publication: [hal-04035305](#)

Contact: Pierre Genevès

6.1.4 MedAnalytics

Keywords: Big data, Predictive analytics, Distributed systems

Functional Description: We implemented a method for the automatic detection of at-risk profiles based on a fine-grained analysis of prescription data at the time of admission. The system relies on an optimized distributed architecture adapted for processing very large volumes of medical records and clinical data. We conducted practical experiments with real data of millions of patients and hundreds of hospitals. We demonstrated how the various perspectives of big data improve the detection of at-risk patients, making it possible to construct predictive models that benefit from volume and variety.

Publications: [hal-01517087](#), [hal-01877742](#), [hal-03124966](#), [hal-03125018](#), [hal-03160473](#), [hal-03066941](#), [hal-03266004](#)

Contact: Pierre Genevès

Partner: CHU Grenoble

7 New results

7.1 Foundations of next-generation data management systems

Schema-Based Query Optimisation for Graph Databases

Participants: Chandan Sharma, Pierre Genevès, Nils Gesbert, Nabil Layaïda.

Recursive graph queries are increasingly popular for extracting information from interconnected data found in various domains such as social networks, life sciences, and business analytics. Graph data often come with schema information that describe how nodes and edges are organized. We propose a type inference mechanism that enriches recursive graph queries with relevant structural information contained in a graph schema. We show that this schema information can be useful in order to improve the performance when evaluating recursive graph queries. Furthermore, we prove that the proposed method is sound and complete, ensuring that the semantics of the query is preserved during the schema-enrichment process. Experimental results with a complete implementation of the approach show significant performance gains for query evaluations over property graphs, with several evaluation backend. These results were presented at SIGMOD 2025 in Berlin [7].

Distributed Evaluation of Graph Queries Using Recursive Relational Algebra

Participants: Sarah Chlyah, Pierre Genevès, Nabil Layaïda.

We present a system, Dist- μ -RA, for the distributed evaluation of recursive graph queries. Dist- μ -RA builds on the recursive relational algebra [2, 1] and extends it with evaluation plans suited for the distributed setting. The goal is to offer expressivity for high-level queries while providing efficiency at scale and reducing communication costs. Specifically, we propose a new approach for the evaluation of recursive algebraic terms in a distributed manner. The method enables generating independent parallel loops on the worker nodes in a cluster of machines instead of executing a global loop on the driver node. The advantage of the parallel local loops is a minimization of the amount of data shuffled between worker nodes. This reduces communication costs and significantly improves overall query evaluation time. We applied this approach to recursive graph queries on real and synthetic datasets. Experimental results on both real and synthetic graphs show the effectiveness of the proposed approach compared to existing systems. These results were presented at ICDE 2025 in Hong Kong [5] [6.1.1].

Efficient Iterative Programs with Distributed Data Collections

Participants: Sarah Chlyah, Nils Gesbert, Pierre Genevès, Nabil Layaïda.

Big data programming frameworks have become increasingly important for the development of applications for which performance and scalability are critical. In those complex frameworks, optimizing code by hand is hard and time-consuming, making automated optimization particularly necessary. In order to automate optimization, a prerequisite is to find suitable abstractions to represent programs; for instance, algebras based

on monads or monoids to represent distributed data collections. Currently, however, such algebras do not represent recursive programs in a way which allows for analyzing or rewriting them. In this paper, we extend a monoid algebra with a fixpoint operator for representing recursion as a first class citizen and show how it enables new optimizations. Experiments with the Spark platform illustrate performance gains brought by these systematic optimizations. These results have been published in the Journal of Logical and Algebraic Methods in Programming [3].

An Enterprise Marketplace for Unified Access to Multi-Cloud and Enterprise Products in a Large Banking Infrastructure

Participants: Richard Casetta, Nils Gesbert, Pierre Genevès.

We present the design and evaluation of an Enterprise Marketplace that unifies web-based access to multi-cloud and enterprise products within a large banking infrastructure. Building such a platform poses significant technical and organizational challenges, including the generalization of diverse APIs and the accommodation of heterogeneous user profiles. Our solution enables autonomous product publishing via a no-code interface, enforces multi-layered governance to ensure security and compliance, and integrates disparate providers through a standardized API. We demonstrate how this approach enhances product quality, producer autonomy, and user experience, supported by adoption metrics and operational data from a real-world deployment. Finally, we reflect on key lessons learned and persistent challenges after a decade of production use, serving over 200,000 users. These results will be presented at ICSE 2026 in Rio de Janeiro [4].

COSMetyc: OpenStreetMap in OCaml

Participants: Ugo Comignani.

We introduce COSMetyc, an OCaml library for manipulating OpenStreetMap (OSM) data, a large-scale collaborative geographic database. Given the heterogeneity of usages, formats, and representations in the OSM ecosystem, COSMetyc adopts a modular design that leverages OCaml's type system to statically guarantee data validity. The library supports importing and exporting data from multiple formats (notably GeoJSON and OSM XML) through several typed representations, provides conversions between coordinate systems, and enables efficient spatial queries. We report on the design choices underlying the library and show how OCaml features such as functors, phantom types, and polymorphic variants can be used to manage this complexity in a principled and scalable way [8].

7.2 Neurosymbolic AI

A Comparative Analysis of Neuro-symbolic Methods for Link Prediction

Participants: Guillaume Delplanque, Pierre Genevès, Luisa Werner, Nabil Layaida.

Link prediction on knowledge graphs is relevant to various applications, such as recommendation systems, question answering, and entity search. This task has been approached from different perspectives: symbolic methods leverage rule-based reasoning but struggle with scalability and noise, while knowledge graph embeddings (KGE) represent entities and relations in a continuous space, enabling scalability but often neglecting logical constraints from ontologies. Recently, neurosymbolic approaches have emerged to bridge this gap by integrating embedding-based learning with symbolic reasoning. This paper provides a structured review of state-of-the-art neurosymbolic methods for link prediction. Beyond a qualitative analysis, a key contribution of this work is a comprehensive experimental benchmarking, where we systematically compare

these methods on the same datasets using the same metrics. This unified experimental setup allows for a fair assessment of their strengths and limitations, bringing elements of answers to following key questions: How accurate are these methods? How scalable are they? How beneficial are they for different levels of provided knowledge and to which extent are they robust to incorrect knowledge? These results have been presented at the NeSy 2025 conference in Santa Cruz [6].

On Scaling Neurosymbolic Programming through Guided Logical Inference

Participants: Thomas Valentin, Pierre Genevès, Luisa Werner, Sarah Chlyah.

Probabilistic neurosymbolic learning seeks to integrate neural networks with symbolic programming. Many state-of-the-art systems rely on a reduction to the Probabilistic Weighted Model Counting Problem (PWMC), which requires computing a Boolean formula called the logical provenance. However, PWMC is #P-hard, and the number of clauses in the logical provenance formula can grow exponentially, creating a major bottleneck that significantly limits the applicability of PNL solutions in practice. We propose a new approach centered around an exact algorithm DPNL, that enables bypassing the computation of the logical provenance. The DPNL approach relies on the principles of an oracle and a recursive DPPL-like decomposition in order to guide and speed up logical inference. Furthermore, we show that this approach can be adapted for approximate reasoning guarantees, called ApproxDPNL. Experiments show significant performance gains. In particular, DPNL enables scaling exact inference further, resulting in more accurate models [9].

OntoKGE: A Framework for Injecting Ontology Rules into Knowledge Graph Embedding Training.

Participants: Luisa Werner, Nabil Layaïda, Pierre Genevès.

Link prediction is a key task in knowledge graphs that involves inferring new links between entities based on existing ones. Knowledge graph embedding models address this task by representing entities and relations as points in a geometric space and using distance or similarity functions to predict missing links. However, traditional knowledge graph embedding models are trained only on assertional facts and ignore semantic information, encoded for example in the ontology. We aim to improve the quality of knowledge graph embedding models for link prediction by leveraging prior knowledge. To this end, we propose OntoKGE, an embedding-agnostic framework that integrates a reasoning module into the training process of knowledge graph embedding models. This module is based on Datalog and derives additional training facts, enriching the training set and enhancing the learned embeddings. Experiments on multiple benchmarks show that OntoKGE improves link prediction performance across multiple knowledge graph embedding models [10].

8 Bilateral contracts and grants with industry

8.1 Bilateral contracts with industry

Participants: Pierre Genevès, Maroua Zeblah, Richard Casetta, Nils Gesbert, Sarah Chlyah.

We collaborate with BNP Paribas in Paris, a major international financial group, on the development of logical and algebraic methods to support the design and verification of robust cloud architectures, within the framework of a CIFRE-funded PhD thesis.

In addition, we work with a Paris-based French fintech startup on query optimization techniques for multidimensional data, also through a CIFRE-funded PhD thesis.

9 Partnerships and cooperations

9.1 International initiatives

9.1.1 Visits to international teams

Research stays abroad Luisa Werner carried out a short research visit to Luc De Raedt’s group at KU Leuven.

9.2 National initiatives

9.2.1 ANR

GraphRec

Participants: Pierre Genevès, Nabil Layaïda, Nils Gesbert, Sarah Chlyah, Ugo Comignani, Luisa Werner, Chandan Sharma.

- Title: GraphRec: Efficient and Scalable Recursive Programming with Graphs
- ANR, Appel à projets générique 2023 – CE23 – Intelligence artificielle et science des données, PRME
- Coordinator: Pierre Genevès
- Abstract: This project seeks to design and develop novel methods for expressive and efficient information extraction from graphs, based on recursive graph queries and neuro-symbolic programming.
- **GraphRec website:** <https://tyrex.inria.fr/graphrec>

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: selection

Member of the conference program committees

- Pierre Genevès has been PC member of SIGMOD 2025.
- Ugo Comignani has been PC member of SIGMOD 2025 and IEEE BigData 2025.

10.1.2 Scientific expertise

Pierre Genevès has been referee for the Agence Nationale de la recherche (ANR) and Agence Nationale de la Recherche et de la Technologie (ANRT), in charge of reviewing research project proposals.

Pierre Genevès has been expert reviewer for the Qatar National Research Fund.

Chandan Sharma and Pierre Genevès have been expert reviewers for the National Agency for Research and Development (ANID), Ministry of Science, Technology, Knowledge and Innovation, Chile.

10.1.3 Research administration

Pierre Genevès is responsible for the Computer Science Specialty at the MSTII Doctoral School of University Grenoble Alpes (ED 217).

Pierre Genevès is member of the board at Grenoble Informatics Laboratory (LIG), responsible for the research axis on formal methods, models and languages regrouping 4 research teams (CAPP, CONVECS, SPADES, TYREX).

10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

- Master: P. Genevès is co-responsible and teacher of the M2-level course “Accès à l’information: du web des données au web sémantique” in the ENSIMAG ISI 3A program at Grenoble-INP (30h)
- Master : N. Gesbert, “Analyse et Conception Objet de Logiciels”, 30 h eq TD, M1, Grenoble INP
- Master : N. Gesbert, “Construction d’applications Web”, 27 h eq TD, M1, Grenoble INP
- Master : N. Gesbert, “Principes des systèmes de gestion des bases de données”, 54 h eq TD, M1, Grenoble INP
- Licence : N. Gesbert, “Logique pour l’informatique”, 45 h eq TD, L3, Grenoble INP
- Licence : N. Gesbert is in charge of the L3-level course “logique pour l’informatique” at Grenoble INP Ensimag.
- N. Gesbert is responsible of the pedagogical team “Gestion de données” at Grenoble INP Ensimag.
- Master : U. Comignani is co-responsible of the “BigData” master, co-accredited between Grenoble Ecole de Management and Grenoble INP
- Master : U. Comignani is in charge of the “Projets fil rouge”, 10 h eq TD, MS BigData, Grenoble INP
- Master : U. Comignani, “Principes des systèmes de gestion de bases de données”, 99.5 h eq TD, M1, Grenoble INP
- Master : U. Comignani is in charge of the “Projet BD”, 64 h eq TD, M1, Grenoble INP
- Master : U. Comignani, “Stockage et traitement de données à grande échelle”, 34 h eq TD, M2, Grenoble INP
- Master : U. Comignani, academic tutorship of an apprentice, 10 h eq TD, M1, Grenoble INP

10.2.1 Supervision

- PhD in progress: Maroua Zeblah, Query Optimisation for column oriented databases, PhD started in April 2023, co-supervised by Pierre Genevès and Nabil Layaïda.
- PhD in progress: Guillaume Delplanque, Differentiable programming for Knowledge Graphs, PhD started in September 2023, co-supervised by Pierre Genevès and Nabil Layaïda.
- PhD in progress: Richard Casetta, Formal verification of cloud applications, PhD started in 2024, co-supervised by Nils Gesbert and Pierre Genevès.

10.2.2 Juries

Pierre Genevès has been president of the jury for the PhD thesis of Hadi Dayekh, entitled “Passive and Active Learning of Switched Nonlinear Dynamical Systems”, Université Grenoble Alpes, and defended in April 2025. <https://theses.hal.science/tel-05073907>

11 Scientific production

11.1 Major publications

- [1] A. Fejza, P. Genevès and N. Layaïda. ‘Efficient Enumeration of Recursive Plans in Transformation-based Query Optimizers’. In: *Proceedings of the VLDB Endowment (PVLDB)* 17.11 (7th July 2024), pp. 3095–3108. DOI: [doi:10.14778/3681954.3681986](https://doi.org/10.14778/3681954.3681986). URL: <https://inria.hal.science/hal-03692274> (cit. on p. 8).

- [2] L. Jachiet, P. Genevès, N. Gesbert and N. Layaïda. ‘On the Optimization of Recursive Relational Queries: Application to Graph Queries’. In: *SIGMOD 2020 - ACM International Conference on Management of Data*. Portland, United States, June 2020, pp. 1–23. doi: [10.1145/3318464.3380567](https://doi.org/10.1145/3318464.3380567). URL: <https://hal.inria.fr/hal-01673025> (cit. on p. 8).

11.2 Publications of the year

International journals

- [3] S. Chlyah, N. Gesbert, P. Genevès and N. Layaïda. ‘Efficient Iterative Programs with Distributed Data Collections’. In: *Journal of Logical and Algebraic Methods in Programming* 144 (Mar. 2025), pp. 1–36. doi: [10.1016/j.jlamp.2025.101047](https://doi.org/10.1016/j.jlamp.2025.101047). URL: <https://inria.hal.science/hal-04108082> (cit. on p. 9).

International peer-reviewed conferences

- [4] R. Casetta, T. Brisbout, J.-F. Tur, M. Barry, J. Veybel, J.-M. Garcia, N. Gesbert and P. Genevès. ‘An Enterprise Marketplace for Unified Access to Multi-Cloud and Enterprise Products in a Large Banking Infrastructure’. In: *Proceedings of the International Conference on Software Engineering (ICSE)*. ICSE 2026 - IEEE/ACM International Conference on Software Engineering- SEIP. Rio de Janeiro, Brazil, 12th Apr. 2026, pages 1–11. doi: [10.1145/3786583.3786914](https://doi.org/10.1145/3786583.3786914). URL: <https://inria.hal.science/hal-05462351> (cit. on p. 9).
- [5] S. Chlyah, P. Genevès and N. Layaïda. ‘Distributed Evaluation of Graph Queries using Recursive Relational Algebra’. In: *ICDE 2025 - 41 st IEEE International Conference On Data Engineering*. Hong Kong, China, 19th May 2025. URL: <https://inria.hal.science/hal-03295445> (cit. on p. 8).
- [6] G. O. Delplanque, L. S. Werner, N. Layaïda and P. Genevès. ‘A Comparative Analysis of Neuro-symbolic Methods for Link Prediction’. In: *Proceedings of Machine Learning Research (PMLR)*, vol. 284, pages 674–696, 2025. NeSy 2025 - Proceedings of the 19th International Conference on Neurosymbolic Learning and Reasoning. Santa Cruz, California, United States, 6th June 2025. URL: <https://hal.science/hal-04991391> (cit. on p. 10).
- [7] C. Sharma, P. Genevès, N. Gesbert and N. Layaïda. ‘Schema-Based Query Optimisation for Graph Databases’. In: *Proceedings of the ACM on Management of Data*. SIGMOD 2025 - ACM International Conference on Management of Data. Vol. 3. 1. Berlin, Germany: ACM, 11th Feb. 2025, pp. 1–29. doi: [10.1145/3709722](https://doi.org/10.1145/3709722). URL: <https://inria.hal.science/hal-04485125> (cit. on p. 8).

National peer-reviewed Conferences

- [8] T. Baleras, M. Bodin and U. Comignani. ‘COSMetyc : OpenStreetMap en OCaml’. In: *JFLA 2026 – 37es Journées Francophones des Langages Applicatifs*. Vol. JFLA 2026 – 37es Journées Francophones des Langages Applicatifs. Oberbronn, Alsace, France, 27th Jan. 2026. URL: <https://hal.science/hal-05427963> (cit. on p. 9).

Reports & preprints

- [9] T. J.-M. Valentin, L. S. Werner, P. Genevès and N. Layaïda. *On Scaling Neurosymbolic Programming through Guided Logical Inference*. 2025. URL: <https://hal.science/hal-04914004> (cit. on p. 10).
- [10] L. S. Werner, P. Genevès and N. Layaïda. *OntoKGE: A Framework for Injecting Ontology Rules into Knowledge Graph Embedding Training*. 14th May 2025. URL: <https://inria.hal.science/hal-05075151> (cit. on p. 10).