

2025 Activity Report


RESEARCH CENTRE: Inria Branch at the University of Montpellier
IN PARTNERSHIP WITH: Université de Montpellier, INRAE

Project-Team

BOREAL

Knowledge Representation and Rule-Based
Languages for Reasoning on Data

In collaboration with Laboratoire d'informatique, de robotique et de microélectronique de Montpellier (LIRMM), Ingénierie des Agropolymères et Technologies Emergentes (IATE)



Project-Team BOREAL

Creation of the Project-Team: 2022 June 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

- A3.1. – Data
- A3.2. – Knowledge
- A7.1.3. – Graph algorithms
- A7.2. – Logic in Computer Science
- A9. – Artificial intelligence
 - A9.1. – Knowledge
 - A9.8. – Reasoning

Other research topics and application domains

- B3.5. – Agronomy
- B6.5. – Information systems

Contents

Project-Team BOREAL	1
1 Team members, visitors, external collaborators	4
2 Overall objectives	5
3 Research program	5
4 Application domains	6
4.1 Agronomy and agroecology	6
5 Highlights of the year	7
6 Latest software developments, platforms, open data	7
6.1 Latest software developments	7
6.1.1 InteGraal	7
6.1.2 TreeForce	8
6.1.3 B-Runner	8
6.1.4 IRIRefs	9
6.1.5 Py4Graal	9
6.1.6 DLGPE	10
6.1.7 NanoParse	10
7 New results	11
7.1 Preliminaires about Knowledge-Based Data Management with Existential Rules	11
7.2 Contributions	12
7.3 Foundations of Databases and Knowledge Bases	12
7.4 Applications of Rule-Based Reasoning	13
8 Partnerships and cooperations	14
8.1 International initiatives	14
8.1.1 Participation in other International Programs	14
8.2 International research visitors	15
8.2.1 Visits of international scientists	15
8.3 National initiatives	19
9 Dissemination	19
9.1 Promoting scientific activities	19
9.1.1 Scientific events: organization	19
9.1.2 Scientific events: selection	20
9.1.3 Journal	20
9.1.4 Scientific expertise	20
9.1.5 Research administration	20
9.2 Teaching - Supervision - Juries - Educational and pedagogical outreach	21
9.2.1 Supervision	21
9.2.2 Juries	21
10 Scientific production	21
10.1 Major publications	21
10.2 Publications of the year	22
10.3 Cited publications	23

1 Team members, visitors, external collaborators

Research Scientists

- Federico Ulliana [Team leader, University of Montpellier, Associate Professor Detachment, until Aug 2025]
- Federico Ulliana [Team leader, University of Montpellier, Researcher, from Sep 2025]
- Jean-Francois Baget [Inria, Researcher]
- Pierre Bisquert [INRAE, Researcher]
- Nofar Carmeli [Inria, Researcher]
- David Carral Martinez [Inria, Researcher]

Faculty Members

- Michel Chein [University of Montpellier, Emeritus, HDR]
- Michel Leclère [University of Montpellier, Associate Professor]
- Marie-Laure Mugnier [University of Montpellier, Professor, HDR]

Post-Doctoral Fellow

- Guillaume Perution-Kihli [Inria, Post-Doctoral Fellow, until Aug 2025]

PhD Student

- Akira Charoensit [Inria]

Interns and Apprentices

- Ibrahim Al Ayoubi [Inria, Intern, from Jun 2025 until Jul 2025]
- Carole Beaugeois [Inria, Intern, from Jun 2025 until Jul 2025]
- Jeanne Coschieri [Inria, Intern, from Jun 2025 until Aug 2025]
- Abir Amina Hammoud [Inria, Intern, from Jun 2025 until Jul 2025]
- Bastien Schmitt [Inria, Intern, from Mar 2025 until Jun 2025]

Administrative Assistant

- Sandrine Boute [Inria]

External Collaborators

- Patrice Buche [INRAE, HDR]
- Maxime Buron [Clermont Auvergne University]
- Alain Gutierrez [CNRS]

2 Overall objectives

Current information systems are grounded on the exploitation of data coming from an increasing number of heterogeneous sources. Today, coping with the variety of data requires novel paradigms for effectively accessing and querying information that adapt to the different types of sources, as well as declarative high-level languages to drive the data processing and data quality tasks.

BOREAL is a team working at the crossroads of *knowledge representation and reasoning* and *database theory*. The team focuses on the study of foundational and applied issues of *reasoning* in a context of *data variety*. More specifically, the team aims at deriving a better understanding of the logical fragments that are at the foundations of the frameworks used for exploiting corporate and Web data - and in particular *rule-based languages*. This will pave the way to novel automated-reasoning and graph-based techniques that can be put at service of data-centric applications exploiting heterogeneous and federated data. The team also aims at combining solid foundational and algorithmic work with software development and applications, with an emphasis on the field of agronomy.

3 Research program

The BOREAL team pursues a *knowledge-based data management* (KBDM) approach for tackling the grand challenges posed by data variety, with an important focus on the framework of *existential rules*. The idea of knowledge-based data management is to orchestrate the access to a complex information system made by federated databases through a three-layer architecture - also common to data-integration and ontology-based data access (OBDA). Under this prism, a set of *heterogeneous data sources* is connected to a *knowledge base* via a layer of *mappings*. The idea of KBDM is to define the business logic for data-centric applications at the knowledge base level, and then automatically translate the data-services towards the heterogeneous sources - through reasoning. This approach paves the way to a more principled use of complex information systems, with benefits to both data scientists, data curators, and administrators. What really characterizes the KBDM approach is the leveraging on formalized *domain-specific knowledge*, for abstracting on heterogeneous data and achieving high-quality of data-integration, and on expressive *rule-base languages* like existential rules (and extensions thereof), to drive the effective exploitation of data through reasoning.

Our project focuses on a set of topics related to knowledge-based data management, which we now describe.

Foundations of rule languages A great deal of the power of a KBDM system comes from its rule base. A prominent research direction for the team is the analysis and design of rule languages for reasoning on data. It is well understood that enriching a language with novel features can sensibly increase the complexity of the reasoning tasks. Our goal is hence to identify rules featuring decidable query answering and static analysis, and at the same time find good tradeoffs between their expressivity and complexity, so as to devise novel and practically useful rule-based frameworks.

Algorithms and optimizations for query answering Reasoning-driven data management needs optimization to effectively exploit large data. We target the design of efficient and scalable algorithms for query answering. Our goal is to devise novel hybrid approaches that combine materialization and virtualization strategies and account for the interplay between the components of the KBDM system (data, mappings, rules). Our ambition is also to build new bridges between knowledge representation and data-management by exploring the range of possibilities opened by the reuse of existing database technology to develop new reasoning systems.

Fine-grained complexity of query answering The query answering problem is at the heart of many reasoning tasks in KBDM. From a complexity analysis point of view, since the database to query can be voluminous, it is not always enough to know that a certain task can be done in polynomial time. Hence, an important goal for us is to study the fine-grained complexity (that is, to find the degree of the polynomial that bounds the number of operations required) as well as the enumeration complexity of the query answering problem. The aim of this research direction is to obtain the theoretical knowledge required for practical query optimization.

Architectures for knowledge-based data integration The realm of possibilities in heterogeneous data integration leads to the offspring of a family of KBDM architectures, one for each applicative context. Our goal is to study architectures inspired from emerging practical use-cases, including *federations* of independent sources as well as *multi-level* architectures where KBDM systems are stacked to progressively distill information and achieve high-value data. We also focus on the type of mappings required to cope with heterogeneity, because data may differ along several dimensions such as its format, refinement, dynamicity, and certainty; this is required to build a unified view of a complex information system.

Quality of knowledge-based data integration Knowledge-based data management can result in high quality data for users and applications. Yet, they also need mechanisms to assist data curators to constantly evaluate and improve all of their components towards the ultimate goal of matching the desired data integration level. Our aim is to investigate explanation mechanisms able to justify answers to queries and to point out inconsistencies in the data. We are also interested in techniques for deriving, within a knowledge-base, equivalent formulations of queries that are expressed outside of it, at the source level; these are critical for the verification of mappings and rules.

4 Application domains

4.1 Agronomy and agroecology

Agronomy is more and more at the center of important debates around questions of environmental impact related to the practice of intensive agriculture, especially at large scale. Through our research collaborations with INRAE (National Research Institute for Agriculture, Food and Environment) and DFKI (German Institute for Artificial Intelligence) our goal is to contribute and to define new models, techniques, and applications, enabling a better exploitation of data generated in these fields so as to put it at the service of decision-making processes.

Agronomy is a strong expertise domain in the area of Montpellier. And indeed, BOREAL is a joint team with INRAE, and the team has established closed collaborations with two Montpellier research laboratories (UMR, “Unités Mixte de Recherche”), namely IATE and ABSys. These collaborations can also reach a larger extent, for example, in the context of the #DigitAg (Institute Convergences Agriculture Numérique, Section 8.3) our team participated to the joint Inria-INRAE “White Book” on digital agriculture which can be considered a manifesto of the current challenges posed by digital agriculture [16].

A major issue for IATE (Engineering of Agro-polymers and Emerging Technologies) is to model the transformation of products in agrifood chains (i.e., the chain of all processes leading from some raw material, such as plants, to the final products, including waste treatment). This modeling has several objectives. It provides better understanding of the processes from start to finish, which aids in decision making, with the aim of improving the quality of the products and decreasing the environmental impact (e.g., reducing waste, choosing right food packaging). There is a need for tools for making easier for data scientists to integrate and analyze the heterogeneous data resulting from agrifood chains.

A major issue for ABSys (Biodiversified Agrosystems) is the study of sustainable farming systems. It is now established that the restoration of sustainable farming systems requires the adoption of agroecological practices supporting the reintroduction of biodiversity in agroecosystems. Indeed, an agroecosystem should provide not only cash crops but also ecosystem services that support the durability of the farming systems itself. This leads to more complex agroecosystems including a higher number of plant species. There is thus a crucial need for tools that would assist users in the design of such new agroecosystems, from researchers in agronomy to agricultural advisors and farmers.

Beside INRAE, our team collaborates with two DFKI teams located in Osnabrück and Kaiserslautern in the context of a bilateral project Inria-DFKI (“R4Agri”, Section 8). From an applicative perspective, the major issue targeted by this project is the development of monitoring tools based on reasoning which can equip robotic or mechanic devices used in agricultural farms. This can be used to enhance the agricultural processes but also to enforce regulations, for instance by assessing that the spraying of chemicals remains at a safe distance from river borders. In this context, there is a need for tools allowing one to interpret and analyze the number of types of sensor data that are generated.

5 Highlights of the year

- The team published three works in top-tier venues (Core Ranking A*/A) targeting topics in knowledge representation and reasoning (KR) and database theory (PODS, LMCS).
- Marie-Laure Mugnier was named Program Chair of KR 2026 (with F. Baader at TU Dresden), the leading conference in the domain of knowledge representation and reasoning. Nofar Carmeli has been involved in the organization of two top venues in databases: PODS 2025 (proceeding chair) and EDBT/ICDT 2026 (local organizer).
- David Carral organized the 1st European Workshop on Formal Logic At Montpellier ANd database Theory (FLAMANT 2025). As part of the workshop, 12 researchers visited our team and engaged in many small-group discussions with the goal of fostering collaborations.
- The team continued the software development activity of InteGraal, and introduced a number of satellite software libraries dedicated to working with existential rules (Py4Graal, DLGPE, NanoParse, IRIRef).

6 Latest software developments, platforms, open data

6.1 Latest software developments

6.1.1 InteGraal

Name: InteGraal : Knowledge-Representation and Reasoning for Data Integration

Keywords: Knowledge Bases, Data integration, Knowledge representation, Automated Reasoning, Heterogeneous Data, Knowledge Graphs

Scientific Description: InteGraal is a tool for integrating and reasoning on heterogeneous and federated data. The tool embodies algorithms and techniques developed at the crossroads between the fields of knowledge representation and reasoning and data management. From the historic point of view, this tool is the result of a complete re-engineering of the Graal tool, whose API and functionalities have been completely updated. Also, with respect to Graal, the tool is very much oriented towards data integration.

Functional Description: InteGraal has been designed in a modular way, in order to facilitate software reuse and extension. It should make it easy to test new scenarios and techniques, in particular by combining algorithms. The main features of Graal are the following: (1) internal storage to store data by using a SQL or RDF representation (Postgres, MySQL, HSQL, SQLite, Remote SPARQL endpoints, Local in-memory triplestores) as well as a native in-memory representation (2) data-integration capabilities for exploiting federated heterogeneous data-sources through mappings able to target systems such as SQL, RDF, and black-box (eg. Web-APIs) (3) algorithms for query-answering over heterogeneous and federated data based on query rewriting and/or forward chaining (or chase)

Release Contributions: 2025. Refactored query and reasoning interfaces to handle heterogenous data-sources. Added a number algorithms for explaining reasoning. Worked on APIs for applications building on InteGaral.

2024. Added reasoning with stratified negation. Advanced usability and features for mappings for integrating heterogeneous data. Extended the command line interface.

2023. Mappings for integrating heterogeneous data. Compilation-based query rewriting. Command line interface.

2022: First release, software deposit with Apache 2 licence.

2021: Functional specification, design and development of a major improved version of the tool. Started refactoring of the API, and of several modules for knowledge base representation, data storage, query answering and forward-chaining reasoning (chase). Started the development of new modules for handling heterogeneous data: mappings and federations.

News of the Year: This year we refactored the query interface to provide more flexible support for reasoning on heterogeneous data. We also added an important module which includes a number of algorithms for explaining queries and reasoning. We worked on the APIs of the tool to ease the development of applications on top of InteGraal. Finally, the tool has been used in several internships and thesis of the team, as well as for collaborations with our research partners.

URL: <https://gitlab.inria.fr/rules/integraal>

Publication: [INRIA2/lirmm-04304601v1](#)

Contact: Federico Ulliana

Participants: Akira Charoensit, Jean-Francois Baget, Pierre Bisquert, Guillaume Perution-Kihli, Michel Leclère, Marie-Laure Mugnier, Florent Tornil, Federico Ulliana

6.1.2 TreeForce

Keywords: JSon, Databases, Knowledge Bases, Automated Reasoning, Rewriting, NoSQL, Data integration, Knowledge representation, Heterogeneous Data

Scientific Description: TreeForce is a java tool for reasoning on tree data. It leverages on query rewriting techniques and NoSQL document oriented key-value stores. This library can be seen as a general toolbox for implementing reasoning techniques tailored for tree-shaped data and rules. It is composed of two main modules. The first includes generic data structures and algorithms for trees and tree-automata. The second includes automata-based query rewriting techniques as well as efficient evaluation techniques for large sets of rewritings.

Functional Description: TreeForce is a java tool for reasoning on tree data. It leverages on query rewriting techniques and NoSQL document oriented key-value stores. This library can be seen as a general toolbox for implementing reasoning techniques tailored for tree-shaped data and rules. It is composed of two main modules. The first includes generic data structures and algorithms for trees and tree-automata. The second includes automata-based query rewriting techniques as well as efficient evaluation techniques for large sets of rewritings.

Release Contributions: 2023. ArangoDB wrapper. Code improvement. 2022. Novel instance-aware rewriting and evaluation algorithms. Introduced summarization, partitioning and parallelization techniques. 2021: First version of TreeForce. Automata for unordered tree languages. Automata-based query-rewriting algorithms. MongoDB wrapper.

Contact: Federico Ulliana

Participants: Olivier Rodriguez, Federico Ulliana

6.1.3 B-Runner

Name: B-Runner

Keywords: Benchmarking, Experimentation, Java, Automated Reasoning, Databases, Knowledge Graphs

Scientific Description: B-Runner is a Java tool for the conduction of experimental analysis.

Functional Description: B-Runner is a library for collaborative benchmarking on knowledge and rule-based reasoners. The motivation for this project was to systematize testing, both on on InteGraal and other reasoners. The goal of B-Runner is to enable benchmarking for reasoning tools with a small cost, high robustness, and repeatability guarantees. The tool can be used as a best-practice for realizing and communicating on experimental analysis.

Release Contributions: 2025. Added support for testing explanations with the OWL-API and for InteGraal. 2024. Core module for experiment conduction

News of the Year: This year we developed new modus of B-Runner for testing the explanation of reasoning with tools such as the OWL-API and InteGraal.

URL: <https://gitlab.inria.fr/rules/brunner>

Publication: [lirmm-04646842v2](#)

Contact: Federico Ulliana

Participants: Federico Ulliana, Quentin Yeche, Pierre Bisquert, Akira Charoensit, Florent Tornil, Renaud Colin

6.1.4 IRIRefs

Name: IRIRefs

Keyword: Standard

Scientific Description: Contains a full RFC 3987 compliant parser of IRI References. Allows resolution of relative IRIs, recomposition, and all normalization schemes suggested in the standard. Also allows for relativisation, the inverse of resolution. To the best of our knowledge, this is the only java library fully compliant with the standard. Also contains an IRIManager, basis for the management of IRIs in DLGPE and Integraal.

Functional Description: This project aims at a java implementation of RFC 3987 Internationalized Resource Identifiers (IRIs). It relies upon a parser written in nanoparse, and offers the possibility to build irirefs (relative or not) from strings, recompose (display) them, resolve a relative against a base or normalize a full IRI, all according to the specifications in RFC 3987. A relativization mechanism is also offered to display short versions of a full IRI.

Release Contributions: See Changelog.

News of the Year: Launch of irirefs. Available in gitlab, github (as a mirror) and maven central.

URL: <https://gitlab.inria.fr/jfbaget/irirefs>

Contact: Jean-Francois Baget

Participant: Jean-Francois Baget

6.1.5 Py4Gaal

Name: Py4Gaal

Keyword: Knowledge representation

Scientific Description: py4graal is a python library that communicates with a java server running Integraal. It allows a lightweight, intuitive access to Integraal reasoning mechanisms. It allows a simplified access to Integraal for developers that do not want to get involved in the complexity of the Integraal java library, even with the simplified access provided by the external API. We believe this python library to be a pre-requisite for Integraal to be adapted, for instance, in a Data Science environment.

Functional Description: Py4Gaal is a Python interface to the Integraal reasoning engine, implemented in Java and accessed through Py4J. It brings rule-based reasoning into Python, allowing you to create fact bases, define rules, and evaluate queries while delegating the heavy lifting to a high-performance Java backend.

Release Contributions: Port of the version developed for Graal to Integraal

News of the Year: Launch of py4graal2, a completely rewritten version of the first release written in 2021 by Tom Salembien. The current version, written by Carole Beaugois, communicates with Integraal, whereas the 2021 version communicated with its predecessor, Graal. Since then, py4graal has been used to showcase Integraal's reasoning capabilities: during the LIRMM evaluation as well as in tutorials for potential partners, the simplicity of its use has been highlighted.

URL: <https://gitlab.inria.fr/jfbaget/py4graal2>

Contact: Jean-Francois Baget

Participants: Jean-Francois Baget, Carole Beaugois, Tom Salembien

6.1.6 DLGPE

Name: Datalog Plus Extended

Keywords: Knowledge representation, Parser

Scientific Description: DLGPE is both a language that generalizes the DLGP language used by Integraal and paves the way for future developments, and a Java library that can be used to: * parse DLGPE (using ANTLR4 and NanoParse for IRI references), * visit the AST and generate Java objects (for example those used in Integraal), * run an LSP server to communicate with editing tools, * use a semantics-based syntax-highlighting editor.

Functional Description: DLGPE is a java library consisting in an ANTLR4 grammar for the DLGPE language, a visitor allowing to build the parsed objects in Integraal, a LSP server and an editor with syntax highlighting.

News of the Year: Initial launch of DLGPE. This is a preliminary version that will be used in 2026 as the basis for an Inria ADT.

URL: <https://gitlab.inria.fr/jfbaget/dlgpe/>

Contact: Jean-Francois Baget

Participant: Jean-Francois Baget

6.1.7 NanoParse

Name: NanoParse

Keyword: Parser

Scientific Description: Nanoparse allows to define grammars with Java constructs (no external grammar files required). It is composed of modular readers (regex, string, sequence, choice, repetition, optional, etc.) and generates fully navigable parse trees. It supports recursive grammars, is lightweight & fast when no deep look-ahead is required, and is ready for integration with domain-specific languages.

Functional Description: NanoParse is a lightweight, composable parsing library written in Java. It lets you define grammars directly in code and parse complex structures with minimal boilerplate.

News of the Year: Nanoparse published as a translation of a former Python version. Now available on gitlab, github (as a mirror), and maven central. Developed to be the parsing engine for irirefs.

URL: <https://gitlab.inria.fr/jfbaget/nanoparse>

Contact: Jean-Francois Baget

Participant: Jean-Francois Baget

7 New results

Before presenting this year’s results, we first introduce some general preliminary notions in Section 7.1 to provide context for the results discussed later in this section. Moreover, we provide a summary of this year’s contributions in Section 7.2 and then discuss each of them in a dedicated section.

7.1 Preliminaires about Knowledge-Based Data Management with Existential Rules

This broad topic encompasses research areas such as ontology-mediated query answering (OMQA), data integration (DI), and ontology-based data access (OBDA) because of the expressivity of existential rule languages and the complexity of integration architectures it embraces.

Existential rules. Existential rules are first-order-logic formulas representing implications of the form $\forall X \forall Y [Body(X, Y) \rightarrow \exists Z Head(X, Z)]$ where *Body* and *Head* are positive conjunctions of atoms without functional symbols, and *Head* can have existentially quantified variables. These rules allow one to model complex relationships over the domains of interest, and at the same time dispose of a value invention mechanism through existentially quantified variables. This makes them suitable for many data and knowledge tasks on both open and closed domains. As a result, existential rules are ubiquitous in many fields. They are used to model dependencies, schema mappings, and expressive queries in databases. They are used as ontological languages as a valid complement to Description Logics, and at the same time as a generalization of so-called Horn Description Logics which lay at the foundations of important Semantic Web standards.

Rule-based query answering. Given a query Q , a database D , and a set of rules R , query answering asks to determine whether $D, R \models Q$ (where \models denotes standard first-order logic entailment), that is if the query Q is a logical consequence of the knowledge base made by the database D and the rules R . In the field of knowledge representation and reasoning, rule-based query answering is studied for rules expressing ontologies and referred as ontology-mediated query answering (OMQA). Formalisms such as Description Logics and Existential Rules (a.k.a, Tuple-Generating-Dependencies, or Datalog \pm) are typically targeted for expressing ontologies. Overall, the main emphasis of this topic is in the study of *rule languages* and the role they play in query answering.

Rule-based query answering over heterogeneous and federated data. In this context, the problem formulation remains similar, however the database D is replaced by a more complex notion of *federation* $(\mathcal{D}, \mathcal{M}, \mathcal{S})$ where \mathcal{D} is a collection of heterogeneous data sources, \mathcal{S} is a global integration schema, and \mathcal{M} is a set of mappings linking the datasources in \mathcal{D} to the global schema \mathcal{S} . This framework is at the foundations of data integration (DI) in databases and of ontology-based data access (OBDA) in knowledge representation and reasoning. OBDA focuses on global integration schemes and rules built on ontologies enabling query rewriting, while DI is more concerned with rules representing data-dependencies. Overall, both give more emphasis to *heterogeneous and federated data* in rule-based query answering.

Reasoning strategies for query answering. The two prominent strategies for rule-based query answering are *materialization* (also known as saturation, or forward-chaining) and *virtualization* (also known as query rewriting, or backward chaining). Both can be seen as ways of reducing query answering (which involves reasoning) into classical query evaluation. Materialization amounts to storing the inferences enabled by rules, thereby obtaining an extended database, on which queries are evaluated. Query rewriting amounts to compiling relevant rules into the query, thereby obtaining a rewritten query (usually a union of queries), which is evaluated on the (unaltered) database. Both approaches have their own strengths, and at the basis of this duality is the fact that while materialization is independent of queries, rewriting is independent of the database. Hence, each strategy better suits certain applicative scenarios, and both can possibly be combined thereby resulting in hybrid approaches.

7.2 Contributions

This year, we studied a number of theoretical, algorithmic, and applied questions of knowledge-based data management and database theory. Our main contributions cover the following topics:

- Foundational issues (Section 7.3) related to the termination of reasoning strategies using existential rules, abstracting data queries into the ontology, and the fine-grained complexity of answering queries;
- Applications, using InteGraal to develop knowledge-based applications in the context of enforcing regulations in agriculture and machine learning for extreme event prediction, and using argumentation techniques towards justified decision-making in agri-food systems.

In addition to our main publications, presented next, it is worth noting that the team also supervised a number of student internships from ENS Paris, INSA Toulouse, and University of Montpellier (Section 9.2.1), investigating other foundational and applied issues of reasoning and database theory. Finally, complementing methodological work, we also pursued an important team effort in the development of tools for rule-based query answering over heterogeneous and federated data (see Section 6.1).

7.3 Foundations of Databases and Knowledge Bases

Participants: Jean-Francois Baget, Pierre Biquert, Nofar Carmeli, David Caral, Akira Charoensit, Michel Leclère, Marie-Laure Mugnier, Guillaume Perution-Kihli, Federico Ulliana.

Restricted Chase Termination: You Want More than Fairness. The chase is a fundamental algorithm with ubiquitous uses in database theory. Given a database and a set of existential rules, it iteratively extends the database to ensure that the rules are satisfied in a most general way. This process may not terminate, and a major problem is to decide whether it does. This problem has been studied for a large number of chase variants, which differ by the conditions under which a rule is applied to extend the database. Surprisingly, the complexity of the universal termination of the restricted (a.k.a. standard) chase is not fully understood. We close this gap by placing universal restricted chase termination in the analytical hierarchy. This higher hardness is due to the fairness condition, and we propose an alternative condition to reduce the hardness of universal termination.

- Published at the Principles of Database Systems conference (PODS 2025) [8], with Lukas Gerlach (TU Dresden), Lucas Larroque, and Michaël Thomazo (VALDA, DI-ENS, PSL university, CNRS).

Abstractions of Queries in Ontology-Based Data Access. In ontology-based data access (OBDA), multiple data sources are integrated via mappings to an ontology. We consider an OBDA setting based on existential rules, hence a single formalism to encode both the mappings and the ontology. Query answering relies on the standard semantics of *certain answers*. We address the recent issue of *query abstraction*, which consists of abstracting data queries by translating them to the ontology layer. Such issue arises in a range of relevant scenarios, related to the design of OBDA systems or the automatic characterization of the semantics of data services implemented at the data level. Since a perfect abstraction may not exist, the notions of minimally-complete and maximally-sound abstractions have been introduced. These can be seen as approximations of perfect abstractions.

We study query abstractions within an extension of (unions of) conjunctive queries with a limited form of inequality and a special predicate marking database constants. While this extension does not lead to an increased complexity of the problems of interest, we show that it is able to express minimally-complete abstractions, and so also perfect abstractions when they exist. We also characterize maximally-sound abstractions by making a new connection with a notion stemming from data exchange (namely, that of maximum recovery).

- Published at the International Conference on Principles of Knowledge Representation and Reasoning (KR 2025) [9].

Enumeration fine-grained complexity of unions of conjunctive queries. We study the enumeration of answers to Unions of Conjunctive Queries (UCQs) with optimal time guarantees. More precisely, we wish to identify the queries that can be solved with linear preprocessing time and constant delay. Despite the basic nature of this problem, it was shown only recently that UCQs can be solved within these time bounds if they admit free-connex union extensions, even if all individual CQs in the union are intractable with respect to the same complexity measure [15]. Our goal is to understand whether there exist additional tractable UCQs, not covered by the currently known algorithms. As a first step, we show that some previously unclassified UCQs are hard using the classic 3SUM hypothesis, via a known reduction from 3SUM to triangle listing in graphs. As a second step, we identify a question about a variant of this graph task that is unavoidable if we want to classify all self-join-free UCQs: is it possible to decide the existence of a triangle in a vertex-unbalanced tripartite graph in linear time? We prove that this task is equivalent in hardness to some family of UCQs. Finally, we show a dichotomy for unions of two self-join-free CQs if we assume the answer to this question is negative. In conclusion, this work pinpoints a computational barrier in the form of a single decision problem that is key to advancing our understanding of the enumeration complexity of many UCQs. Without a breakthrough for unbalanced triangle detection, we have no hope of finding an efficient algorithm for additional unions of two self-join-free CQs. On the other hand, a sufficiently efficient unbalanced triangle detection algorithm can be turned into an efficient algorithm for a family of UCQs currently not known to be tractable.

- Published at Logical Methods in Computer Science (LMCS) [7], with Karl Bringmann (Max Planck Institute for Informatics, Saarland University).

7.4 Applications of Rule-Based Reasoning

Participants: Jean-Francois Baget, Pierre Bisquert, David Carral, Michel Leclère, Marie-Laure Mugnier, Guillaume Perution-Kihli, Akira Charoensit, Federico Ulliana.

R4Agri: Integrating Environmental Regulations Into Autonomous Agricultural Robotics: A Case for Waterbody-Aware Fertilization. As part of the R4Agri project with DFKI (see Section 8.1), we tackled the issue of operating autonomous robots in the agricultural domain while integrating and reasoning on background knowledge. Specifically, operating such robots requires compliance with the regulatory aspects of the process. For instance, the improper spreading of chemicals near water bodies (e.g., fertilizers) may cause significant environmental damage and, therefore, is strictly regulated. As an emblematic case study, we considered the decision-making process of a mobile robot spreading fertilizers near a protected water body. This requires the integration of contextual information obtained from multimodal sensors and high-level knowledge about regulatory aspects. We proposed an intrinsically-declarative framework for such an integration. Our framework leverages and extends semantic web vocabularies to integrate regulatory constraints and environmental conditions where the robot is operating. Then, it uses rules and reasoning to detect risks of violation on real-time data generated by the autonomous robot. The inference of a risk of violation subsequently triggers actions controlling the robot behavior. Such inferences can be transparently explained, thereby avoiding the robot to behave as a black-box to the supervising technician.

This framework was implemented using our InteGraal tool and demonstrated on the following scenario on a physically-based virtual environment: two vehicles, an unmanned ground vehicle (the spraying robot) and an unmanned aerial vehicle (a drone), are cooperating to build a spatial representation on their environment, using different sensors; the sensed data is fused and interpreted, producing higher-level data such as the position of the spraying engine, the field slope gradient at this position, or the distance to the border of the water body. Such data is then provided as a stream to the reasoning engine.

- Published in three venues (each paper focusing on a different aspect):
 - the European Conference on Mobile Robots (ECMR 2025) [10]
 - the International Joint Conference on Rules and Reasoning (RuleML+RR-2025) [11]
 - the Künstliche Intelligenz in der Umweltinformatik workshop (KIU-2025) [12]

with Ahmad Kadi, Ansgar Bernardi, Martin Atzmueller, and Nikolas Müller (DFKI, Bremen).

- [Companion repository](#) for the paper at RuleML+RR-2025.

Gypscie-KG: Building a Logic-Based Approach for Knowledge Graph Data Integration View in ML Systems. In the context of a collaboration with the Iroko team and LNCC (Brazil), we contributed to the development of Gypscie-KG, a system that integrates heterogeneous machine learning (ML) data into a knowledge graph (KG) using logic based-rules to enable semantic queries and reasoning. Gypscie is a web platform for managing machine learning tasks (i.e., training and running models) dedicated to the prediction of extreme meteorological events. By relying on the InteGraal tool developed by our team, we built a knowledge graph capturing all of the processes and the data handled by Gypscie. This allows data-scientists to have new tools to explain extreme event predictions, such as analyzing the input data that led to a particular prediction, the datasets used to train a model, as well the transformation that a dataset has suffered.

- Published in the LAGO 2025 Workshop [13] organized within the Brazilian Database Conference (SBBDB 2025), with Gabriela Moraes, Fabio Porto, Bernardo Gonçalves (LNCC / MCT, Rio de Janeiro), and Patrick Valduriez (Iroko).

Justified Preference Aggregation in Agri-Food Systems: An Approach, Argumentation Methods, and Tools. In recent years, there has been a growing recognition of the need to ensure sustainability of agri-food systems, covering a variety of stakeholders and activities from production to waste management. Multi-Criteria Decision Assessment (MCDA) have emerged as crucial tools for sustainability assessment, but many do not consider stakeholder perspectives, which prevents their adoption and use. Our projects, NoAW (No Agricultural Waste) and AgriLoop (High-value products from agricultural residues through sustainable chains), address this issue by focusing on innovative valorization routes for agricultural waste by assessing stakeholder impact categories through participatory decision-making. This article proposes a novel methodology integrating computational social choice and argumentation techniques for achieving justified collective decision-making in agri-food systems. This methodology includes a theoretical framework and related tools, which facilitate the identification, analysis, and aggregation of preferences based on justifications.

- Published in the International Journal of Agricultural and Environmental Information Systems [14], with Patrice Buche, Maksim Koptelov (IATE)

8 Partnerships and cooperations

8.1 International initiatives

8.1.1 Participation in other International Programs

Bilateral project R4Agri

Participants: Pierre Bisquert, David Carral, Akira Charoensit, Marie-Laure Mugnier, Guillaume Pérution, Federico Ulliana.

Title: “R4Agri”- Reasoning on Agricultural Data: Integrating Metrics and Qualitative Perspectives

Partner Institution(s):

- Inria
- DFKI, Germany

Date/Duration: 01/01/2022-30/06/2025 (42 months)

Website: project.inria.fr/r4agri/

Additional info: AI tools supporting competitive and sustainable agriculture need to exploit highly diverse kinds of data and knowledge, from raw data provided by sensors to high level expertise knowledge. Taking numerical agriculture as the targeted application domain, the overall goal of the R4Agri project is to provide frameworks for reasoning about knowledge based on heterogeneous data, with a focus on multi-modal and multi-scale sensor data. Main challenges include context-dependent interpretation of sensor data, which involves reasoning about prior knowledge, and query answering techniques that exploit domain knowledge and accommodate the specificities of data sources in a flexible manner.

On the Boreal side, we extended our tool Integraal to support such frameworks, which required to extend the rule language, to develop new kinds of data mappings, and to make the knowledge base architecture and query answering mechanisms more flexible in order to handle dynamic facts. We also developed different sorts of explanation facilities in order to justify the output of reasoning. Together with our DFKI partners, we demonstrated the application potential of our framework in a realistic use case (see Section 7.4).

8.2 International research visitors

8.2.1 Visits of international scientists

Many of the team's visitors this year came to attend the FLAMANT workshop. Moreover, the final meeting of the R4Agri project was held in Montpellier in May 2025.

Piotr Ostropolski-Nalewaja

Status: Associate professor

Institution of origin: University of Wroclaw

Country: Poland

Dates: From February 5 to February 21

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Andreas Pieris

Status: Professor

Institution of origin: University of Cyprus and of Edinburgh

Country: UK and Cyprus

Dates: From February 10 to February 13

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Lucas Larroque

Status: PhD

Institution of origin: Valda research team at Inria Paris

Country: France

Dates: From February 10 to February 14

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Carsten Lutz

Status: Associate professor

Institution of origin: Leipzig University

Country: Germany

Dates: From February 10 to February 14

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Michaël Thomazo

Status: Inria CRCN researcher

Institution of origin: Valda research team at Inria Paris

Country: France

Dates: From February 10 to February 14

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Jerzy Marcinkowsky

Status: Professor

Institution of origin: University of Wroclaw

Country: Poland

Dates: From February 11 to February 20

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Sebastian Rudolph

Status: Professor

Institution of origin: TU Dresden

Country: Germany

Dates: From February 11 to February 21

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Lukas Gerlach

Status: PhD

Institution of origin: TU Dresden

Country: Germany

Dates: From February 11 to February 20

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Meghyn Bienvenu

Status: Directeur de recherche at CNRS

Institution of origin: CNRS - LaBRI

Country: France

Dates: From February 13 to February 21

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Lucia Gomez Alvarez

Status: CRCN researcher

Institution of origin: Moex research team at Inria Grenoble

Country: France

Dates: From February 17 to February 21

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Timothy Stephen Lyon

Status: Postdoctoral researcher

Institution of origin: TU Dresden

Country: Germany

Dates: from February 17 to February 19

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: research stay

Stefan Mengel

Status: CNRS researcher

Institution of origin: Centre de Recherche en Informatique de Lens (CRIL)

Country: France

Dates: From February 24 to February 28

Context of the visit: Collaboration with Nofar Carmeli

Mobility program/type of mobility: Research stay

Ansgar Bernardi

Status: Research scientist

Institution of origin: DFKI

Country: Germany

Dates: From May 21 to May 23

Context of the visit: R4Agri

Mobility program/type of mobility: Bilateral project

Ahmad Kadi

Status: Engineer

Institution of origin: DFKI

Country: Germany

Dates: From May 21 to May 23

Context of the visit: R4Agri

Mobility program/type of mobility: Bilateral project

Michaël Thomazo

Status: Inria CRCN researcher

Institution of origin: Valda research team at Inria Paris

Country: France

Dates: From July 21 to July 25

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

Piotr Ostropolski-Nalewaja

Status: Associate professor

Institution of origin: University of Wrocław

Country: Poland

Dates: From September 15 to October 1

Context of the visit: Ongoing collaboration

Mobility program/type of mobility: Research stay

Quentin Manière

Status: Postdoctoral researcher

Institution of origin: Center for Scalable Data Analytics and AI at Leipzig

Country: Germany

Dates: From September 29 to October 10

Context of the visit: FLAMANT 2025

Mobility program/type of mobility: Research stay

8.3 National initiatives

EXPAND ANR Project (ANR-25-CE23-1215) (2025-2030)

Participants: Jean-François Baget, Pierre Bisquert, Nofar Carmeli, David Carral, Michel Leclère, Marie-Laure Mugnier, Federico Ulliana.

EXPAND (Expanding the reach of ontology-based data access: EXpressivity, exPlanation, and Algorithms) is an ANR project accepted in 2025 in the scientific axis "Artificial Intelligence and Data Sciences". This project, led by Inria Paris, brings together three major French teams in this field. The main goal of the project is to expand the applicability of ontology-based query answering by allowing for enhanced query languages, and provide richer ways to manipulate and understand query answers. The team participates in all project tasks, and all members of the team are involved in this project. One and a half years of post-doctoral work and two years of engineering work are planned for our team.

project.inria.fr/expand/

Convergence institute #DigitAg (2017-2026)

Participants: Jean-François Baget, Marie-Laure Mugnier, Federico Ulliana.

Located in Montpellier, #DigitAg (for Digital Agriculture) gathers 17 founding members: research institutes, including Inria, the University of Montpellier and higher-education institutes in agronomy, transfer structures and companies. Its objective is to support the development of digital agriculture. BOREAL is involved in this project on the issues of designing data and knowledge management systems adapted to agricultural information systems, and of developing methods for integrating different types of information and knowledge (generated from data, experts, models). A PhD thesis (Elie Najm, 2019-2022) investigated knowledge representation and reasoning for the design of new agroecological systems, in collaboration with the research laboratory ABSys - Biodiversified Agrosystems (formerly UMR SYSTEM).

www.hdigitag.fr/eng/

9 Dissemination

9.1 Promoting scientific activities

9.1.1 Scientific events: organization

- David Carral organized the 1st European Workshop on Formal Logic At Montpellier ANd database Theory (FLAMANT 2025), which took place in Montpellier from February 5 to February 21, 2025. The workshop aimed to foster collaborations between researchers in computational logic and database theory. Rather than scheduling a large number of talks, the emphasis was on small-group meetings, giving participants ample time to work together. Further information is available at the [seminar website](#).

General chair, scientific chair

- Proceedings chair of the Symposium on Principles of Database Systems (PODS 2025): Nofar Carmeli

Member of the organizing committees

- Local organization of EDBT/ICDT 2027 in Lille (Joint conference: International Conference on Extending Database Technology and International Conference on Database Theory): Nofar Carmeli

9.1.2 Scientific events: selection

Chair of conference program committees

- Program co-chair of the 23rd International Conference on Principles of Knowledge Representation and Reasoning (KR 2026): Marie-Laure Mugnier

Member of the conference program committees (PC)

- Area chair of the 22nd International Conference on Principles of Knowledge Representation and Reasoning (KR 2025): Marie-Laure Mugnier
- PC member of the 22nd International Conference on Principles of Knowledge Representation and Reasoning (KR 2025): David Carral, Michel Leclère
- PC member of the Symposium on Theoretical Aspects of Computer Science (STACS 2025): Nofar Carmeli
- PC member of the International Joint Conference on Rules and Reasoning (RuleML+RR 2025): Pierre Bisquert, David Carral, and Federico Ulliana
- PC member of the 38th International Workshop On Description Logics (DL 2025): David Carral

9.1.3 Journal

Reviewer - reviewing activities

- Reviewer for Artificial Intelligence Journal (AIJ) : Jean-François Baget, Marie-Laure Mugnier
- Reviewer for Transactions on Graph Data and Knowledge (TGDK): David Carral
- Reviewer for Information Processing Letters (IPL): Nofar Carmeli

9.1.4 Scientific expertise

Evaluation of scientific projects

- Member of the ANR Evaluation Committee “Artificial Intelligence and Data Science” (ANR CES 23 - AAP 25 - 154 submitted proposals): Marie-Laure Mugnier

Academic recruitment committees

- Member of a recruitment committee for a Professor position at the University of Montpellier (“repyramidage”): Marie-Laure Mugnier
- Member of a recruitment committee for an Assistant Professor position at the University Côte d’Azur: Michel Leclère, Marie-Laure Mugnier

9.1.5 Research administration

- President of the “Section 27 Committee” (Computer Science) of the University of Montpellier (July 2021 - November 2025): Marie-Laure Mugnier
- Member of the “Section 27 Committee” (Computer Science) of the University of Montpellier (July 2021 - November 2025): Michel Leclère
- Member of the Council and Human Resources Commission of the Scientific Pole MIPS (Mathematics Informatics Physics and Systems) of the University of Montpellier (since its creation): Marie-Laure Mugnier
- Member of the scientific animation group of INRAE’s Transform department (since April 2025): Pierre Bisquert
- Scientific leader at the local level (Inria Sophia) of the ANR Project EXPAND (ANR-25-CE23-1215): Michel Leclère

9.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

- Michel Leclère, Marie-Laure Mugnier, and Federico Ulliana teach at the Computer Science department of the Science Faculty. They are in charge of courses in Programming and Logics (Licence), as well as Symbolic Artificial Intelligence, Semantic Management of Data, Datawarehouses, Big-Data and NoSQL systems, and Theory of Data and Knowledge Bases (Master).
- Concerning full-time researchers in 2025, Jean-François Baget, Nofar Carmeli, and David Carral, taught in the Computer Science Master about Database Theory and Knowledge Bases (6 to 9 hours per person).

9.2.1 Supervision

- Pierre Bisquert, David Carral, and Federico Ulliana continue to co-supervise Akira Charoensit, a PhD student who began her doctoral research in May 2023. Her work focuses on the development of efficient algorithms for computing explanations of entailments in rule-based languages.
- David Carral continues to co-supervise Lucas Larroque with Michaël Thomazo. Lucas is a PhD student at ENS Paris who began his doctorate in September 2023. His research focuses on rewriting techniques aimed at obtaining decidable algorithms for reasoning in first-order logic.
- David Carral and Federico Ulliana supervised Jeanne Coschieri, an L3 student at ENS Paris, who completed a six-week research internship in our group on topics related to knowledge representation.
- David Carral co-supervised Laura Gruson during her M2 research thesis, together with Michaël Thomazo. Laura is an M2 student at ENS Paris, and her work focused on topics related to knowledge representation and dynamic complexity.
- Michel Leclère, Pierre Bisquert and Federico Ulliana supervised Abir Amina Hammoud and Ibrahim Al Ayoubi, a licence and master's student at the University of Montpellier, who worked together on a development project with InteGraal in collaboration with Iroko and LNCC.
- Jean-Francois Baget supervised Carole Beaugeois, student at INSA Toulouse, who worked on a Python API for InteGraal.

9.2.2 Juries

- Member of the PhD committee for David Camarazo at Bourgogne University in December 2025: Federico Ulliana
- Member of the PhD committee of Thomas Munoz, University of Hasselt, Belgium (The defense is scheduled in January 20th 2026, but reviewing the thesis was done in 2025): Nofar Carmeli

10 Scientific production

10.1 Major publications

- [1] K. Bringmann, N. Carmeli and S. Mengel. ‘Tight Fine-Grained Bounds for Direct Access on Join Queries’. In: *ACM Transactions on Database Systems* (6th Dec. 2024). DOI: [10.1145/3707448](https://doi.org/10.1145/3707448). URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-04834303>.
- [2] N. Carmeli, M. Grohe, B. Kimelfeld, E. Livshits and M. Tibi. ‘Database Repairing with Soft Functional Dependencies’. In: *ACM Transactions on Database Systems* 49.2 (2024), pp. 1–34/8. DOI: [10.1145/3651156](https://doi.org/10.1145/3651156). URL: <https://hal.science/hal-04617777>.
- [3] D. Carral, L. Gerlach, L. Larroque and M. Thomazo. ‘Restricted Chase Termination: You Want More than Fairness’. In: *ACM digital library*. PODS 2025 - ACM SIGMOD/PODS International Conference on Management of Data. Vol. 3. Proceedings of the ACM on management of data 2. Berlin, Germany, 22nd June 2025, pp. 1–17. DOI: [10.1145/3725246](https://doi.org/10.1145/3725246). URL: <https://hal.science/hal-05240721>.

- [4] D. Carral, L. Larroque and M. Thomazo. ‘Ontology-Based Query Answering over Datalog-Expressible Rule Sets is Undecidable’. In: KR 2024 - 21st International Conference on Principles of Knowledge Representation and Reasoning. Hanoi, Vietnam, 2nd Nov. 2024. URL: <https://hal.science/hal-04710719>.
- [5] I. Eldar, N. Carmeli and B. Kimelfeld. ‘Direct Access for Answers to Conjunctive Queries with Aggregation’. In: ICDT 2024 - 27th International Conference on Database Theory. Leibniz International Proceedings in Informatics. Paestum, Italy: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 4:1–4:20. DOI: [10.4230/LIPIcs.ICDT.2024.4](https://doi.org/10.4230/LIPIcs.ICDT.2024.4). URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-04617768>.
- [6] M. Leclère, M.-L. Mugnier and G. Pérution-Kihli. ‘Abstractions of Queries in Ontology-Based Data Access’. In: *Proceedings of KR’25*. KR 2025 - 22nd International Conference on Principles of Knowledge Representation and Reasoning. Melbourne, Australia, 11th Nov. 2025. URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-05246623>.

10.2 Publications of the year

International journals

- [7] K. Bringmann and N. Carmeli. ‘Unbalanced Triangle Detection and Enumeration Hardness for Unions of Conjunctive Queries’. In: *Logical Methods in Computer Science* 21.1 (27th Mar. 2025). DOI: [10.46298/lmcs-21\(1:29\)2025](https://doi.org/10.46298/lmcs-21(1:29)2025). URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-05246520> (cit. on p. 13).

International peer-reviewed conferences

- [8] D. Carral, L. Gerlach, L. Larroque and M. Thomazo. ‘Restricted Chase Termination: You Want More than Fairness’. In: *ACM digital library*. PODS 2025 - ACM SIGMOD/PODS International Conference on Management of Data. Vol. 3. Proceedings of the ACM on management of data 2. Berlin, Germany, 22nd June 2025, pp. 1–17. DOI: [10.1145/3725246](https://doi.org/10.1145/3725246). URL: <https://hal.science/hal-05240721> (cit. on p. 12).
- [9] M. Leclère, M.-L. Mugnier and G. Pérution-Kihli. ‘Abstractions of Queries in Ontology-Based Data Access’. In: *Proceedings of KR’25*. KR 2025 - 22nd International Conference on Principles of Knowledge Representation and Reasoning. Melbourne, Australia, 11th Nov. 2025. URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-05246623> (cit. on p. 12).
- [10] N. Müller, A. Kadi, M.-L. Mugnier, G. Pérution-Kihli, F. Ulliana and M. Atzmueller. ‘A Legislation-Aware Robotic Framework for Autonomous Fertilization Near Protected Water Bodies’. In: *IEEE Xplore*. ECMR 2025 - 12th European Conference on Mobile Robots. Padua, Italy, 2nd Sept. 2025. URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-05165190> (cit. on p. 13).
- [11] G. Pérution-Kihli, A. Kadi, N. Müller, A. Charoensit, D. Carral, P. Bisquert, F. Ulliana, A. Bernardi and M.-L. Mugnier. ‘Integrating Environmental Regulations Into Autonomous Agricultural Robotics: A Case for Waterbody-Aware Fertilization’. In: *9th RuleML+RR 2024: Istanbul, Turkey*. RuleML+RR 2025 - Rules and Reasoning, 9th International Joint Conference, RuleML+RR 2025. Istanbul, Turkey, 22nd Sept. 2025. URL: <https://hal.science/hal-05242346> (cit. on p. 13).

National peer-reviewed Conferences

- [12] A. Kadi, N. Müller, A. Bernardi, F. Ulliana and G. Pérution-Kihli. ‘Combining Open Data and Formal Reasoning for Autonomously Controlled Spreading near Water Bodies’. In: *KIU Workshopt 2025 - Künstliche Intelligenz in der Umweltinformatik*. Potsdam, Germany, 17th Sept. 2025. URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-05165170> (cit. on p. 13).

Conferences without proceedings

- [13] G. Moraes, F. Porto, F. Ulliana, J.-F. Baget, M. Leclère, P. Bisquert, B. Gonçalves and P. Valduriez. ‘Gypscie-KG: Building a Logic-Based Approach for Knowledge Graph Data Integration View in ML Systems’. In: LAGO 2025 – LLMs, Análise de Grafos e Ontologias (SBBD). Fortaleza, Brazil, 29th Sept. 2025. URL: <https://hal-lirmm.ccsd.cnrs.fr/lirmm-05229486> (cit. on p. 14).

Reports & preprints

- [14] P. Bisquert, P. Buche, A. Hecham and M. Koptelov. *Justified Preference Aggregation in Agri-Food Systems: An approach, Argumentation Methods and Tools (Supplementary Material)*. 28th Oct. 2025. URL: <https://hal.science/hal-04607020> (cit. on p. 14).

10.3 Cited publications

- [15] N. Carmeli and M. Kröll. ‘On the enumeration complexity of unions of conjunctive queries’. In: *ACM Transactions on Database Systems (TODS)* 46.2 (2021), pp. 1–41 (cit. on p. 13).
- [16] N. Mitton, L. Brossard, T. Bouadi, F. Garcia, R. Gautron, N. Hilgert, D. Ienco, C. Largouët, E. Lutton, V. Masson, R. Martin-Clouaire, M.-L. Mugnier, P. Neveu, P. Preux, H. Raynal, C. Roussey, A. Termier and V. Bellon Maurel. ‘Foundations and state of the art’. In: *Agriculture and Digital Technology: Getting the most out of digital technology to contribute to the transition to sustainable agriculture and food systems*. White book Inrira 6. Acknowledgements (contribution, proofreading, editing) – Isabelle Piot-Lepetit. INRIA, 2022, pp. 30–75. URL: <https://hal.inrae.fr/hal-03609470> (cit. on p. 6).