

2025 Activity Report

RESEARCH CENTRE: Inria Centre at Université Côte d'Azur


IN PARTNERSHIP WITH: CNRS, Université Côte d'Azur

Project-Team

WIMMICS

Web-Instrumented huMan-Machine Interactions,
Communities and Semantics

In collaboration with Laboratoire informatique, signaux systèmes de Sophia Antipolis (I3S)



Project-Team WIMMICS

Creation of the Project-Team: 2025 February 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

- A1.3.1. – Web
- A3.1.1. – Modeling, representation
- A3.1.2. – Data management, quering and storage
- A3.1.3. – Distributed data
- A3.1.4. – Uncertain data
- A3.1.6. – Query optimization
- A3.1.7. – Open data
- A3.1.10. – Heterogeneous data
- A3.1.11. – Structured data
- A3.2. – Knowledge
 - A3.2.1. – Knowledge bases
 - A3.2.2. – Knowledge extraction, cleaning
 - A3.2.3. – Inference
 - A3.2.4. – Semantic Web
 - A3.2.5. – Ontologies
 - A3.2.6. – Linked data
- A3.3.2. – Data mining
- A3.4. – Machine learning and statistics
- A5.1. – Human-Computer Interaction
 - A5.1.1. – Engineering of interactive systems
 - A5.1.2. – Evaluation of interactive systems
 - A5.1.8. – 3D User Interfaces
 - A5.1.9. – User and perceptual studies
- A5.2. – Data visualization
- A5.8. – Natural language processing
- A7.1.3. – Graph algorithms
- A7.2.2. – Automated Theorem Proving
- A8.2.2. – Evolutionary algorithms
- A9.1. – Knowledge
- A9.2. – Machine learning
 - A9.2.1. – Supervised learning
 - A9.2.2. – Unsupervised learning
 - A9.2.3. – Reinforcement learning
 - A9.2.6. – Neural networks
 - A9.2.8. – Deep learning
- A9.4. – Natural language processing
- A9.6. – Decision support

- A9.7. – AI algorithmics
- A9.8. – Reasoning
- A9.9. – Distributed AI, Multi-agent
- A9.10. – Hybrid approaches for AI
- A9.11. – Generative AI
- A9.13. – Agentic AI
- A9.15. – Symbolic AI
- A9.16. – Societal impact of AI

Other research topics and application domains

- B1.1. – Biology
- B1.2.2. – Cognitive science
- B2. – Digital health
- B5.8. – Learning and training
- B5.9. – Industrial maintenance
- B6.3.1. – Web
- B6.3.2. – Network protocols
- B6.5. – Information systems
- B9.5.1. – Computer science
- B9.5.4. – Chemistry
- B9.5.6. – Data science
- B9.6.2. – Juridical science
- B9.6.4. – Management science
- B9.6.6. – Archeology, History
- B9.6.7. – Geography
- B9.6.10. – Digital humanities
- B9.7. – Knowledge dissemination
- B9.7.1. – Open access
- B9.7.2. – Open data

Contents

Project-Team WIMMICS	1
1 Team members, visitors, external collaborators	7
2 Overall objectives	8
2.1 Context and Objectives	8
2.2 Research Topics	9
3 Research program	9
3.1 Knowledge Graph Life Cycle with a view on Data Integration	9
3.2 Combined intelligent methods for heterogeneous knowledge graphs	11
3.3 Interaction design for decision making on and with knowledge graphs	12
4 Application domains	13
5 Social and environmental responsibility	14
5.1 Footprint of research activities	14
5.2 Impact of research results	14
6 Highlights of the year	14
6.1 Awards	14
7 Latest software developments, platforms, open data	14
7.1 Latest software developments	14
7.1.1 CORESE-Core	14
7.1.2 CORESE-GUI	15
7.1.3 CORESE-Server	15
7.1.4 CORESE-Command	16
7.1.5 Datalens	16
7.1.6 DBpedia	16
7.1.7 Gen ² KGBot	17
7.1.8 GUsT-3D	17
7.1.9 IndeGx	17
7.1.10 KartoGraphI	18
7.1.11 Metadatamatic	18
7.1.12 MGExplorer	19
7.1.13 Muvin	19
7.1.14 Olivaw	20
7.1.15 Q ² Forge	21
7.1.16 RDFminer	21
7.1.17 SciLEX	21
7.1.18 SemWebRAG	22
7.1.19 Zoomathia KG Pipeline	22
7.1.20 Zoomathia KG Web Application	22
7.2 Open data	23
8 New results	24
8.1 Knowledge Graph Life Cycle with a view on Data Integration	24
8.1.1 Semantic Web for the Integration of Pharmacogenomics Knowledge	24
8.1.2 A Semantic Web Ontology for Psychosocial Factors of Dysfunctional Eating Attitudes and Behaviors in Sport	25
8.1.3 An Ontology for Modeling User Activity within Visualization Interfaces	25
8.1.4 Semantic Annotation of a Corpus of Texts in Ancient Zoology as a Knowledge Graph	25
8.1.5 Semantic Annotation of Scientific Litterature in Agriculture as a Knowledge Graph	26

8.1.6	Agile Ontology Engineering: Tooling and Methodology	26
8.1.7	Corese Semantic Web Factory	26
8.1.8	W3C Data activity and AC Rep	27
8.2	Combined intelligent methods for heterogeneous knowledge graphs	27
8.2.1	Sentiment Analysis with Fuzzy Polarity Propagation	27
8.2.2	IndeGx	28
8.2.3	Learning Pattern-Based Extractors from Natural Language and Knowledge Graphs: Applying Large Language Models to Wikipedia and Linked Open Data	28
8.2.4	Hybridizing machine learning and knowledge graphs: injection of relation signatures	29
8.2.5	Stability of knowledge graph embedding models for link prediction	29
8.2.6	Temporal Graph Modeling and Hybrid Learning for Health State Prediction	29
8.2.7	Selecting relevant pairs for analogy-based pruning of knowledge graphs	30
8.2.8	How to reason with probabilistic information on argumentation graphs?	30
8.2.9	Investigating Language Model Capabilities to Represent and Process Formal Knowledge in Order to Assist Ontology Engineering	30
8.2.10	Hypermedia Multi-Agent Systems	31
8.3	Interaction design for decision making on and with knowledge graphs	31
8.3.1	Visual Exploration of Individual Mobility Data	31
8.3.2	Temporal Exploration of Knowledge Graph-based Collaboration Networks	32
8.3.3	Interaction with extended reality	32
8.3.4	Supporting Dataset Discovery via Network Topologies and Visualization	33
8.3.5	AI agent to convert natural language questions into SPARQL queries	33
8.4	Study of the Attention Economy, its detrimental impacts and leads for regulation	34
9	Bilateral contracts and grants with industry	34
9.1	Bilateral contracts with industry	34
10	Partnerships and cooperations	37
10.1	International initiatives	37
10.1.1	Participation in other International Programs	37
10.2	International research visitors	38
10.2.1	Visits of international scientists	38
10.2.2	Visits to international teams	39
10.3	European initiatives	39
10.3.1	Other european programs/initiatives	39
10.4	National initiatives	40
10.5	Regional initiatives	44
11	Dissemination	45
11.1	Promoting scientific activities	45
11.1.1	Scientific events: organization	45
11.1.2	Scientific events: selection	46
11.1.3	Journal	46
11.1.4	Invited talks	47
11.1.5	Scientific expertise	47
11.1.6	Research administration	48
11.2	Teaching - Supervision - Juries - Educational and pedagogical outreach	48
11.2.1	Teaching	48
11.2.2	Supervision	50
11.2.3	Juries	50
11.3	Popularization	51
11.3.1	Productions (articles, videos, podcasts, serious games, ...)	51
11.3.2	Participation in Live events	52

12 Scientific production	52
12.1 Major publications	52
12.2 Publications of the year	53
12.3 Cited publications	58

1 Team members, visitors, external collaborators

Research Scientists

- Pierre-Antoine Champin [INRIA, Associate Professor Detachement]
- Fabien Gandon [INRIA, Senior Researcher, HDR]
- Franck Michel [CNRS, Researcher]
- Pierre Monnin [INRIA, Researcher, from Oct 2025]

Faculty Members

- Catherine Faron [Team leader, UNIV COTE D'AZUR, Professor, HDR]
- Hajer Akid [UNIV COTE D'AZUR, Associate Professor, from Sep 2025]
- Marco Alba Winckler [UNIV COTE D'AZUR, Professor, HDR]
- Aline Menin [UNIV COTE D'AZUR, Associate Professor]
- Pierre Monnin [UNIV COTE D'AZUR, Associate Professor, until Sep 2025]
- Andrea Tettamanzi [UNIV COTE D'AZUR, Professor, HDR]

Post-Doctoral Fellows

- Yousouf Taghzouti [CNRS, from Nov 2025]
- Yousouf Taghzouti [UNIV COTE D'AZUR, Post-Doctoral Fellow, until Oct 2025]

PhD Students

- Hanna Abi Akl [DSTI School of Engineering, from Dec 2025]
- Matthieu Feraud [CNRS, from Oct 2025]
- Ndeye Emilie Mbengue [UNIV COTE D'AZUR, from Oct 2025]
- Guillaume Meroue [INRIA]
- Genesis Coromoto Montenegro Uribe [BERGER-LEVRAULT, CIFRE, from Sep 2025]
- Benjamin Navet [CNRS]
- Clement Quere [UNIV COTE D'AZUR]
- Celian Ringwald [INRIA]
- Nicolas Robert [UNIV COTE D'AZUR]

Technical Staff

- Abdessamad Abdoun [INRIA, Engineer, from May 2025]
- Remi Ceres [INRIA, Engineer]
- Florent Jaillet [CNRS]
- Pierre Maillot [INRIA, Engineer]

Interns and Apprentices

- Krysto Dagues De La Hellerie [INRIA, Intern, from Mar 2025 until Aug 2025]
- Minh Huy Do [INRIA, Intern, from Feb 2025 until Jun 2025]
- Erwan Hain [INRIA, Apprentice, until Jan 2025]
- Matteo Lacheny [UNIV COTE D'AZUR, Intern, from Apr 2025 until Jun 2025]
- Jeremy Moncada [INRIA, Intern, from Apr 2025 until Aug 2025]
- Sajal Paudyal [INRIA, Intern, from Apr 2025 until Sep 2025]
- Killian Michel Liam Piel [UNIV COTE D'AZUR, Intern, from Apr 2025 until Jun 2025]
- Salma Talib [UNIV COTE D'AZUR, Intern, from Mar 2025 until Aug 2025]
- Lyam Thibaud [INRIA, Intern, from Jun 2025 until Aug 2025]

Administrative Assistant

- Jane Desplanques [INRIA]

Visiting Scientists

- Andrea Nasi [Univ Torino, from Sep 2025]
- Bryan Elliott Tam [UNIV GAND, from Sep 2025 until Oct 2025]

External Collaborators

- Hanna Abi Akl [DSTI School of Engineering, until Nov 2025]
- Andrei Ciortea [UNIV SAINT-Gall]
- Olivier Corby [retired]
- Nicolas Delaforge [Probabl]
- Molka Dhouib [Freelance, from Nov 2025]
- Alain Giboin [retired]

2 Overall objectives

2.1 Context and Objectives

With more than 5 billion direct users, the Web is one of the most successful architectures for public and private information systems. With more than 30 years of existence, the Web architecture and standards have also passed the test of time. In parallel, in today's increasingly complex digital landscape, knowledge graphs play a pivotal role for organizing and contextualizing data within information systems: A knowledge graph is a labeled and oriented multi-graph, representing entities and relationships between them, and constrained by formal logic vocabularies (schematas, ontologies, thesaurii) grounding their semantics. It provides a structured, flexible, and interconnected representation of knowledge that facilitates heterogeneous data integration, efficient data retrieval, analysis, and decision-making processes. In this context, we study knowledge-based information system (KBIS), i.e. information systems that leverage knowledge to provide enhanced data management and decision support capabilities. More precisely we focus on Web-based information systems. We rely on and contribute to knowledge graph methods, open standard formalisms and

human-centered approaches. We deploy and evaluate the software implementing our models and methods in different application domains (e.g., biomedical and healthcare, enterprise management).

The first Wimmics team was created on July 2013. It proposed to study models and methods to bridge formal semantics (i.e. with logical foundations, e.g. FOL logics) and social semantics (i.e. emerging from social interactions, e.g. as a social network) on the Web. Its research topic was initially focused on graph-oriented knowledge representation, reasoning and operationalization to model and support actors, actions and interactions in web-based epistemic communities. The main application was to support and foster interactions in online communities and manage their resources. Until 2025, the team has kept its core topics (representing, processing and interacting with knowledge graphs on the Web) and developed new ones (argumentation, natural language processing, online music communities). This enabled to create the MARIANNE team in February 2026 on argumentation and natural language processing.

The present Wimmics team created in February 2026 is refocused on and extends the core topics of the first Wimmics team (and of its ancestor teams Edelweiss and Acacia): graph-based Knowledge Representation and Reasoning (KRR). It is at the crossroad of several research fields related to Knowledge-based Information Systems (KBIS) on the Web. Compared to Wimmics 1.0, we move from KRR focusing on logical consistency and reasoning over well-defined ontologies and schemas, to Knowledge Representation and Artificial Intelligence (KRAI) methods. We focus on (ecosystems of) knowledge graph-based KBIS on the Web. Leveraging and combining the variety of AI methods, schemata, and data, we develop new (meta) models and new hybrid AI methods for such KBIS to support all the different stages of their life-cycle and human interaction.

2.2 Research Topics

Overall we target models and methods for improving information systems on the Web in all their tasks: extracting, storing, validating, querying, exploring and enriching knowledge. However, such information systems and their related tasks are associated with several challenges, either arising from data (e.g., heterogeneity, incompleteness), processes (e.g., uncertainty), or human interaction (e.g., explainability, traceability). Such challenges are addressed in the three topics: Topic 1 focuses on Knowledge Engineering methods and data models and representations; Topic 2 focuses on intelligent data processing techniques; and Topic 3 focuses on the interactions techniques, with the shared common objective of contributing to the provision and exploitation of knowledge graphs in information systems.

It should be noted that these three topics share common foundations, the first one being the use of graph-based formalisms (data models, schemas and syntaxes) and in particular oriented multi-graphs labeled by ontologies. This model is common to all three axes and to all Wimmics members. A second common base is the systematic use of Web architecture and standards in the design of our methods and their implementation. From this stems a third common ground, which is the compatibility and therefore the possible combination of our different methods. To perform or improve methods for the different tasks and for the different stages in the life-cycle of knowledge graphs, we propose AI techniques leveraging and combining methods from both symbolic and non symbolic AI (topic 2). We explore new KRAI approaches such as graphs of graphs models (topic 1) and hybrid graph processing methods (topic 2). To perform their role as information systems, the developed applications must provide efficient human interaction means with their data and processes (topic 3) for which, again, we leverage knowledge graphs features (topic 1) and adequate AI techniques (topic 2). All the symbiotic AI-human interactions envisaged are mediated by the Web architecture, be it a public World Wide Web or a Company Web, which opens up the challenges of decentralization, access rights and federation, addressed in topic 1 and topic 2. In short, our work on the life cycle of knowledge graphs (topic 1), combined intelligent methods (topic 2) and interaction with these graphs (topic 3) complement and reinforce each other.

3 Research program

3.1 Knowledge Graph Life Cycle with a view on Data Integration

Knowledge graphs (KGs) follow a more or less stable life-cycle: they are modeled, populated, validated, published, exploited and maintained. Our first topic of research are the models and methods required to

support that life cycle. Since the initial methods for ontology design and knowledge modeling of the 90s, many innovations happened including crowd-sourcing approaches, methods inspired from agile development and, more recently, new methods based on Large Language Models (LLMs) [112, 83]. There is a need to review and an opportunity to combine the latest and past contributions to offer new and more integrated methods leveraging the best of all approaches to scale, speed-up and reduce the cost of knowledge representation, acquisition and maintenance from heterogeneous sources.

New approaches to Knowledge Engineering (KE): We aim to enable the capitalization of uniform and standard-based methods for KG construction and management. Modularizing and parameterizing KE pipelines to construct and incrementally refine KGs is key to reusing them, when the data sources and processes are well characterized, such as in scientific domains. We investigate various promising lines of research among which the development of a collaboration model providing “Git-for-KG” [93] features (branch, diff, merge, etc.) to manage the KG life-cycle. This implies the construction of a KG metamodel that integrates a logical contract on the data. We envision a Web of multi-modal KGs on several abstraction levels: (i) KGs describing specific knowledge units, (ii) KGs interconnecting such descriptions, (iii) meta-KGs describing knowledge sources [101] with meta-knowledge units or summaries, (iv) KGs interlinking such meta-KGs. Such an ecosystem of KGs facilitate the development of intelligent processing throughout the KG life-cycle by efficiently capturing the relevant components. Relatedly, we investigate how to handle in a unified view both declarative and procedural knowledge (inference or validation rules, transformation rules, and more generally KG pipelines). This entails developing meta-KGs that annotate procedural knowledge, which can then be processed at a higher level of abstraction to facilitate the reuse of this knowledge. Moreover, we aim to develop human-centric KE methods that are key to explainable and trustworthy AI. To this aim, we contribute to KG evaluation methods, considering transparency, provenance, FAIRness (Findability, Accessibility, Interoperability, Reusability), accountability, trustworthiness of KGs, completeness, and facilitate the generation of (PROV-O) traces for the different stages of KE.

Knowledge Extraction from various data types: A KG life cycle is initiated by the extraction of schemata and facts from various types of data. Knowledge extraction can also be a means to enrich existing KGs. We explore the automatic extraction of relations and the generation of triples in the Resource Description Framework (RDF) from text using generative pre-trained language models (PLM). Another topic of interest lies in knowledge extraction from heterogeneous data (tabular data, texts, code) to construct or enrich KGs. In particular, tabular data require the development of specific extraction methodologies to address their specificities (e.g., limited context, different forms of tables). A promising approach consists in performing a synergic extraction from related tabular and textual data, since a text can provide additional context (e.g., describing the structure of a table), while tables can help focus on entities of interest in text.

Open-standard representation for Web-based Knowledge Representation (KR): We specialize in ontology-based and graph-based knowledge representation. We consider KR languages with theoretical foundations in logics (e.g., Description Logics) and graph structures (e.g., Conceptual Graphs). Furthermore, we are situated in Web-based KR, meaning that we use, contribute and extend Web standard languages for KR: RDF (directed labeled multi-graph data model), SPARQL (query language), SHACL (validation schema model), RDFS/OWL (ontological and inferential schema), and extensions (e.g., RDF-star, Canonical RDF) [75]. This also means that we adopt Linked Data best practices (e.g., URI-based identification and access) and FAIR data principles in the production and publication of KGs in the perspective of open reproducible science. We are also concerned with handling local close-world assumption on a KG (contextual knowledge) vs the open world assumption on the Linked Data, depending on use case requirements.

Uncertainty and Neuro-Symbolic-Enhanced Knowledge Representation: One special line of research we investigate is the representation of uncertainty, which is pervasive in real-world applications. There are different types of uncertainty to account for in KGs: epistemic uncertainty stemming from lack of sufficient knowledge and tightly related to incompleteness, and ontic uncertainty stemming from a phenomenon or system being inherently random. Uncertainty can lie in the factual or domain knowledge to be represented in a KG [82], arise from the integration of multiple and potentially conflicting or incoherent data sources [87], or from the knowledge extraction methods used to construct it. We investigate the many facets of uncertainty and propose extensions to the semantic Web standards leveraging probability theory and its extensions to explicitly account for uncertainty and make its treatment possible [86]. Relatedly, we also explore the conceptualization of ontological models that combine both symbolic and neuro-symbolic aspects (e.g., definition by analogy, by perception) and propose extensions to semantic Web standards to represent it. This facilitate KG processing in scenarios where uncertainty, incompleteness or inconsistency must be captured

(Topic 2).

Decentralized provision of Knowledge graphs: In many real-world use cases, information systems integrate data distributed among several providers. RDF is used as a unifying data model for integrating heterogeneous sources, either by materialization (converting the original source into an RDF native storage) or virtualization (querying the original data in SPARQL). Knowledge may be integrated into a unique central KG or distributed among several KGs for maintainability or privacy purposes, or because of the coexistence of multiple viewpoints (e.g., personal KGs). We study the publication and availability (via indexation [100] and discoverability [95]) of KGs with a view on decentralized approaches for the publication and sharing of both graphs datasets and schemata, including Solid and multi-agent systems architectures [113]. We also study how to propose several RDF views of the same data (e.g. JSON data with multiple JSON-LD contexts), while maintaining a form of consistency across them. Another challenge lies in managing various types of schemata associated with the same data, designed for entailment (RDFS, OWL, rules) or validation (XML Schema, JSON Schema, SHACL, etc.). It requires to organize, reconcile and synchronize them, possibly by considering existing well-adopted languages like UML.

3.2 Combined intelligent methods for heterogeneous knowledge graphs

Knowledge graphs operations: We have an extensive and continued experience on the classical tasks performed on knowledge graphs. **Querying** KGs is a core task in all use cases and we have been early contributors of methods to query RDF data, starting from preliminary works towards the birth of the SPARQL standard and the implementation of the latter in the Corese semantic engine [76][80]. We address many aspects of querying, e.g. approximate querying [77], querying heterogeneous data [105][106] or Web API [104], and federated querying to handle use cases with decentralized KGs [78]. We also develop languages on top of SPARQL to define and execute transformations and functions on RDF data [79][81]. A second classical task is **validating** KGs against domain knowledge in the form of constraints, for which we have a longstanding research line. We early contributed to develop methods to validate RDF graphs, e.g., [119], and we implemented the SHACL standard in Corese. We also propose extensions, e.g., to deal with uncertainty [88]. Another classical task is **KG mining** that may be viewed as a special case of knowledge discovery from data, where the data is a KG and the new knowledge can take the form of OWL axioms, SWRL rules, or SHACL shapes. Again we early contributed to develop methods for ontology mining [84] and continued to contribute to this day, dealing with errors in the KG [115], progressively considering more complex OWL axioms [109] and enlarging the scope to SHACL constraints [89]. Finally another task is **KG enrichment**. This can be done by using external data sources with approaches close to those for KG construction (Topic 1), by discovering links with other KGs (KG interlinking, ontology alignment) or within the KG itself, using various reasoning or learning approaches [85][92][97].

Intelligent methods: We are clearly moving from Knowledge Representation and Reasoning (KRR) to Knowledge Representation and Artificial Intelligence (KRAI) because real applications need more than classical logical reasoning. Our aim is to contribute to a variety of AI methods based on reasoning (logical, analogical, approximate, plausible, etc.) or learning (symbolic, statistical, neural networks, genetic programming), and their combinations in hybrid AI methods, including neuro-symbolic approaches. We focus on AI methods compatible with core ontology-based and graph-based KR. We contribute to intelligent methods for the above described classical tasks on KGs, studying the learning, alignment and application of validation or inferential schemata to a KG, and the integration, interaction and enrichment of KGs with different types of intelligent processing. The development of hybrid AI methods combining well known AI methods is a long term research direction that we want to investigate as a priority.

We contribute to the development of neuro-symbolic AI combining symbolic reasoning and machine learning. Such a combination may involve enhanced loss functions, the analysis of the semantics of latent spaces to retrieve and discover formal knowledge, and the combination of LLMs with symbolic knowledge and associated deductive reasoning. We also design neural-based approaches for efficient and flexible KG querying by retrieving elements from the latent space. Additionally, we aim to enable interactions between different types of intelligent processing and characterizing their common or emerging properties, e.g. similarly to studies showing the interest of applying deductive rules before learning graph embedding spaces. We study such interactions by modeling intelligent processing as agents, enabling their interactions through knowledge graphs and with the objective to refine them. We envision an orchestrator module for intelligent processors. Given a task or a query, it should assess the capacity of various intelligent processings to answer,

transferring them all or parts of the query, and fusing their answers. In this view, KGs constitute both the input of intelligent processors but also the pivotal exchange structure for them to share their inferences, and consolidate their respective results. To make KGs the unifying structure, we also contribute to Topic 1 by extending KR languages to represent both knowledge and intelligent processors within the same space.

Transversal characteristics addressed: We study knowledge-graph based information systems on the Web. Consequently, we address specific challenges when developing AI methods. We aim to develop **Web oriented** AI methods and implementations, compliant with Web software architectures, Web languages and standards, Web protocols, etc. The Web is natively decentralized, with decentralized data, decentralized services, and decentralized stakeholders. It therefore requires decentralized AI (decentralized processing, architectures, etc.). In this decentralized landscape, the Web of Linked Data is designed to be decentralized and heterogeneous, and yet most RDF-based applications still rely on a central SPARQL endpoint and relatively homogeneous data. We explore the challenges of building truly decentralized applications on top of the Web of Linked Data. We aim to explore alternative **decentralized approaches** such as graph traversal queries solving on decentralized linked data sources (e.g., querying over a Solid ecosystem) or hypermedia multi-agent systems (hMAS) [73, 91] for collaborative task solving. The **heterogeneity** of knowledge graphs on the Web comes from many aspects (size, dynamics, data source quality, processing quality, degree of uncertainty, etc.) which are as many dimensions as we have to take into account in developing intelligent methods. Hence, we develop uncertainty-aware intelligent methods for KGs processing, with algorithms to integrate uncertainty in querying, deduction and embedding of KGs. We also aim to establish criteria for selecting KGs to use based on uncertainty and provenance metadata, as well as other types of metadata, enabling users and agents to make informed decisions regarding trust and data application.

On another note, the imperative for **explainable and trustworthy** AI systems is now a well known and still open question. It is paramount to ensure transparency, accountability, and user confidence in the decisions and actions driven by these systems. These are challenges that we address when developing AI methods for KG-based information systems. The heterogeneity of the data sources and methods used to construct or enrich knowledge graphs ask for even more traceability and explainability means. Automatically discovered knowledge can only be provisional and should be trusted only to a certain degree. We investigate how to take into account the provenance and trustworthiness of knowledge graphs when processing them (logical or approximate reasoning, mining, etc.). We rely on a multi-level KG ecosystem (see Topic 1) where KGs themselves are annotated. We also investigate how to make explainable and trustworthy the knowledge engineering pipelines to construct, refine and process KGs. This involves tracing and representing KG processes into meta-KGs within the multi-level KG ecosystem (see Topic 1). We leverage KGs to study the semantic understanding of different AI models, e.g., by mining latent spaces to analyze the ontological knowledge that was captured and discover new knowledge. Models are also evaluated based on several properties (e.g., type hierarchy, composition) measured on a post-hoc basis from their predictions. Additionally, we propose knowledge mining techniques that can be explained to humans (e.g. learn to define by analogy, learn rules).

We investigate using KGs as pivotal structure toward symbiotic AI-human interactions. KGs thus form the exchange structure between humans and artificial agents to represent, describe, and explain both their knowledge and their intelligent processing. In this view, humans are able to use, understand, and trust AI results and AI agents have access to human knowledge to exploit and analyze it, e.g., to guide their own learning.

3.3 Interaction design for decision making on and with knowledge graphs

Visualizations, interfaces and interaction design are vital to information systems as they are here to support humans in their tasks. Because they support information systems, knowledge graphs and their processing require interactions with users (construction, visualization, contribution, validation, analysis, etc.). Users must play a central role by expressing their requirements for interacting with data sources and by communicating the findings obtained through the interaction with knowledge graphs. This last topic focuses on knowledge graph-based and AI-leveraging techniques to design these interactions and to support decision-making with knowledge graphs (providing explanations, justifications, traces, provenance, etc.).

Visual representation and exploration of knowledge graphs: We capitalize our past experiences in various research projects, e.g., [103][72], where visualizations techniques were needed to communicate the content of KGs to the users and support end-user interaction with KGs. We aim to design and develop a

generic library for Semantic Web developers to help designing visualizations of the results of queries over KG, acting as an interface with possibly various data visualization libraries. We also design and develop specific visualization methods for special kinds of KGs, among which KGs representing **spatio-temporal** data, and KGs representing validation schemata (e.g. SHACL) and inferential schemata (RDFS, OWL, rules) (longer term). We study the usage of chained visualizations to explore KGs [102]: our aim is to develop methods and tools for describing and implementing complex KG visualization **pipelines** to support KG exploration. We also investigate the visualization of knowledge graphs **ecosystems**, e.g. the combined visualization of graph data and their annotation (e.g. provenance information) to support decision making through the explanation of the processes and sources that led to a result.

Natural language (NL) access and multimodal access to knowledge graphs: We study AI-based techniques for designing interactions with KGs, starting with combining latest NLP methods and interaction design to support natural language based querying and manipulation (summarizing, validation, mining, etc.) of knowledge graphs, possibly combining and synchronizing several views/widgets such as an LLM-based chatbot [117], and one or more specialized views (charts, maps, graph visualizations, tables, etc.).

Knowledge acquisition through user interaction: Users might contribute to the different steps of the knowledge-graph life-cycle, in particular by supporting the acquisition of knowledge represented in graphs. A typical example of user contributions are user-made annotation and labeling, which are still the most reliable means to obtain quality data to create knowledge [116]. It is worth to notice that annotation tasks (i.e. an explicit method for knowledge acquisition) can be coupled with the analysis of direct interaction with data (e.g., analysis of user queries, frequent concepts/path explored during the interaction, etc.) to reveal implicit user intentions and behavior patterns. We study knowledge acquisition through user interaction with knowledge graphs, which has a two-fold contribution: on one hand, it deepens the understanding of how users can express and understand knowledge embedded into knowledge graphs; on the other hand, it can improve our understanding of the user needs towards interacting with knowledge graphs, which can ultimately be used to develop usable interactive tools allowing users to contribute to personalized experiences with data [90]. Identifying user behavior patterns also contributes to knowledge acquisition and formalization of how users interact with systems, which has a multi-fold implication in the design of Web applications [114].

Key characteristics addressed: Generally, as a counterpart of our scopes in Topics 1 and 2, our focus here is web-based interfaces and interaction design for the processing of decentralized KGs with decentralized approaches. In particular, we investigate interactions with personal KGs. The **Solid** project aims to change the way we build Web applications, by putting the data under the control of the user. In this view, every user has a personal KG (Solid pod) which different applications use and contribute to. But to effectively apply that control, the average user must be able to understand the global structure of their personal KG and to monitor how applications use and change it. Our aim is to develop interaction models and tools so that users can fully benefit from Solid.

On another note, we adopt a reflexive approach of interaction design for KG visualization, exploration and processing: We represent, in a dedicated special type of meta-KG of our ecosystem of multi-level knowledge graphs, user information describing their exploration processes, interactions and insights obtained during the exploration process within visualizations either in single or collaborative environments. Then we propose intelligent methods, leveraging reasoning and learning, to process the resulting KG to support the improvement of interaction and visualization techniques as well as to facilitate knowledge sharing and communication within collaborative environments.

4 Application domains

Our application domain is the one of information systems (IS) defined as “sociotechnical, organizational systems designed to collect, process, store, and distribute information [and] used to provide information, contribute to knowledge as well as digital products that facilitate decision making” in the [Wikipedia article on Information System](#). More precisely we contribute to Web-based information systems relying on graph knowledge bases.

We study these knowledge-based information systems (KBIS), and deploy and evaluate the software implementing our methods in different scenarios and contexts including data integration (project D2KAB), federated services (projects Startin’ Blox, Dekalog, EAESI), scientific data publication and access (projects ISSA for Digital Libraries, D2KAB for Scientific Literature Review, HISINUM for Digital Humanities), life

science data (projects TaxRef on MNHP taxonomic referential, D2KAB in genomics and agrometeorology, MetaboLinkAI in metabolomics), biomedical data (PGxLOD knowledge graph in pharmacogenomics), real estate intelligence (project MUSE), analytical provenance, etc.

5 Social and environmental responsibility

5.1 Footprint of research activities

The team now integrates footprint metrics in its evaluations and comparisons of methods. For instance, in [111], we monitored the training time and the carbon cost for training a knowledge graph extraction model.

5.2 Impact of research results

We are especially interested in identifying Knowledge Graphs tasks for which SMLs (Small Language Models) can be efficiently used.

6 Highlights of the year

The present Wimmics team was created on February 2025, with Catherine Faron as its scientific leader. Fabien Gandon became co-chair of the W3C [Web & AI Interest Group](#) created on October 2025.

6.1 Awards

- Best paper award at the RAGE-KG 2025 Workshop for "User Interface and Agent Interface for Online Generation of Knowledge Graph's Competency Questions and Question-Query Training Sets" by Yousof Taghzouti, Franck Michel, Tao Jiang, Louis-Felix Nothias and Fabien Gandon. [69].
- Best paper award at the 6th IEEE International Symposium on the Internet of Sounds, for the paper "Interactive Audio Sculpting Plugin Customization and UI Affordances in Immersive Environments" by Michel Buffa, Marco Winckler, Quentin Escobar, Samuel Demont, Ayoub Hofr and Adam Mir-Sadjadi. [39].

7 Latest software developments, platforms, open data

7.1 Latest software developments

7.1.1 CORESE-Core

Name: COnceptual REsource Search Engine - Core

Keywords: Semantic Web, RDF, RDFS, SPARQL, OWL, SHACL, Automated Reasoning, Validation, Interoperability, Linked Data, Knowledge Graphs, Knowledge Bases, Knowledge representation, Querying, Ontologies

Scientific Description: CORESE-Core is a library used in research to apply and evaluate Semantic Web standards and the algorithms they require. It is also the basis for proposing and prototyping extensions to these standards and their processing.

Functional Description: CORESE-Core is a library that implements and extends the Semantic Web standards established by the W3C, such as RDF, RDFS, SPARQL1.1 Query & Update, OWL RL, SHACL, and others.

This library offers a wide range of features for creating, manipulating, parsing, serializing, querying, reasoning and validating RDF data.

In addition, it offers advanced extensions such as STTL, SPARQL Rule and LDScript, which extend the functionality and processing capabilities of the data.

NB: CORESE-Core is a library derived from the earlier CORESE software.

News of the Year: <https://github.com/corese-stack/corese-core/blob/develop/CHANGELOG.md>

URL: <http://project.inria.fr/corese>

Contact: Remi Ceres

Participants: Remi Ceres, Fabien Gandon

7.1.2 CORESE-GUI

Name: COnceptual REsource Search Engine - Graphical User Interface

Keywords: GUI (Graphical User Interface), User Interfaces, Knowledge Bases, Knowledge Graphs, Knowledge graph, Knowledge representation, Ontologies, Linked Data, Validation, Automated Reasoning, SHACL, OWL, SPARQL, RDFS, RDF, Querying, Applications

Scientific Description: CORESE-GUI is a graphical user interface developed to interact with the CORESE-Core library. It provides users, especially those less experienced in programming, with an intuitive and visual access to the functionalities of CORESE-Core. This interface includes tools for visualizing semantic data, editing SPARQL queries, and monitoring data processing results. CORESE-GUI also serves as a platform for experimenting with new extensions and processing methods in the field of semantic web, thereby making these technologies more accessible to researchers and practitioners.

Functional Description: This desktop application allows the user to call up CORESE-Core features for creating, manipulating, parsing, serializing, querying, reasoning and validating RDF data.

The application enables direct use of Semantic Web languages standardized by the W3C, such as RDF and its syntaxes, RDFS, SPARQL1.1 Query & Update, OWL RL, SHACL, and others.

News of the Year: <https://github.com/corese-stack/corese-gui-swing/blob/develop/CHANGELOG.md>

URL: <http://project.inria.fr/corese>

Contact: Remi Ceres

Participants: Remi Ceres, Fabien Gandon

7.1.3 CORESE-Server

Name: COnceptual REsource Search Engine - Server

Keywords: Server, Linked Data, Semantic Web, Ontologies, Knowledge Graphs, Knowledge Bases, RDF, RDFS, SPARQL, SHACL, Querying, Validation, Automated Reasoning

Scientific Description: This server version allows remote applications to access CORESE-Core functionalities for creating, manipulating, analyzing, serializing, querying, reasoning, and validating RDF data. The server facilitates remote use of W3C-standardized Semantic Web languages, such as RDF and its syntaxes, RDFS, SPARQL1.1 Query & Update, OWL RL, SHACL, and more.

Functional Description: This server version enables a remote application to call CORESE-Core's functions for creating, manipulating, analyzing, serializing, querying, reasoning and validating RDF data.

The server enables remote use of Semantic Web languages standardized by the W3C, such as RDF and its syntaxes, RDFS, SPARQL1.1 Query & Update, OWL RL, SHACL, and others.

News of the Year: <https://github.com/corese-stack/corese-server/blob/develop/CHANGELOG.md>

URL: <https://project.inria.fr/corese/>

Contact: Remi Ceres

Participants: Remi Ceres, Fabien Gandon

7.1.4 CORESE-Command

Name: COnceptual REsource Search Engine - Command Line

Keywords: Command, RDF, RDFS, SPARQL, SHACL, Knowledge acquisition

Scientific Description: This command-line version of CORESE enables users to incorporate CORESE-Core functionalities into scripts, workflows, and consoles for creating, manipulating, analyzing, serializing, querying, reasoning, and validating RDF data. It allows direct use of W3C-standardized Semantic Web languages, such as RDF and its syntaxes, RDFS, SPARQL1.1 Query & Update, OWL RL, SHACL, and more.

Functional Description: This command-line version enables users to call CORESE-Core's functionality in scripts, workflows and console mode for the creation, manipulation, analysis, serialization, querying, reasoning and validation of RDF data.

The command enables direct use of W3C-standardized Semantic Web languages, such as RDF and its syntaxes, RDFS, SPARQL1.1 Query & Update, OWL RL, SHACL, and others.

News of the Year: <https://github.com/corese-stack/corese-command/blob/develop/CHANGELOG.md>

URL: <https://project.inria.fr/corese/>

Contact: Remi Ceres

Participants: Remi Ceres, Fabien Gandon

7.1.5 Datalens

Keywords: Data visualization, Artificial intelligence

Functional Description: Datalens leverages custom network topologies, multi-faceted filters, and advanced visualization techniques to help users discover relevant datasets published online for their specific tasks. It harnesses the visualization capabilities of MGExplorer to enable a multi-perspective exploration of data. Currently, the tool supports navigation through datasets and models available on HuggingFace.

URL: <https://dataviz.i3s.unice.fr/datalens/>

Contact: Aline Menin

7.1.6 DBpedia

Name: DBpedia

Keywords: RDF, SPARQL

Functional Description: DBpedia is an international crowd-sourced community effort to extract structured information from Wikipedia and make this information available on the semantic Web as linked open data. The DBpedia triple stores then allow anyone to solve sophisticated queries against Wikipedia extracted data, and to link the different data sets on these data. The French chapter of DBpedia was created and deployed by Wimmics and is now an online running platform providing data to several projects such as: QAKIS, Izipedia, zone47, Sépage, HdA Lab., JocondeLab, etc.

URL: <http://wiki.dbpedia.org/>

Contact: Fabien Gandon

Participants: Fabien Gandon, Elmahdi Korfed

7.1.7 Gen²KGBot

Name: Gen²KGBot

Keywords: Knowledge graph, LLM, SPARQL

Functional Description: Gen²KGBot intends to allow users to "speak to a knowledge graph", that is, use natural language to query knowledge graphs in a generic manner, with the help of generative large language models (LLM).

It provides a generic framework to translate a natural-language (NL) question into its counterpart SPARQL query, execute the query and interpret the SPARQL results.

News of the Year: The latest evolutions include the coupling with Q²Forge and the addition of new agentic scenarios to improve the quality of generated queries.

Publications: [hal-05332768](#), [hal-05070442](#), [hal-05399933](#), [hal-05333244](#)

Contact: Franck Michel

Participants: Yousouf Taghzouti, Franck Michel, Tao Jiang, Louis-Felix Nothias, Fabien Gandon

Partner: Institut de Chimie de Nice

7.1.8 GUsT-3D

Name: Guided User Tasks Unity plugin for 3D virtual reality environments

Keywords: 3D, Virtual reality, Interactive Scenarios, Ontologies, User study

Functional Description: We present the GUsT-3D framework for designing Guided User Tasks in embodied VR experiences, i.e., tasks that require the user to carry out a series of interactions guided by the constraints of the 3D scene. GUsT-3D is implemented as a set of tools that support a 4-step workflow to : (1) annotate entities in the scene with names, navigation, and interaction possibilities, (2) define user tasks with interactive and timing constraints, (3) manage scene changes, task progress, and user behavior logging in real-time, and (4) conduct post-scenario analysis through spatio-temporal queries on user logs, and visualizing scene entity relations through a scene graph.

The software also includes a set of tools for processing gaze tracking data, including: cleaning and synchronization of the data, calculation of fixations with I-VT, I-DT, IDTVR, IS5T, Remodnav, and IDVT algorithms, and visualization of the data (points of regard and fixations) in both real time and collectively.

URL: <https://project.inria.fr/creattive3d/gust-3d/>

Publications: [hal-04102737](#), [hal-04446066](#), [hal-03635452](#)

Contact: Hui-Yin Wu

Participants: Hui-Yin Wu, Marco Alba Winckler, Lucile Sassatelli, Florent Robert

Partner: I3S

7.1.9 IndeGx

Keywords: Semantic Web, Indexation, Metadata

Functional Description: IndeGx is a framework for the creation of an index of a set of SPARQL endpoints. The framework relies only on available semantic web technologies and the index appears as an RDF database. The index is primarily composed of the self-description available in the endpoint. This original description is verified and expanded by the framework, using SPARQL queries.

Release Contributions: The previous version was a Java application coded with Apache Jena, this version uses an engine coded in Typescript with rdflib, graphy, sparqljs, coupled with a Corese Server, in a docker application. - Treatment of endpoints in parallel -The automatic pagination of simple queries to avoid overwhelming SPARQL endpoints. - The usage of Corese as an interface with SPARQL endpoints to reduce missing data due to errors coming from incorrect standard compliance in distant SPARQL endpoints. - Rules are now expected to make heavy use of federated querying, with the SERVICE clause. - Possibility to define the application of several rules as a prerequisite to the application of another. - End of the difference between CONSTRUCT and UPDATE rules to differentiate between the application of local and distant queries. Only test queries are supposed to be SELECT, ASK, or CONSTRUCT. All action queries are expected to be UPDATE queries. - Possibility to define a set of rules as a pre-treatment or a post-treatment on the extracted data. In this case, the endpoint URL becomes the URL of the local corese server (not accessible from the outside of the docker) - Handling many different errors in the RDF format of data found in remote endpoints - Possibility of disabling the query logging of the framework - Possibility of using the query logging of the framework to avoid repeating rule application in case of an execution interruption - Integration of LDscript in rules possible.

We also offer two automatically refreshed catalogs, - The catalog of endpoints taken from numerous sources, updated daily - The catalog of endpoints and their statuses, refreshed hourly

News of the Year: Full rewriting of the software. See details in release contribution.

URL: <https://github.com/Wimmics/IndeGx>

Publication: [hal-03946680v1](#)

Contact: Pierre Maillot

Participants: Pierre Maillot, Fabien Gandon, Catherine Faron, Olivier Corby, Franck Michel

7.1.10 KartoGraphI

Keywords: SPARQL, Linked Data, Indexing

Functional Description: Website displaying a screenshot of the state of the Linked Data web according to the description retrieved by the IndeGx software

News of the Year: Documentation updates

URL: <https://github.com/Wimmics/KartoGraphI>

Publication: [hal-03652865](#)

Contact: Pierre Maillot

Participants: Pierre Maillot, Fabien Gandon, Catherine Faron, Olivier Corby, Franck Michel

7.1.11 Metadatamatic

Keywords: RDF, Semantic Web, Metadata

Functional Description: Website offering a form to generate in RDF the description of an RDF base.

URL: <https://wimmics.github.io/voidmatic/>

Contact: Pierre Maillot

Participants: Fabien Gandon, Franck Michel, Olivier Corby, Catherine Faron

7.1.12 MGExplorer

Name: Multivariate Graph Explorer

Keywords: Information visualization, Linked Data

Scientific Description: MGExplorer (Multidimensional Graph Explorer) allows users to explore different perspectives to a dataset by modifying the input graph topology, choosing visualization techniques, arranging the visualization space in meaningful ways to the ongoing analysis and retracing their analytical actions. The tool combines multiple visualization techniques and visual querying while representing provenance information as segments connecting views, which each supports selection operations that help define subsets of the current dataset to be explored by a different view. The adopted exploratory process is based on the concept of chained views to support the incremental exploration of large, multidimensional datasets. Our goal is to provide visual representation of provenance information to enable users to retrace their analytical actions and to discover alternative exploratory paths without losing information on previous analyses.

Functional Description: MGExplorer is an information visualization tool designed for exploring multivariate graphs, integrating various visualization techniques. It allows users to select and combine these techniques into a graph that traces the exploration path of a database. Developed with the D3.JS library, MGExplorer runs directly in a web browser. The tool is available online and can be customized using SPARQL queries created and managed within the LDViz software, which facilitates the creation, storage, and management of such queries. Additionally, MGExplorer can be integrated into any web project as an npm package, providing a modular solution for data visualization.

Release Contributions: MGExplorer is now available as a web component, making it easy to integrate into any web project via an npm package, accessible at <https://www.npmjs.com/package/mgexplorer>. It can be customized to visualize either local datasets or results from SPARQL queries.

News of the Year: The software has been restructured and converted into an npm package (<https://www.npmjs.com/package/mgexplorer>) that can be integrated into any web page. It now allows data from any knowledge graph to be visualised, whether public or private (with local use and configuration provided by the developer).

URL: <https://dataviz.i3s.unice.fr/mgexplorer>

Publications: [hal-03518845](#), [hal-03292172](#), [hal-03404572](#), [hal-03404580](#)

Contact: Aline Menin

Participants: Aline Menin, Marco Alba Winckler, Olivier Corby

Partner: Universidade Federal do Rio Grande do Sul

7.1.13 Muvin

Name: Multidimensional Visualization of Networks over Time

Keywords: Data visualization, LOD - Linked open data, Temporal Networks

Scientific Description: Muvin addresses the challenges of visualizing complex collaboration networks by implementing an incremental approach tailored for exploring co-authorship networks composed of multivariate entities distributed over time. Traditional representations of such networks can become visually cluttered, making it difficult to focus on relevant information. To tackle this, Muvin employs a focus+context technique, allowing users to zoom in on specific data points while maintaining an overview of the broader network. By enabling incremental data exploration and supporting multi-layered linked open data (LOD), Muvin effectively handles the complexity and scalability issues of collaboration networks. This approach intends to facilitate domain-specific tasks, such as identifying influential collaborators and understanding knowledge dissemination in co-authorship networks.

Functional Description: Muvin facilitates the exploration of a two-layer network that captures collaborations among entities such as researchers, artists, keywords, and more, as well as the temporal evolution of related elements, including scientific publications or songs. The tool adopts an incremental approach, enabling users to dynamically import data from a SPARQL endpoint into the exploration workflow. SPARQL queries can be created and adjusted on the fly using the LDViz query management tool, allowing users to experiment with different queries to address specific data-related questions. Developed with the D3.js library for visualization, Muvin is designed primarily for exploring data from knowledge graphs. The tool is accessible online at <https://dataviz.i3s.unice.fr/muvin>.

News of the Year: The software has been restructured and converted into an npm package (<https://www.npmjs.com/package/muvin>) that can be integrated into any web page. It now allows data from any knowledge graph to be visualised, whether public or private (with local use and configuration provided by the developer). In addition, the software has been the subject of a user study aimed at evaluating the value of the approach implemented for exploring collaborative networks that evolve over time.

URL: <http://dataviz.i3s.unice.fr/muvin/>

Publications: [hal-05302463](#), [hal-03748134](#)

Contact: Aline Menin

Participants: Aline Menin, Marco Alba Winckler

7.1.14 Olivaw

Name: Ontology Long-lived Integration Via ACIMOV Workflow

Keywords: Ontologies, Ontology engineering, Semantic Web, Git svn, Linked Data, LOD - Linked open data, Web

Scientific Description: Olivaw proposes: (1) command lines that make an Acimov ontology development easier, (2) composite actions that can directly be called in workflows from any Acimov project, (3) a pre-commit hook that prevents mistakes from being pushed to an Acimov repository. The test reports are first represented using the EARL vocabulary and then exported in the markdown format to fit a github environment. A template repository also exists in order for an ontology project to begin with the accurate repository architecture, workflows and special files.

Functional Description: Agile and collaborative approaches to ontology development are crucial because they contribute to making them user-driven, up-to-date, and able to evolve alongside the systems they support, hence proper continuous validation tooling is required to ensure ontologies match these standards all along their development. We propose OLIVAW (Ontology Long-lived Integration Via ACIMOV Workflow), a tool supporting the ACIMOV methodology on GitHub. It relies on W3C Standards to assist the development of modular ontologies through GitHub Composite Actions, pre-commit hooks, or a command line interface. OLIVAW was tested on several ontology projects to ensure its usefulness, genericity and reusability. A template repository is available for a quick start. OLIVAW is published under the LGPL-2.1 license and archived on Software Heritage and Zenodo.

URL: <https://github.com/Wimmics/olivaw>

Publications: [hal-04823080](#), [hal-05291789](#)

Contact: Nicolas Robert

Partner: IMT - MINES Saint-Étienne

7.1.15 Q²Forge

Name: Question-Query Forge

Keywords: Natural language, SPARQL, Knowledge graph, LLM

Functional Description: This project provides an end-to-end pipeline to generate a dataset of (natural language question, SPARQL query) pairs for a Knowledge Graph (KG).

News of the Year: The latest evolutions are toward supporting multiple users and multiple KGs.

Publications: [hal-05070442](#), [hal-05399933](#), [hal-05333244](#), [hal-05332768](#)

Contact: Franck Michel

Participants: Yousouf Taghzouti, Franck Michel, Tao Jiang, Louis-Felix Nothias, Fabien Gandon

Partner: Institut de Chimie de Nice

7.1.16 RDFminer

Keywords: Evolutionary Algorithms, Semantic Web, Web API, Dashboard

Functional Description: RDFminer is an open source Web application to automatically discover SHACL shapes through an evolutionary process. It takes an RDF data graph as input, from which shapes are mined and assessed using a probabilistic validation framework. The user can interact with RDFminer through a dashboard where they can launch and monitor the mining of shapes, and analyse the results in real time.

URL: <http://ns.inria.fr/rdfminer>

Publication: [hal-04566981v1](#)

Contact: Andrea Tettamanzi

Participants: Remi Felin, Thu Nguyen, Andrea Tettamanzi, Catherine Faron, Fabien Gandon

7.1.17 SciLEX

Name: Science Literature Exploration

Keywords: Textmining, Systematic review, Collaborative science, Linked Data

Functional Description: Scilex is a tool allows to start a scientific paper collect to analyse the state of art of a given domain. It also allows the annotation, the enrichment as well as the analysis of the results.

News of the Year: Multiple improvements in the interface with scientific libraries and Zotera

Publication: [hal-04175511](#)

Contact: Celian Ringwald

Participants: Celian Ringwald, Anaïs Ollagnier, Benjamin Navet

7.1.18 SemWebRAG

Name: Semantic Web Retrieval Augmented Generation

Keywords: Semantic Web, Ontologies, Knowledge Graphs, Large Language Models, Retrieval Augmented Generation

Functional Description: SemWebRAG implements a pipeline relying on a Knowledge Graph (KG) for Retrieval Augmented Generation (RAG) with Large Language Models. The KG is built via entity extraction from a corpus of documents, and enrichment with interlinks to ontologies or Wikidata.

Contact: Pierre Monnin

Participants: Pierre Monnin, Fabien Gandon, Krysto Dagues De La Hellerie

7.1.19 Zoomathia KG Pipeline

Name: Automatic annotation of an ancient zoological corpus

Keywords: Zoology, NLP, Semantic annotation, Semantic Web

Functional Description: The Zoomathia corpus contains texts on animals compiled within the framework of the Zoomathia GDRI funded by the CNRS. It aims to support the study of the transmission of zoological knowledge from antiquity to the Middle Ages. This project provides a text processing pipeline for the Zoomathia corpus. It adapts and combines methods from NLP and knowledge engineering to analyze, classify and automate the semantic annotation of the texts. The result is the Zoomathia Knowledge Graph.

News of the Year: Refactoring and significant improvement of the document

URL: <https://github.com/Wimmics/zoomathia>

Contact: Catherine Faron

Participants: Arnaud Barbe, Catherine Faron, Molka Dhouib, Franck Michel

Partner: CEPAM (Cultures, Environnements, Préhistoire, Antiquité, Moyen Âge)

7.1.20 Zoomathia KG Web Application

Name: Web Application for the exploitation of the Zoomathia Knowledge Graph

Keywords: Zoology, NLP, Semantic annotation, Semantic Web

Functional Description: The Zoomathia corpus contains texts on animals compiled within the framework of the Zoomathia GDRI funded by the CNRS. It aims to support the study of the transmission of zoological knowledge from antiquity to the Middle Ages. The Zoomathia KG is a knowledge graph annotating the Zoomathia corpus with concepts from the TheZoo thesaurus. This project provides a web application that allows researchers to explore the Zoomathia KG via a search for works by concept, explore a selected work while visualizing the concepts annotating each of its parts, and visualize the results of queries implementing competency questions on a selected work from the corpus.

News of the Year: Refactoring and significant improvement of the document

URL: <http://zoomathia.i3s.unice.fr/>

Contact: Catherine Faron

Participants: Arnaud Barbe, Catherine Faron, Molka Dhouib, Franck Michel

Partner: CEPAM (Cultures, Environnements, Préhistoire, Antiquité, Moyen Âge)

7.2 Open data

TAXREF-LD: Knowledge Graph of the French taxonomic registry

Contributors: Franck Michel, Catherine Faron

Description: TAXREF-LD is a Linked Data knowledge graph representing TAXREF, the French national taxonomical register for fauna, flora and fungus, that covers mainland France and overseas territories. TAXREF-LD is a joint initiative of the UMS Patrinat of the National Museum of Natural History, and the I3S laboratory, University Côte d’Azur, Inria, CNRS.

Dataset PID (DOI,...): DOI:10.5281/zenodo.12733630

Project link: [TAXREF-LD, 10.5281/zenodo.12733630](https://doi.org/10.5281/zenodo.12733630)

Publications: [108]

Contact: Franck Michel

Release contributions: version 17.0 implements new SKOS collections and better management of vernacular. See full description at <https://github.com/frmichel/taxref-ld/blob/master/CHANGELOG.md>

WheatGenomicsSLKG

Contributors: Nadia Yacoubi Ayadi, Franck Michel, Catherine Faron

Description: Wheat Genomics Scientific Literature Knowledge Graph is a FAIR knowledge graph that exploits the Semantic Web technologies to integrate information about Named Entities (NE) extracted automatically from a corpus of PubMed scientific papers on wheat genetics and genomics. This work is supported by the French National Research Agency under grant ANR-18-CE23-0017 (project [D2KAB](#)).

Dataset PID (DOI,...): DOI:10.5281/zenodo.10420888

Project link: [WheatGenomicsSLKG, 10.5281/zenodo.10420888](https://doi.org/10.5281/zenodo.10420888)

Publications: [118]

Contact: Franck Michel

Release contributions: (first release)

Pharmacogenomics datasets for Ontology Matching

Contributors: Pierre Monnin

Description: These datasets constitute benchmarks to evaluate Ontology Matching algorithms on a complex structure-based instance matching task from the domain of pharmacogenomics. Pharmacogenomics involves n -ary tuples representing so-called “pharmacogenomic relationships” and their components of three distinct types: drugs, genetic factors, and phenotypes. The goal resides in matching such tuples. These datasets were extracted from the [PGxLOD](#) knowledge graph.

Dataset PID (DOI,...): DOI:10.5281/zenodo.8419361

Project link: [PGxLOD, 10.5281/zenodo.8419361](https://doi.org/10.5281/zenodo.8419361)

Contact: Pierre Monnin

Release contributions: this is the first published version.

Semantically Enriched Datasets for Link Prediction: DB100k+, NELL-995+ and YAGO3-10+

Contributors: Nicolas Robert, Pierre Monnin, Catherine Faron

Description: Starting from the widely accepted datasets DB100k, NELL-995 and YAGO3-10, we semantically enriched them with ontological knowledge, namely class hierarchy and relation signatures (domains and ranges), and inferred new entity type assertions to create DB100k+, NELL-995+ and YAGO3-10+. We also provide a generic masking script to generate sub-graphs with variable proportions of triples with signed/partially signed (no domain or no range)/unsigned (no domain and no range) relations, to evaluate the impact of semantic information on the performance of Machine Learning models.

Dataset PID (DOI,...): DOI:10.5281/zenodo.15834518

Project link: [GitHub repository for datasets NELL-995+, DB100k+ and YAGO3-10+](#)

Publications: [52]

Contact: Nicolas Robert

Release contributions: this is the first published version.

DBpedia.fr : French chapter of the DBpedia knowledge graph dataset

Contributors: Fabien Gandon, Franck Michel, Celian Ringwald

Description: The [DBpedia.fr](#) project ensures the creation and maintenance of a French chapter of the DBpedia knowledge base a crowd-sourced community effort to extract structured content from the information created in various Wikimedia projects. Statistics indicate very high usage rate: the server processed 1.8+ billion queries over the year. This represents a 3.86 million daily average and 32.5 million daily max.

Dataset PID (DOI,...):

Project link: <https://fr.dbpedia.org/>

Contact: Célian Ringwald

Release contributions: No new release was done this year but we carried out continuous monitoring and support to ensure a high-availability service.

8 New results

8.1 Knowledge Graph Life Cycle with a view on Data Integration

8.1.1 Semantic Web for the Integration of Pharmacogenomics Knowledge

Participants: Pierre Monnin.

Life sciences produce and consume vast amounts of scientific data. The graph-structured nature of these data naturally leads to data-driven research efforts leveraging Semantic Web and Knowledge Graph technologies. Among such usages, knowledge graph construction and management is a well established topic. One subtask lies in matching similar or related units across datasets to identify possible overlaps and merge multiple sources of knowledge. In this direction, this year again, we proposed the track “**Pharmacogenomics**” in the international challenge “**Ontology Alignment Evaluation Initiative**”. This track focuses on the matching of pharmacogenomic knowledge units, which are n -ary tuples involving components of three distinct types (drugs, genetic factors, and phenotypes). This year again, none of the approaches participating in the 2025 campaign were able to produce alignments [31]. These results highlight once more the interest in considering domain-specific problems, bringing additional challenges to the field of ontology matching. Given the inadequacy of ontology matching systems to produce valid alignments, such challenges are currently unaddressed and require to design new methods or enrich existing ones.

8.1.2 A Semantic Web Ontology for Psychosocial Factors of Dysfunctional Eating Attitudes and Behaviors in Sport

Participants: Molka Dhouib, Catherine Faron.

Dysfunctional eating attitudes and behaviors (DEAB) represent complex phenomena influenced by multiple psychosocial factors within the sport context. Current research in this field is characterized by a high level of complexity and notable inconsistencies regarding the theoretical concepts employed, the measurement tools used, and the relationships between these concepts. In this context, we developed an ontology of psychosocial factors related to DEAB in sport, aimed at improving the understanding of the complexity of interactions among these factors. This work was carried out in collaboration with researchers from the LAMHESS laboratory of University Côte d'Azur. Based on a systematic review of the scientific literature, a knowledge graph was constructed. In addition, a web-based application was developed and made available online to allow researchers to query and explore the knowledge graph [55].

8.1.3 An Ontology for Modeling User Activity within Visualization Interfaces

Participants: Aline Menin, Catherine Faron.

The study of user activity supports evaluation of visualization systems, recommendation of suitable views or tasks, guidance of interaction, and validation of analytical results. It enables researchers to understand how these visualization systems are used and to gain insight into user's reasoning processes during data exploration. However, there is a lack of structured frameworks for systematically collecting and reasoning over such data. Thus, we build upon Semantic Web standards to model and represent user activity as knowledge graphs. We introduced an ontology-based model for representing user activity within visualization systems. The DIVA ontology and thesaurus, built upon established RDF vocabularies and domain-specific taxonomies, captures the provenance of data, user interactions, visualizations, and analytical activities arising from exploratory processes within visualization systems [43]. We validated the model by generating a KG from system logs obtained during user experiments with a visualization tool for urban mobility data exploration, and by demonstrating its applicability through SPARQL queries and visualizations of query results designed to address a set of competency questions. The results demonstrate the model's potential to support researchers in understanding and comparing exploratory behaviors across visualizations. For the sake of reproducibility and reuse of the proposed model and KG by the community, the source code is publicly available at github.com/Wimmics/diva.

8.1.4 Semantic Annotation of a Corpus of Texts in Ancient Zoology as a Knowledge Graph

Participants: Catherine Faron, Franck Michel.

This work was carried out as part of the HISINUM project and related to the Zoomathia international research network which aims to study the constitution and transmission of zoological knowledge from Antiquity to the Middle Ages. The aim is to produce a corpus of texts on ancient zoology semantically annotated by a knowledge graph, respecting semantic web standards, interoperable and published on the open data web. The resulting knowledge graph allows the integration and the interrogation of relevant knowledge in order to support epistemologists, historians and philologists in their analysis of these texts and knowledge transmission through them [26]. A pipeline was developed and set up to process the corpus, extract relevant information and produce the knowledge graph (7.1.19). Additionally, a web interface allows researchers to explore the corpus via a search by concept, explore a selected work while visualizing the concepts annotating each of its parts, and visualize the results of queries implementing competency questions on a selected work from the corpus. (7.1.20).

8.1.5 Semantic Annotation of Scientific Litterature in Agriculture as a Knowledge Graph

Participants: Catherine Faron, Franck Michel.

This work was carried out as part of the D2KAB project (Data to Knowledge in Agriculture and Biodiversity), which aims to develop semantic web-based tools to describe and make agronomical data actionable and accessible following the FAIR principles. We focus on constructing domain-specific Knowledge Graphs (KGs) from textual data sources, using Natural Language Processing (NLP) techniques to extract and structure relevant entities. Our approach is based on the formalization of a semantic data model using common linked open vocabularies such as the Web Annotation Ontology (OA) and the Provenance Ontology (PROV). The model was developed by formulating motivating scenarios and competency questions from domain experts. This model has been used to construct three different KGs from three distinct corpora: PubMed scientific publications on wheat and one rice genetics and phenotyping, and French agricultural alert bulletins. The named entities to be recognized include genes, phenotypes, traits, genetic markers, taxa and phenological stages normalized using semantic resources such as the Wheat Trait and Phenotype Ontology (WTO), the French Crop Usage (FCU) thesaurus and the Plant Phenological Description Ontology (PPDO). Named entities were extracted using different NLP approaches and tools. The relevance of the semantic model was validated by implementing experts questions as SPARQL queries to be answered on the constructed RDF knowledge graphs. Our work demonstrates how domain-specific vocabularies and systematic querying of KGs can reveal hidden interactions and support agronomists in navigating vast amounts of data. The resources and transformation pipelines developed are publicly available in Git repositories [30].

8.1.6 Agile Ontology Engineering: Tooling and Methodology

Participants: Fabien Gandon, Nicolas Robert.

We contributed to the Agile and Continuous Integration for Modular Ontologies and Vocabularies (ACIMOV) [94] ontology engineering methodology for developing ontologies and vocabularies. ACIMOV extends the SAMOD agile methodology to (1) ensure alignment to selected reference ontologies; (2) plan module development based on dependencies; (3) define ontology modules that can be specialized for specific domains; (4) empower active collaboration among ontology engineers and domain experts; (5) enable application developers to select views of the ontology for their specific domain and use case. ACIMOV adopts the standard git-based approach for coding, leveraging agility and DevOps principles. It was implemented in OLIVAW [51] using the collaborative software development platforms Github tooling with continuous integration and continuous deployment workflows (CI/CD workflows) that run syntactic and semantic checks on the repository, specialize modules, generate and publish the ontology documentation. The software was also enhanced to include the possibility to use it out of the ACIMOV project architecture.

8.1.7 Corese Semantic Web Factory

Participants: Rémi Ceres, Fabien Gandon, Olivier Corby.

Corese [74], an open-source Semantic Web platform, implements W3C languages such as RDF, RDFS, OWL RL, SHACL, SPARQL, and extensions including SPARQL Function, SPARQL Transformation, and SPARQL Rule.

In the enhancement of Corese's distribution, two new interfaces, Corese-GUI and Corese-Command, were launched on Flathub. Additionally, a one-click installation script for Corese-Command is now available for Linux and MacOS.

The [documentation of Corese](#) has been fully updated.

The new interface, Corese-Command, supplements existing ones such as Corese-Library, Corese-GUI, Corese-Server, and Corese-Python. Corese-Command, evolving from the previous Corese-CLI, enables command-line usage of Corese. It encompasses subcommands for converting RDF file formats, running SPARQL queries, performing SHACL validation on RDF datasets, and executing SPARQL queries on remote endpoints. Improvements in file loading now allow handling of local files, URLs, or directories.

All interfaces have been unified to support Corese configuration files in properties format.

Enhancements include bug fixes in Corese-Python, addition of Markdown result format for SPARQL, and N-Quads RDF serialization.

Relevant websites include the Corese project site at [Corese Web site](#) and the GitHub repository at [Corese github URL](#).

8.1.8 W3C Data activity and AC Rep

Participants: Pierre-Antoine Champin, Youssef Taghzouti, Rémi Ceres, Fabien Gandon, Franck Michel, Olivier Corby.

Semantic Web technologies are based on a set of standards developed by the [World Wide Web consortium \(W3C\)](#). Participation in these standardization groups gives to researcher the opportunity to promote their results towards a broad audience, and to keep in touch with an international community of experts. Wimmics has a long history of being involved in W3C groups.

As W3C fellow, Pierre-Antoine Champin also works within the W3C team to support Semantic Web related working groups and promote the emergence of new ones, to ensure the necessary evolutions of these technologies. In 2024, the new [Linked Web Storage](#) Working Group was chartered, to standardize the Solid protocol and continued in 2025. The [Solid project](#) was started by Tim Berners-Lee, inventor of the Web, and builds on Semantic Web standards to promote the (re-)decentralization of the Web. Solid has been a research topic for Wimmics in the past years, including in the collaboration with Startin'Blox. The [RDF-star](#) Working Group is pursuing its efforts to publish the new version of RDF and SPARQL, extending them with the ability to make statements about statements. A new [Data Shapes](#) Working Group was created in December 2024 to adapt SHACL to those changes in RDF. We intend to reflect those changes into Corese (see Section 7.1.1); in fact, Corese already implements an experimental version of RDF-star.

We contribute to the the [W3C Data Shapes WG](#). The mission of this group is to update data shapes standards (SHACL) in line with the versions of core Semantic Web standards that cater for RDF-star and to extend the applications of data shapes with new packaging and use specifications. We are involved in the edition of [SHACL 1.2 Core](#) and [SHACL 1.2 Profiling](#)

We also contribute to the [Dataset Exchange WG](#). The mission of this group is (1) to maintain and revise the [Data Catalog Vocabulary \(DCAT\)](#), taking into account feature requests from the DCAT user community and (2) to define and publish guidance on the specification and use of application profiles when requesting and serving data on the Web. We are co-editors of the [Content Negotiation by Profile](#). We also participated in the Dagstuhl seminar on Metadata Models and Services Typologies in Digital Resource-Sharing Frameworks and agreed to be editor for the future Vocabulary of Variable Description (VVD) recommendation, also lead by the Dataset Exchange WG.

Finally, Fabien Gandon remains the W3C AC Rep for Inria representing institute in all standardization processes and W3C meetings (annual W3C TPAC conference and W3C AC Meeting) and he became the co-chair of the [Web and AI Interest Group at W3C \(WebAI IG\)](#).

8.2 Combined intelligent methods for heterogeneous knowledge graphs

8.2.1 Sentiment Analysis with Fuzzy Polarity Propagation

Participants: Andrea Tettamanzi.

We proposed a novel refinement of a gradual polarity propagation method that we had previously introduced, to learn the polarities of concepts and their uncertainties with respect to various domains from a labeled corpus. This year’s contribution consists in introducing a positive correction term in the polarity propagation equation to counterbalance an ubiquitous negative psychological bias in reviews. The proposed approach has been successfully evaluated using a standard benchmark, showing an improved performance relative to the state of the art, good cross-domain transfer and excellent coverage [41].

8.2.2 An Open Platform for Quality Measures in a Linked Data Index

Participants: Pierre Maillot, Olivier Corby, Catherine Faron, Fabien Gandon, Franck Michel.

There is a great diversity of RDF datasets publicly available on the web. Choosing among them requires assessing their “fitness for use” for a particular use case, and thus, finding the right quality measures and evaluating data sources according to them. However, this is not an easy task due to the large number of possible quality measures, and the multiplicity of implementation and assessment platforms. Therefore, there is a need for a common way to define measures and evaluate RDF datasets, using open standards and tools.

Developed in the context of the ANR DeKaloG, IndeGx is a SPARQL-based framework to design indexes of Knowledge Graphs declaratively [100]. We extended it to support more advanced data quality measures. We demonstrated our approach by reproducing two existing measures, showing how one can formalize and add measures using such an open declarative framework. This work was presented at the Web Conference 2024 [99]. We also reported on the use of KR models and methods to index Semantic Web Endpoints and Knowledge Graphs [59].

8.2.3 Learning Pattern-Based Extractors from Natural Language and Knowledge Graphs: Applying Large Language Models to Wikipedia and Linked Open Data

Participants: Célian Ringwald, Fabien Gandon, Catherine Faron, Franck Michel, Hanna Abi Akl.

Whether automatically extracted from structured elements of articles or manually populated, the open and linked data published in DBpedia, and Wikidata offer rich and structured complementary views of the textual descriptions found in Wikipedia. However, the unstructured text of Wikipedia articles contains a lot of information that is still missing in DBpedia and Wikidata. Extracting them would be interesting to improve the coverage and quality of these knowledge graphs (KG) and this would have an important impact on all downstream tasks.

This work proposes to exploit the dual bases formed from Wikipedia pages and Linked Open Data (LOD) bases covering the same subjects in natural language and in RDF, to produce RDF extractors targeting specific RDF patterns and tuned for a given language. Therefore, the main research question is: Can we learn efficient customized extractors targeting specific RDF patterns from the dual base formed by Wikipedia on one hand, and DBpedia and Wikidata on the other hand?

The landscape of the research field drawn at the intersection of language models and knowledge graphs is very dynamic and quickly evolving. For this reason, as the first step of this work, we designed an extended systematic review of the latest NLP approaches to KG extraction [61].

In a second step, we started the design a first dataset focused on datatype properties. We restricted the selection of our training to facts respecting a given SHACL shape and information that could be found in the Wikipedia abstract. Then, to learn how to extract relations with datatype properties from natural language, we exploited pre-trained encoder-decoder models, and more precisely BART (denoising autoencoder sequence-to-sequence model). We explored several aspects of the task formulation that could impact the generation of triples in this context: the size of the model, the size of the learning sample needed to learn a given SHACL pattern, and the syntax of the triples [110, 112].

We continued the work by questioning the impact of the syntax chosen for representing the generated output by benchmarking 12 variations of RDF syntaxes, but also by comparing two small language models (T5 and BART), and demonstrated the performances of a light Turtle syntax [111].

Finally, we designed and evaluated a method to fine-tune small language models for shape-based active relation extraction [48] and we extended this work with techniques for overcoming the generalization limits of SLM finetuning for both shape-based Extraction of datatype and object properties [49].

8.2.4 Hybridizing machine learning and knowledge graphs: injection of relation signatures

Participants: Nicolas Robert, Pierre Monnin, Catherine Faron.

Knowledge graphs (KGs) are nowadays largely adopted, representing a successful paradigm of how symbolic and transparent AI can scale on the World Wide Web. However, they are generally tackled by Machine Learning (ML) and mostly numeric based methods such as graph embedding models (KGEMs) and deep neural networks (DNNs). The latter methods have been proved efficient but lack major characteristics such as interpretability and explainability. Conversely, these characteristics are intrinsically supported by symbolic AI methods and artefacts, thus motivating a research effort to hybridize machine learning and knowledge graphs.

Towards such an hybridization, we investigated the improvement of KGEMs with symbolic knowledge for the task of link prediction which aims at predicting the missing tail of a triple $(h, r, ?)$ or the missing head of a triple $(?, r, t)$. To train such methods, batches of positive and negative triples are considered in loss functions. However, different kinds of negative triples exist: considering signatures of relations (domain and range), some negative triples may be semantically valid (e.g., `(PierreMonnin, workingPlace, Sophia Antipolis)`), while others may be semantically invalid (e.g., `(PierreMonnin, workingPlace, Philosophy)`).

From the proposed enriched datasets [52], we extended the symbolic knowledge injection method in loss functions proposed by Hubert *et al.* [96]. We experimented their method on a more realistic scenario, where only a part of the KG schema is known and perform ablative studies to assess the performance of such approaches *w.r.t.* the proportion of the schema that is known.

8.2.5 Stability of knowledge graph embedding models for link prediction

Participants: Guillaume Méroué, Pierre Monnin, Fabien Gandon.

While current evaluations of Knowledge Graph Embedding Models (KGEMs) for link prediction mainly focus on global metrics such as MRR or Hits@K, they overlook the impact of randomness at the triple level. Our experiments show that high-performing models can yield divergent predictions and embedding spaces across runs. By isolating several sources of stochasticity, we demonstrate that each contributes similarly to instability, with no clear link between performance and stability. Ensemble voting offers only limited gains, highlighting important limitations of current benchmarking practices.

8.2.6 Temporal Graph Modeling and Hybrid Learning for Health State Prediction

Participants: Hajer Akid.

Predicting the State of Health (SOH) of lithium-ion batteries and detecting abnormal degradation behaviors require data representations and learning methods capable of capturing both local variations within charge-discharge cycles and long-term degradation trends across cycles. Battery monitoring data are often handled as flat multivariate time series, which limits the exploitation of structural information emerging from the temporal organization of repeated cycling processes. To address this limitation, we developed BATT2GRAPH,

a hybrid data-driven approach that combines a temporal graph-based data representation of battery cycling data with a deep learning predictive model. Battery cycling data are modeled as a temporal property graph, enabling the structured organization of multivariate time-series signals and historical indicators. Building on this graph-based representation, we designed a hybrid architecture that jointly exploits raw sequential signals and aggregated statistical features extracted from the graph structure: a Convolutional Neural Network (CNN) captures local temporal patterns, while a Long Short-Term Memory (LSTM) models long-term dependencies across sequences. Experimental results obtained on a large-scale benchmark dataset show that the proposed approach improves prediction accuracy compared to state-of-the-art baselines, while enabling flexible monitoring and analysis supported by the graph-based data organization [32].

8.2.7 Selecting relevant pairs for analogy-based pruning of knowledge graphs

Participants: Ndeye-Emilie Mbengue, Pierre Monnin.

More and more knowledge graphs are publicly published and accessible on the Web of data, covering a widening array of domains. This allows their reuse in other downstream applications or in the construction of other knowledge graphs. However, not all represented knowledge is useful or pertaining in such cases. This is particularly the case for general large-scale knowledge graphs such as Wikidata. Additionally, the sheer size of such knowledge graphs entails scalability issues. These two aspects ask for efficient methods to extract subgraphs of interest from existing knowledge graphs. We previously developed a frugal analogy-based algorithm that, given seed entities of interest and properties to traverse, extracts their neighboring subgraphs from Wikidata, only keeping relevant neighbors while pruning irrelevant ones [98].

Analogical reasoning is based on proportions of the form "A is to B as C is to D" (e.g., "Paris is to France as Berlin is to Germany"). Recent studies in Analogy theory have shown the critical importance of selecting appropriate known (A, B) pairs to extrapolate the outcome to new and unknown (C, D) pairs to enhance the model robustness or reduce the computational cost induced by the number of required analogies generated. We studied this selection issue, formalized through the notion of Competency, on the tasks of KG pruning, selecting competent (A, B) left-pairs during the inference phase and comparing with random-based approach. The empirical study demonstrates promising results in the characterization of competency within knowledge graphs, with some methodologies succeeding in both increasing model performance and reducing the number of required analogies [42].

8.2.8 How to reason with probabilistic information on argumentation graphs?

Participants: Pierre Monnin.

The paper [37] presents fast and exact methods for computing the probability of an argument's acceptance using Dung's semantics in the Constellation paradigm of Abstract Argumentation. For (directed) Singly-Connected Graphs (SCGs), the problem can now be solved in linearithmic time instead of being exponential in the number of attacks, as reported in the literature. Moreover, in the more general case of Directed Acyclic Graphs (DAGs), we provide an algorithm whose time complexity is linearithmic in the product of the out-degree of dependent arguments, i.e., arguments reaching the argument considered for acceptance through multiple paths in the graph. We theoretically show that this complexity is lower than the lower-bound of the (exact) Constellation method, which is also supported by empirical results. We also compare our approach on DAGs with the (approximate) Monte-Carlo method, which is stopped when our approach obtains the exact results. Within this time constraint, Monte-Carlo still outputs significant errors, underlying the fast computation of our approach.

8.2.9 Investigating Language Model Capabilities to Represent and Process Formal Knowledge in Order to Assist Ontology Engineering

Participants: Hanna Abi Akl, Fabien Gandon, Catherine Faron, Pierre Monnin.

Since their introduction, Language Models (LMs) have excelled in many Natural Language Processing (NLP) tasks but still struggle with reasoning. Explicit and implicit reasoning capabilities are essential for complex tasks such as knowledge graph construction or ontology engineering. In our work [33], we assess the effects of different representations of reasoning problems on LM performance. We focus specifically on translating first-order logic (FOL) syllogisms into other formal languages on the one hand and the impact these translations have on the reasoning capabilities of Small Language Models (SLMs) on the other. The research allowed us to develop the Common Logic Grammar Construction (CLGC) pipeline to automatically translate FOL statements into formal languages like the Common Logic Interchange Format (CLIF) or Tensor Function Logic (TFL). From our experiments pairing different knowledge representations with LMs in a variety of learning settings, we showed that it is possible to substitute Natural Language (NL) with more compact and less ambiguous logical problem representations with only a minor trade-off in reasoning performance. We also introduced the Syllogistic Evaluation Framework (SEF) to classify logical problems into syllogistic categories and generate a reasoning trace for a given logical problem. Our preliminary results will allow us to apply the SEF-CLGC pipeline to reasoning tasks in ontology engineering like aiding ontology completion.

8.2.10 Hypermedia Multi-Agent Systems

Participants: Andrei Ciortea, Fabien Gandon.

Recent years have brought renewed interest in Web-based Multi-agent Systems (MAS), primarily motivated by progress in the Web of Things, Distributed Knowledge Graphs, and Generative AI. Central to these developments is the flexible, autonomous use of hypermedia—for example, to discover knowledge, invoke device functionality, or use tools. However, existing frameworks for Web-based MAS typically lack support for working with hypermedia abstractions and controls. To fill this gap, we contributed to a model and framework for hypermedia-based MAS. Our specific focus is on the environment as a first-class abstraction in MAS: agents are situated and embodied in a distributed hypermedia environment that (i) provides them with a uniform abstraction of the system and (ii) is instrumented with tools and resources they can discover and use to achieve their goals. Given a single entry point into the hypermedia environment, agents are enabled to “arrive-and-operate”: they use their prior knowledge and experience to achieve their goals by browsing hypermedia and exploiting action possibilities discovered at run time—mimicking how humans are supported by well-designed hypermedia environments today. We illustrate our approach through a demonstrator and discussed its benefits and drawbacks against an equivalent implementation without hypermedia [56].

8.3 Interaction design for decision making on and with knowledge graphs

8.3.1 Visual Exploration of Individual Mobility Data

Participants: Aline Menin, Marco Winckler, Clement Quere.

The study of spatio-temporal data, and particularly, human mobility data has always been of great importance in supporting decision-making processes affecting the everyday lives of individuals. In this context, visualization plays a crucial role in revealing tendencies and helping making sense of large and varied datasets. While the literature extensively discusses the visual exploration of mobility data concerning spatial and temporal aspects, the thematic dimension receives limited attention. The study of thematic aspects is nonetheless essential to understand how individuals’ activities and schedules shape the mobility strategies, and the repercussions of these strategies on land use (such as spatial distribution, location, and density of

various activities). Thus, we analyzed the literature on the visualization of individual mobility data, with a focus on thematic integration [28]. We analyzed 38 papers published between 2010 and 2024 in GIS and VIS venues that describe visualizations of multidimensional data related to individual movements in urban environments, concentrating on individual mobility rather than traffic data. Our findings confirmed that the thematic dimension is only partially represented in the literature, despite its critical significance in transportation. A gap that often stems from the challenge of identifying data sources that inherently provide this information, necessitating visualization designers and developers to navigate multiple, heterogeneous data sources. We identify the strengths and limitations of existing visualizations and suggest potential research directions for the field.

8.3.2 Temporal Exploration of Knowledge Graph-based Collaboration Networks

Participants: Aline Menin, Marco Winckler.

Collaboration networks are widely studied in both natural and social sciences as they reveal collaboration patterns among individuals and institutions. For example, co-authorship networks, which map collaborations between authors, are particularly useful for analyzing influential collaborators and understanding the development and dissemination of knowledge in specific research fields. However, data on collaborations are often complex, involving heterogeneous entities distributed over time, and are typically represented as multiple interconnected nodes. This complexity can lead to visual clutter, making it challenging to create effective visualizations. We investigated the use of an incremental approach to facilitate the exploration of co-authorship networks composed of multivariate entities distributed over time. Particularly, we showed how incremental visualization can assist users in focusing on relevant data while addressing scalability issues through a focus+context technique, enabling users to focus on specific data points while navigating through multiple layers of information [27]. The tool MuvIn (publicly available at dataviz.i3s.unice.fr/muvin/) implements this incremental approach to explore multi-sourced linked open data (LOD) and supports various types of collaboration networks, described via any knowledge graph. A user study involving 19 participants offers insights into how the incremental approach supports domain-specific tasks using co-authorship networks.

8.3.3 Interaction with extended reality

Participants: Aline Menin, Clement Quere, Marco Winckler.

Extended Reality environments are one of the extremes of the reality-virtuality continuum conceptualized by Milgram et al. (Milgram et al. 1995). The other extreme of the continuum is the real world, constrained by physics laws. In between, the continuum includes augmented reality (AR), where feedback from the real world is augmented with simulated cues, and augmented virtuality (AV), which augments a synthetic VR world with cues from the real world. Both AR and AV mix real and virtual cues and are thus two forms of mixed reality (MR). More recently, to encompass the full range of technologies from AR to VR, the term extended reality (XR) is often used.

Mixed, Virtual, and Augmented Reality (XR) environments are increasingly being used for learning and training purposes as they can provide immersive experience which might positively impact learning. Indeed, XR media is seeing a growing usage to tell stories and communicate ideas, offering the possibility to use embodied interactions (i.e., natural interaction gestures to interact with the system), allowing the player to actively participate in the narrative, to simulate more life-like narrative experiences. The advantages of VR storytelling have been seen in many domains, opening possibilities to create meaningful daily experiences in edutainment, rehabilitation, therapy, or professional training. In [50] we present a communication model, which aims at describing the design-usage process of a VR experience and the issues in communication and comprehension potentially affecting the proper communication of the narrative. We presented our taxonomy as well as our structured interview, which was revealed to be suitable for our user study, allowing

the identification and categorization of 127 issues among the eight defined categories. Our work provides a roadmap to identifying sources of miscommunication in VR, a first step to conceiving principles and guidelines for achieving effective communication in storytelling experiences. The data set gathered to perform that study was published as open source in [29].

All XR environments allow interaction in the 3D space where a high volume of data can be displayed. This has fostered a rapid increase in immersive analytics research, particularly in application domains that employ spatio-temporal (ST) data such as mobility data and smart cities [28]. Inspired by these approaches, we investigate the potential of VR technologies to enhance mobility data mining within a visualization technique called space-time cube (STC). More specifically, we have investigated strategies for information retrieval in the STC. While in [47] we present personalized filters that enable dynamic and interactive interrogation of data in space and time, in [46] we investigate the placement of annotations in immersive environments. This work extends to XR environments prior studies conducted on annotations.

Beyond a visual experience, immersive environments allow the creation of sound experiences. In this respect, we have investigated opportunities for real-time collaboration for musical creation and interaction [39][40]. In [40] we have explored strategies and techniques for embedding sound spatialized experiences and in the XR environment and in [39] we have investigated the use of spatialized sound for creating affordances in collaborative music creation. These works open the opportunity for the development of innovative tools for interacting in XR environments. Illustrations of such as potential have been published as demonstrations and posters, [66] [67] [68].

8.3.4 Supporting Dataset Discovery via Network Topologies and Visualization

Participants: Aline Menin.

The rapid growth of publicly available textual resources, such as lexicons and domain-specific corpora, presents challenges in efficiently identifying relevant resources. While repositories are emerging, they often lack advanced search and exploration features. Most search methods rely on keyword queries and metadata filtering, which require prior knowledge and fail to reveal connections between resources. To address this, we present DataLens, a web-based platform that combines faceted search with advanced visualization techniques to enhance resource discovery [60]. DataLens offers network-based visualizations, where the network structure can be adapted to suit the specific analysis task. It also supports a chained views approach, enabling users to explore data from multiple perspectives. A formative evaluation involving six data practitioners has confirmed the promising aspect of visualization to support dataset search, particularly, through the network visualization but also via an egocentric visualization showing detailed pairwise relationships between datasets. Participants appreciate the use of chained views to support the discovery process, despite the initial complexity of the currently employed visualization technique. Future work will build on these insights to help refine our approach to better support dataset search.

8.3.5 AI agent to convert natural language questions into SPARQL queries

Participants: Yousouf Taghzouti, Benjamin Navet, Franck Michel, Fabien Gandon.

An experimental knowledge graph (KG) driven framework (10.26434/chemrxiv-2023-sljbt) was recently introduced to facilitate the integration of heterogeneous data types, encompassing both experimental data (mass spectrometry annotation, results from biological screening and fractionation) as well as meta-data available on the Web (such as taxonomies and metabolite databases). Although this KG efficiently encapsulates the different data structures and semantic relationships, retrieving specific information through structured or visual queries or even programmatically, is not trivial.

In the collaborative projects KG-Bot and MetaboT [34, 65, 64, 63], we designed and implemented an AI agent that can convert natural language questions into SPARQL queries and programmatic data-mining tasks, and generate adapted visualization. By leveraging the potential of emerging Large Language Models

(LLMs) to understand semantic relationships encapsulated in KGs and mentioned in the questions, the agent autonomously iterates to construct a SPARQL query of any submitted natural language question. After retrieving the necessary information from the KG, the agent provides a preliminary interpretation of the results in natural language, along with relevant visualizations and statistics [117].

To consolidate this work, we proposed and obtained an SNF-ANR project called MetabolinkAI consortium, which develops a chatbot-like system using language agents for data exploration and analysis in mass spectrometry research. The project enables querying and processing metabolic data structured as knowledge graphs.

In this context, we generalized our first prototype into Gen²KGBot (generic generative knowledge graph robot) a conversational agent to access (scientific) data and knowledge in natural language [71]. We then focused on the sub problem of building questions and queries datasets for knowledge graphs [69, 70, 62] and minting competency questions and SPARQL Queries for question-answering over knowledge graphs [54]. We proposed a method, a user interface and an agent interface for the online generation of knowledge graph's competency questions and question-query training sets [53]

In parallel, we explore the development AI tools for scientific data exploration and synthesis [45]. We extend the SciLEx software through refactoring and implementing aggregation filters with configurable ranking systems. These enhancements enable optimized article selection, directly integrable with Zotero, streamlining literature review compilation by researchers. We also conduct a state-of-the-art review on memory systems for LLM-based agents in preparation for new work. The SciLEx extension and memory system research provide the technological foundation for a new generation of specialized metabolomics AI assistant.

8.4 Study of the Attention Economy, its detrimental impacts and leads for regulation

Participants: Franck Michel, Fabien Gandon.

During the last two decades, leveraging research in psychology, sociology, neuroscience and other domains, Web platforms have brought the process of capturing attention to an unprecedented scale. With the initial commonplace goal of making targeted advertizing more effective, the generalization of attention-capturing techniques and their use of cognitive biases and emotions have multiple detrimental side effects such as polarizing opinions, spreading false information and threatening public health, economies and democracies.

Aware of the problems raised and our responsibility as a community, since 2023 we have initiated a work meant to warn the computer science community and call for regulation. We brought together contributions from a wide range of disciplines (psychology, sociology, neuroscience, politics, legal domain, computer science, education etc.) to analyze current practices and consequences thereof. We published this work in 2024 at the AI, Ethics and Society conference (AAAI/AIES) [107] in an article that provides a set of propositions and principles that could be used to drive further works, and to call for actions against these practices competing to capture our attention on the Web. As a follow-up, in 2025 we published a new article at the AAA/AIES [35, 36] proposing several actions against attention capture and in particular a Pigouvian tax to regulate the attention market.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Berger-Levrault : integrating knowledge graphs and natural language processing to improve information systems

Participants: Fabien Gandon, Catherine Faron, Franck Michel, Pierre Monnin, Genesis Montenegro, Krysto Dagues de La Hellerie.

Partner: Berger-Levrault.

We have established a strong collaboration with Berger-Levrault including two CIFRE PhDs at the crossroad of knowledge graphs, information systems and natural language processing.

The first thesis is on the following topic: “From raw corpus to lexico-semantic graph: a methodological framework for an agile and efficient industrial GraphRAG”. In this thesis, Berger-Levrault and WIMMICS wish to investigate in greater detail the field of hybridization between large language models (LLMs) and knowledge graphs, in order to design and implement a GraphRAG adapted to Berger-Levrault’s business contexts. The objective is to define strategies for constructing, updating and querying the graph. We also want to select and improve inference and response generation algorithms, while establishing a robust methodology to guarantee the quality, consistency and contextualization of information. This exploration aims operational, reliable and scalable solutions, integrating the specific constraints of Berger-Levrault’s business sectors and anticipating new emerging needs in information management.

The second thesis is on the representation of intervention scenarios and decisions using a knowledge graph with an application to industrial maintenance. Industrial maintenance and its management require enormous resources. In this collaboration with Wimmics, Berger-Levrault wishes to conduct applied research on the representation of business processes in the field of CMMS (Computerized Maintenance Management Systems) using knowledge graphs. In this field, many processes still rely heavily on human experience and oral transmission: incident analysis, capitalization of diagnostics, intervention planning, documentation of decisions. Although crucial, this know-how is difficult to formalize and exploit on a large scale. The absence of a structured memory leads to redundant errors, loss of efficiency and heavy dependence on experts. Knowledge graphs offer a unique opportunity to represent, store and query these past experiences and business rules. They can provide an explicit and verifiable basis for facilitating the execution of industrial actions and capitalizing on knowledge within Berger-Levrault. The objective is to design and evaluate a knowledge graph focused on maintenance experiences. We intend to: model relevant concepts (interventions, validated diagnoses, human decisions, business rules, action scenarios), structure and encode past experiences to make them queryable, usable and adaptable in real contexts and new situations, and evaluate the relevance of the graph in terms of its usefulness for assisting interventions and optimizing maintenance processes.

DSTI Private School: neuro-symbolic and knowledge-based approaches to support reasoning capabilities in language model-based systems

Participants: Fabien Gandon, Catherine Faron, Pierre Monnin, Hanna Abi Akl.

Partner: Data ScienceTech Institute (DSTI) Private School of Engineering

DSTI School of Engineering is a Private Higher Education Institution, a specialist school in France for computer and data engineering, cyber security, data science and artificial intelligence with Bachelor’s & Master’s level programmes fully taught in English. This collaboration started with a focus on the study of neuro-symbolic and knowledge-based approaches to support reasoning capabilities in language model-based systems (LLMs). In a CIFRE thesis we collaborate to solve problems related to LLM reasoning by considering a combination of neuro-symbolic and knowledge-based methods to develop more robust and reliable reasoning capabilities in LLM-based systems. In particular, the research addressed is related to types of knowledge, learning methodologies for LLM-based systems, and model architecture improvement. This involves evaluating and proposing alternatives to natural language text formats, based on semantic web representation formats, as well as means of tracing knowledge acquisition by LLMs in order to validate the learning of key knowledge. In particular, we will focus on ontological knowledge and ontology engineering assisted by LLMs.

See [DSTI](#) Private School of Engineering.

Forgeron3 : Improving the Explainability and Evidence Reporting of Language Models (LLMs)

Participants: Pierre Monnin, Pierre Maillot, Fabien Gandon, Krysto Dagues de La Hellerie.

Partner: Forgeron3.

Language models (LLMs) such as GPT-4 and Llama 3 have revolutionized natural language processing and artificial intelligence (AI). In this context, Forgeron3 has developed Marcus, an AI capable of rapidly deploying intelligent assistants for companies that guarantee total control over internal data. However, explainability and the ability to provide evidence for their responses remain major challenges. The lack of explainability can reduce user confidence, and companies may be reluctant to adopt these technologies without guarantees of explainability and evidence reporting. To address these issues, the integration of approaches such as knowledge graphs into language models at the Retrieval-Augmented Generation (RAG) level may offer promising solutions. The objective of this collaboration is to develop methods and tools to improve the explainability of responses generated by Marcus by introducing knowledge graphs. We work on both theoretical and practical aspects, implementing concrete solutions for evidence retrieval and response explainability.

See [Forgeron3](#) company.

Probabl

Participants: Fabien Gandon, Nicolas Delaforge.

Partner: Probabl. Fabien Gandon is co-scientific advisor of the company Probabl and co-founder with Nicolas Delaforge and other persons of this company supporting open-source tools for data science. Probabl was created with a purpose, a mission to develop, maintain at the state of art, sustain, and disseminate a complete suite of open source tools for data science. Its commercial activities unfold to support the long term mission and the open sources tools in the portfolio of Probabl include the CORESE software by Wimmics.

Probabl was created as a private company in France, registered at the Trade and Companies Register under 979 093 523 (Paris RCS), with a purpose statement "to develop, maintain at the state of art, and sustain a complete suite of open source tools for data science, to benefit France, the EU and the rest of the World". The mission unfolds along 3 axes: (a) A regional mission of technological and economic sovereignty, anchored in France and Europe; (b) a global technological, commercial and industrial mission. and (c) a universal humanist mission, upholding to the highest social and societal standards from France and Europe, with a global reach.

See [Probabl](#) company.

SAP : knowledge graph and a GraphRAG for Customer Communication with Support

Participants: Fabien Gandon, Catherine Faron, Pierre Monnin.

Partner: SAP.

This collaboration with SAP is in the domain of Graph-based knowledge management and natural language access to knowledge graphs. Customer communication with Support could be very generic as SAP knowledge content is large. The answer might be found in many bases (SAP Notes, SAP Knowledge Based Article, SAP Help portal or even in SAP blog posts). Therefore, once a query is raised by the customer, the answer could be simply to give them an SAP Note/KBA number and let the customer find the solution to their issue in the note content. Even if the right article is given as an answer, customers will need to perform extraction steps or make an effort to find out which response is relevant to them. We want to leverage a recommender system to generate curated/personalized answers to user queries, enhancing user experiences across various applications, including customer support, content recommendations, and knowledge management. We intend to represent the SAP documentation system as a Knowledge Graph and to use the LLM and GraphRAG approach to generate the answer and curate it to the customer system context.

See [SAP](#) company.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Participation in other International Programs

SNF-ANR Switzerland-France project MetabolinkAI

Participants: Fabien Gandon, Catherine Faron, Pierre Monnin, Franck Michel.

Title: MetabolinkAI: transforming metabolomics data into actionable insights

Partner Institution(s): • ETH Zurich, Switzerland

- University of Geneva, Switzerland
- University of Fribourg, Switzerland
- University of ZurichETH Zurich
- SIB Swiss Institute of Bioinformatics, Switzerland
- CNRS
- INRAE
- Inria

Date/Duration: from 01/04/2025 to 31/03/2029, 48 months.

ANR Number: ANR-24-CE93-0012.

Description: MetaboLinkAI is a research project co-funded by the Swiss National Science Foundation (SNF 10002786) and the French Agence Nationale de la Recherche (ANR-24-CE93-0012) and dedicated to advancing how researchers collect, manage, and interpret metabolomics data through artificial intelligence (AI). It applies recent advances in AI to analyze complex, heterogeneous data, incorporates knowledge graphs to unify diverse metabolic and chemical information, and employs large language models for more intuitive querying, annotation, exploration, and inference of metabolomics datasets. These approaches aim to address pressing data-interpretation challenges in domains such as biomedical and environmental science.

Central to MetaboLinkAI is the Metabolomics Knowledge Hub, a resource designed to aggregate and harmonize mass spectrometry data, chemical ontologies, pathway information, and related resources. This foundation supports federated queries and semantic enrichment on a large scale.

In parallel, an AI Research Assistant is being developed to interpret user queries via natural language processing and provide automated suggestions and analyses, such as metabolite annotation and pathway inference.

The project also advances graph mining and machine learning methods to handle data gaps and uncertainties, drawing on algorithms for knowledge graph completion, uncertainty modeling, and enhanced visualization.

Finally, MetaboLinkAI will be tested through real-world applications, including biomarker discovery and chemodiversity exploration, bridging the gap between data generation and meaningful scientific insights.

The project is coordinated by ETH Zurich and CNRS. It gathers expertise in metabolomics, knowledge engineering, and artificial intelligence, with a strong commitment to open science, fostering a global community of researchers, innovators, and stakeholders.

Web site: [MetabolinkAI](#)

DFH-UFA (French-German University) Graduate School SeReCo

Participants: Catherine Faron, Aline Menin, Guillaume Meroue, Pierre Monnin, Nicolas Robert.

Title: SEmantic, REasoning and COordination

Partner Institution(s): • FAU Erlangen-Nuremberg, Germany

- Karlsruhe Institute of Technology, Germany
- St. Gallen University, Switzerland
- Mines Saint-Étienne, France
- Jean Monnet University, France

Date/Duration: 2025-2029

Description: The scientific purpose of the SeReCo graduate school is to explore the synergy between Web Science and Artificial Intelligence (AI). Its research topics include Semantic data modeling (Linked Data, ontology, annotations), Distributed data management and semantic integration, Reasoning in open and distributed environments, Multi-Agent-based modeling and programming, Deciding in open and distributed systems, Coordination models and technologies for autonomous agents, Self-organizing systems, Trust in multi-agent systems.

Web site: [SeReCo](#)

10.2 International research visitors**10.2.1 Visits of international scientists****Other international visits to the team**

Andrea Nasi

Status: PhD

Institution of origin: University of Turin

Country: Italy

Dates: from September 1 to December 5, 2025

Context of the visit: 3 months PhD student mobility at the Université Côte d’Azur to obtain the “European Doctorate” Label”. During this mobility, the PhD candidate worked in the team on the development of an Ontology to describe and visualize spatio-temporal data of cultural landscape information.

Mobility program/type of mobility: research stay

Bryan-Elliott Tam

Status: PhD

Institution of origin: University of Ghent, Belgium

Country: Italy

Dates: from September 1 to September 30, 2025

Context of the visit: To ensure the widespread adoption of this decentralized data ecosystem, it is fundamental to develop techniques that allow effective querying of highly decentralized information. Bryan-Elliott Tam’s work focuses on optimizing web queries, treating the web as one large database with a view on SOLID and traversal queries approaches.

Mobility program/type of mobility: research stay

10.2.2 Visits to international teams

Research stays abroad

Pierre Monnin

Visited institution: Università degli Studi di Bari Aldo Moro

Country: Italy

Dates: January 12 – February 1, 2025

Context of the visit: bilateral collaboration on the development of neuro-symbolic approaches for the discovery and injection of symbolic knowledge in knowledge graph embedding models

Mobility program/type of mobility: research stay

Célian Ringwald

Visited institution: King's College London

Country: England

Dates: September 28 – October 10, 2025

Context of the visit: PhD student mobility on the development of an approach to generate SHACL shapes with LLMs.

Mobility program/type of mobility: research stay

10.3 European initiatives

10.3.1 Other european programs/initiatives

SHACKLE, sub-project of ENFIELD EU project

Participants: Pierre Monnin.

Title: SHACKLE: SHape-based pAtterns for Constraining KnowLedge graph Embeddings

Partner Institution(s): Telecom Paris

Date/Duration: February – July 2025, 6 months

Description: The SHACKLE project covers regular mobility to the host institution (Telecom Paris) to work on the research objective of combining Knowledge Graph Embedding Models (KGEMs) and validation schemata, a currently unexplored area.

Project number: Sub-project (oc2-2024-TEs-02) of the ENFIELD project funded by the European Union under the grant agreement No 101120657

Web site: [SHACKLE](#)

COST Action GOBLIN

Participants: Franck Michel, Pierre Monnin, Genesis Montenegro, Célian Ringwald.

Title: Global Network on Large-Scale, Cross-domain and Multilingual Open Knowledge Graphs

Date/Duration: 2024-2028

Description: The goal of the GOBLIN action is to increase and enhance the public open knowledge available in Europe and beyond. The aim is to provide a large-scale, high quality, cross-domain and multilingual knowledge graph technology that is free to use, reuse, and redistribute by bringing people and communities interested in knowledge graphs technologies to work together on topics related to knowledge graphs engineering, knowledge graphs management and knowledge graphs utilization. Pierre Monnin is co-leader of Task 3.1.

Web site: [GOBLIN](#)

5G METRO

Participants: Marco Winckler, Clément Quéré.

Title: Merging Education, Telecommunications and Robotic Outreach

Date/Duration: 2024-2027

Description: The purpose of the project is to enhance inclusion in education through distance learning and the use of fluid, real-time video communication linked with mobile telepresence robots. Focusing on education and healthcare, the 5G METRO project will enable the development of new approaches to education and training in universities. This is an EU Project (101181316) lead by IMREED Université Côté d'Azur, Orange, Awabot. Marco Winckler is co-responsible for the development of WP 5 : case studies.

Web site: [5G Metro](#)

10.4 National initiatives**ANR D2KAB**

Participants: Olivier Corby, Catherine Faron, Fabien Gandon, Franck Michel, Nadia Yacoubi Ayadi.

Title: Data to Knowledge in Agriculture and Biodiversity

Partner Institution(s): • LIRMM

- INRAE
- IRD
- I3S
- Action

Date/Duration: from June 2019 until November 2024, 66 months

Description: The general objective of the D2KAB ANR project is to create a framework to turn agronomy and biodiversity data into knowledge –semantically described, interoperable, actionable, open– and investigate scientific methods and tools to exploit this knowledge for applications in science and agriculture. Within this project the Wimmics team is contributing to the lifting of heterogeneous dataset related to agronomy coming from the different partners of the project and is responsible to develop a unique entry point with semantic querying and navigation services providing a unified view on the lifted data.

Web site: [D2KAB](#)

ANR DeKaloG

Participants: Olivier Corby, Catherine Faron, Fabien Gandon, Pierre Maillot, Franck Michel.

Title: Decentralized Knowledge Graph

Partner Institution(s):

- Université Nantes
- INSA Lyon
- Inria Center at Université Côte d’Azur

Date/Duration: from February 2020, 42 months

Description: DeKaloG aims to: (1) propose a model to provide fair access policies to KGs without quota while ensuring complete answers to any query. Such property is crucial for enabling web automation, i.e. to allow agents or bots to interact with KGs. Preliminary results on web preemption open such perspective, but scalability issues remain; (2) propose models for capturing different levels of transparency, a method to query them efficiently, and especially, techniques to enable web automation of transparency; (3) propose a sustainable index for achieving the findability principle.

Web site: [DeKaloG](#)

ANR CROQUIS

Participants: Andrea Tettamanzi.

Title: Collecte, représentation, complétion, fusion et interrogation de données de réseaux d’eau urbains hétérogènes et incertaines

Partner Institution(s):

- CRIL
- HSM
- I3S

Date/Duration: from March 2022, 48 months + a 12 months extension

Description: The coordinator of the project is Salem Benferhat (CRIL). The local coordinator for Laboratoire I3S is Andrea Tettamanzi. The local unit involves two other members of I3S which are not part of WIMMICS, namely Célia da Costa Pereira and Claude Pasquier.

The contribution of Wimmics is focused on addressing the problem of incomplete and uncertain data.

Web site: [CROQUIS](#)

ANR AT2TA

Participants: Pierre Monnin.

Title: Analogies: from Theory to Tools and Applications

Partner Institution(s): • Université de Lorraine (LORIA)

- HInria Paris (HeKA team)
- Université Paul Sabatier (IRIT)
- IHU Imagine
- Université Côte d'Azur (I3S)
- Infologic

Date/Duration: from February 2023 until February 2026

Description: The coordinator of the AT2TA project is Miguel Couceiro (LORIA, Université de Lorraine). The local coordinator for I3S / Wimmics is Pierre Monnin.

The project AT2TA aims to develop an analogy-based machine learning framework and to demonstrate its usefulness in real case scenarios. Within the project, the Wimmics team is contributing by investigating the potential usages of analogy-based frameworks with and for knowledge graphs, and the associated adequate representation spaces.

Web site: [AT2TA](#)

ANR SAFE-KG

Participants: Fabien Gandon, Catherine Faron, Franck Michel, Pierre Monnin.

Title: Interrogation dans une fédération sécurisée de graphes de connaissances

Partner Institution(s): • Nantes Université

- Insa Lyon
- Inria Center at Université Côte d'Azur
- INSERM

Date/Duration: from November 2025 to March 2029, 42 months

Description: The project coordinator is Hala Skaf-Molli from Nantes Université.

SaFE-KG aims to address these challenges by proposing a secure federation of Knowledge Graphs, integrating trusted authentication and authorization mechanisms. Unlike traditional federations that assume public accessibility, SaFE-KG will develop solutions for secure, scalable, and efficient federations. The key objectives include:

- Unified Access and Usage Representation. In a secure federation, each knowledge graph operates under its own access and usage control policies. SaFE-KG will develop a unified model for representing access and usage across knowledge graphs, allowing for consistent data sharing across organizations while maintaining compliance with individual policies.
- Efficient Federation Engines for Secure Data. Existing federation engines are designed for querying public knowledge graphs. SaFE-KG will redesign key components, such as source selection, query decomposition, summary acquisition, provenance tracking, and access protocols, to support secure federated querying.

- **Secure Federated Query Building.** Existing query-building work targets a single knowledge graph and not a federation. SaFE-KG will develop user-friendly interfaces and tools assisting non-technical users in building queries across federations while preserving security, thus lowering barriers for domain experts and enabling seamless collaboration.

By addressing these challenges, SaFE-KG will enable secure and efficient collaboration across organizations while maintaining strict control over sensitive data. Focusing on both security and performance, SaFE-KG aims to develop solutions that protect data while delivering responsive, scalable systems for real-world applications.

ANR Number: ANR-25-CE23-7852

ECoControl, PEPR Agroécologie et Numérique

Participants: Catherine Faron, Franck Michel, Aline Menin.

Date/Duration: from March 2025 to March 2030, 60 months

Partner Institution(s): • INRAE

- Inria
- CNRS
- CIRAD
- Institut Agro AgroParisTech
- Université Côte d'Azur
- Sorbonne Université
- Université de Rennes

Title: Ecologie des Communautés et Outils Numériques pour augmenTer la RégulatiOn naturelle des insectes ravageurs en agriculture

Description: The coordinator of the Ecocontrol project is Astrid Cruaud (CBGP, Inrae). The local coordinator for the team is Catherine Faron. EcoControl aims to improve our understanding of the regulation services provided by arthropods and to identify agro-ecological levers for strengthening the natural regulation of insect pests in agriculture, at local and regional level, in mainland France, Corsica and Guadeloupe. The key objectives include:

- develop natural language processing methods to extract biological interactions and species traits from the literature;
- combine real-time sequencing and AI-assisted image recognition to characterize insect communities and their trophic interactions on a large scale from field samples;
- combine machine learning approaches to infer missing links in trophic networks, identify local parasitoids that can control an introduced insect, and identify the potential adverse effects of introducing an exogenous biocontrol agent;
- develop an ad hoc theory in community ecology to characterize the regulation process/function and decipher when and how regulation emerges from biotic interactions in arthropod networks;
- adapt artificial intelligence and statistical methods to develop a continuous spatio-temporal understanding of ecological networks and pest regulation, and identify levers that promote natural regulation at the landscape/territorial scale;
- set up a digital platform to share data, protocols and analytical workflows.

Web site: [EcoControl](#)

ISSA (AAP Collex-Persée)

Participants: Franck Michel, Anna Bobasheva, Olivier Corby, Catherine Faron, Aline Menin, Marco Winckler.

Title: Indexation Sémantique d'une archive scientifique et Services Associés pour la science ouverte (2)

Partner Institution(s): • CIRAD

- Mines d'Alès
- Inria

Date/Duration: from October 2020 to January 2025

Description: The ISSA project is led by the CIRAD and aims to set up a framework for the semantic indexing of scientific publications with thematic and geographic keywords from terminological resources. It also intends to demonstrate the interest of this approach by developing innovative search and visualization services capable of exploiting this semantic index. Agritrop, Cirad's open publications archive, serves as a use case and proof of concept throughout the project. In this context, the primarily semantic resources are the Agrovoc thesaurus, Wikidata and GeoNames.

Wimmics team is responsible for (1) the generation and publication of the knowledge graph representing the indexed entities, and (2) the development of search/visualization tools intended for researchers and/or information.

Web site: [ISSA](#)

10.5 Regional initiatives**Chaire 3IA / Cluster IA Université Côte d'Azur - Fabien Gandon**

Participants: Fabien Gandon, Guillaume Méroué, Yousouf Taghzouti.

Description: Created in 2019, the 3IA Côte d'Azur was officially labeled an IA-cluster following its application to the call for expressions of interest (AMI) "IA-cluster: world-class research and training clusters in artificial intelligence" launched as part of the national Artificial Intelligence strategy launched in 2018. Fabien Gandon is Holder of a 3IA Chair at the Interdisciplinary Institute of Artificial Intelligence of University Côte d'Azur on the topic "Knowledge augmentation for human and software agent on the Web". This chair was granted in 2019 and renewed in 2024. It focuses on symbolic, non-symbolic, and hybrid AI methods and models to augment knowledge content, quality, exchange and processing for human and software agent in the context of distributed systems with a specific focus on the Web-based approaches.

Web site: [Chaire 3IA / Cluster IA Université Côte d'Azur - Fabien Gandon](#)

IDEX and DS4H project KG-bot

Participants: Yousouf Taghzouti, Fabien Gandon, Franck Michel.

Date/Duration: from November 2024 to November 2025

Title: Knowledge Graph chatBot: Toward Large Language Model based Interaction with Metabolomics Knowledge Graphs

Description: The KG-bot project is funded by the Academy of Excellence "Networks, Information and Digital Society" and DS4H. It aims to enhance an AI-powered chemistry chatbot prototype designed to improve the accessibility and usability of metabolomics knowledge graphs (KGs). By leveraging mass spectrometry data, the chatbot employs a natural language interface to generate queries (using the SPARQL language), allowing chemists to intuitively explore complex metabolomics the knowledge graph (represented in RDF). Key objectives include broadening the chatbot's compatibility with various large language models (LLMs) and KGs, integrating dynamic tools for data extraction and visualization, and enabling extended dialogical interactions to support iterative queries. The project also seeks to enrich user interactions by providing features such as result visualization, hypothesis generation, and analysis recommendations. Building on the interdisciplinary expertise of the project partners, this initiative fosters transdisciplinary collaboration and aims to deliver scalable solutions applicable across multiple domains. Anticipated outcomes include enhanced access to scientific data and the development of a robust open-source framework to support future academic and industrial applications. Funding will be directed towards supporting postdoctoral researchers and student contributions to the project.

Web site: [KG-bot](#)

HISINUM

Participants: Catherine Faron, Franck Michel.

Date/Duration: from November 2023 to November 2026

Title: Histoire des savoirs et des Idées et pratiques du NUMérique

Partner Institution(s): • GREDEG

- CRHI
- CEPAM
- I3S

Description: HISINUM is a 3-years project funded by the Académie d'Excellence "Homme, Idées et Milieux" of Université Côte d'Aur and led by Muriel Dal Pont Legrand. The aim of this consortium project is to reflect on how digital humanities are renewing research practices and the issue of data in the humanities and social sciences, and on the epistemological impact of the new tools and their ability to change disciplinary boundaries. The project is structured into three programmes and Catherine Faron is co-leading the one on the history of science in ancient and medieval zoology, aiming to develop new methods and new generic intelligent resource analysis services for the valorization and analysis of a corpus of ancient scientific texts.

Web site: [HISINUM](#)

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organization

Member of organizing committees

- Catherine Faron: member of the steering board committee of the [workshop SemDH 2025 at ESWC 2025](#)
- Fabien Gandon: co-organizer of the [workshop HyperAgents 2025 at ECAI 2025](#) and co-editor of the proceedings [57].

11.1.2 Scientific events: selection

Chair of conference program committees

- Fabien Gandon: co-chair the **History of the Web** track at the Web Conference in Sydney April 28 – May 2 2025 (TheWebConf/WWW 2025).
- Marco Winckler: Subcommittee co-chair of the full papers track “Understanding Users” at the IFIP TC13 INTERACT Conference (INTERACT’2025). September 8-12, 2025, Belo Horizonte, Brazil; Doctoral Consortium co-chair at the IFIP TC13 INTERACT Conference (INTERACT’2025). September 8-12, 2025, Belo Horizonte, Brazil.

Member of conference program committees

- Catherine Faron: International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2026; The ACM Web Conference (TheWebConf/WWW) 2026; European Conference on Artificial Intelligence (ECAI) 2025; European Semantic Web Conference (ESWC) 2026, research track PC, resource track senior PC, LLMs for KE special track PC; International Semantic Web Conference (ISWC) 2025, research track, doctoral consortium PC; International conference on Semantic Systems (SemanticS) 2025; International Conference on Agents and Artificial Intelligence (ICAART) 2026; International Workshop on Hypermedia Multi-Agent Systems (HyperAgents 2025); International Workshop on Semantic Digital Humanities (SemDH 2025); International Workshop on Semantic Web and Ontology Design for Cultural Heritage (SWODCH 2025); International Workshop on Natural Scientific Language Processing (NSLP 2025); Journées francophones d’Ingénierie des Connaissances (PFIA-IC) 2025; French conference Extraction et Gestion des Connaissances (EGC) 2026.
- Fabien Gandon: International Semantic Web Conference (ISWC 2025); The AAAI/ACM Conference on AI, Ethics, and Society (AIES 2025); and workshops: GenAIK 2025, SemDev 2025, XAI-KRKG.
- Aline Menin: IEEE Conference on Virtual Reality and 3D User Interfaces (IEEE VR) 2026, Conférence Internationale Francophone sur l’Interaction Homme-Machine (IHM) 2025
- Pierre Monnin: ISWC 2025, ECAI 2025, ECML-PKDD 2025, ESWC 2025, SAC 2025, TheWebConf 2025, MedInfo 2025, and workshops XAI-KG, XAI-KRKG.
- Andrea Tettamanzi: Evo* 2025, GECCO 2025, ESWC 2025, UAI 2025, ECAI 2025, AAMAS 2026, ICAART 2026.
- Marco Winckler: ACM EICS 2025; ICWE 2025, IFIP IOT 2025, IHC 2025, INTERACT 2025.

Reviewer

- Marco Winckler: IEEE VR 2025, IFIP ICEC 2025, WISE 2025, ACM IUI 2025, IHM 2025, IS-EUD 2025, CoPDA workshop 2025, ACM CHI 2025, S-BPM 2025.

11.1.3 Journal

Member of editorial boards

- Catherine Faron: Engineering Applications of Artificial Intelligence, Transactions on Graph Data and Knowledge, Revue Ouverte d’Intelligence Artificielle.
- Pierre Monnin: Transactions on Graph Data and Knowledge
- Marco Winckler: Journal of Web Engineering (River Publishers), Interacting with Computers (Oxford Press), Behaviour and Information Technology, PACM Proceedings on Human-Computer Interaction (ACM Sheridam), IFIP « Advances in Information and Communication Technology » (Springer).

Reviewer - reviewing activities

- Catherine Faron: Semantic Web Journal; Artificial Intelligence in Medicine; Transactions on Graph Data and Knowledge; International Journal of Information Management Data Insights.
- Pierre Monnin: Scientific Reports, International Journal of Approximate Reasoning, Neurocomputing, Semantic Web Journal, Transactions on Graph Data and Knowledge, Engineering Applications of Artificial Intelligence, Pattern Recognition Letters, ACM Computing Surveys
- Aline Menin: Computer Animation and Virtual Worlds Journal, ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS), IFIP TC13 International Conference on Human-Computer Interaction (INTERACT), Interacting with Computers Journal

11.1.4 Invited talks

- Fabien Gandon:
 - “Graphes de Connaissances: logiciels, modèles et standards pour l’échange et le traitement de bases de graphes” at P16 Aniversary, Cyber campus, Paris.
 - “Knowledge Graphs as the Foundation for Interoperable Intelligent Systems” at Yearly seminar of [EcoLab \(laboratoire de l’innovation au service de la transition écologique au sein du ministère de la transition écologique\)](#), Inria, Rocquencourt June 19, 2025.
 - “LLM and Linked Data” at W3C AC yearly meeting.
- Pierre Monnin:
 - “Aligning complex units in knowledge graphs - Symbolic and neuro-symbolic approaches for pharmacogenomics”, SIMDAC 2025, MOBE, Université d’Orléans, June 24, 2025, Orléans, France.
 - "Neuro-symbolic approaches for the knowledge graph lifecycle", DIG team seminar, LTCI, Télécom Paris, March 18, 2025, Paris, France.
 - "The Schema Strikes Back: Refining Knowledge Graphs with Neuro-Symbolic AI", Università degli Studi di Bari Aldo Moro - Seminar of the PhD Program in Computer Science and Mathematics, January 27, 2025, Bari, Italy.
- Marco Winckler:
 - "Interactive Data Visualization: challenges and opportunities for covering the analytical provenance gap", 9th MOMI “Le Monde des Mathématiques Industrielles” / The World of Industrial Mathematics. May 26, 2025, Inria Sophia Antipolis, France.
 - Lacunes dans la compréhension de la conception de récits visuels immersifs. Keynote at the XR2C2 journée d’accélération, July 11, 2025, Paris, France.

11.1.5 Scientific expertise

- Catherine Faron: member of the scientific committee of the [Academy of Excellence 5](#) “People, Ideas and Environments” of the Excellence Initiative of Université Côte d’Azur ; member of the scientific committee of the national research infrastructure [CollEx-Persée](#); member of a review panel for the DFG [National Research Data Infrastructure \(NFDI\)](#); evaluation of projects submitted to the Horizon Infra 2025 call of the European Commission; evaluation of projects submitted to the Horizon-MSCA PF 2025 call of the European Commission.
- Fabien Gandon: evaluation of FNR CORE projects (CORE Multi-Annual Thematic Research Programme) submitted to Fonds National de la Recherche Luxembourg (FNR).
- Pierre Monnin: evaluation of projects submitted to the [2nd Innovation Scheme Open Call of the ENFIELD EU project](#).

- Marco Winckler: steering Committee chair for the IFIP INTERACT and IFIP TC13 vice-chair for conferences; evaluation of projects submitted to the ERC Advanced Grant 2024 Call for the European Research Council.
- Aline Menin: evaluation of project submitted to [ANR AAPG 2025](#).

11.1.6 Research administration

- Catherine Faron: Leader of the Wimmics team.
- Fabien Gandon: Co-president of scientific and pedagogical council of the Data Science Technical Institute ([DSTI](#)); [W3C](#) Advisory Committee Representative (AC Rep) for Inria; Co-scientific advisor of [Mission P16](#) for sovereign digital commons for data science; Inria Representative in the [Web Science Trust](#).
- Pierre Monnin: spokesperson of AfIA, the French Artificial Intelligence Society.
- Marco Winckler: Leader of the SPARKS division of the CNRS laboratory I3S (UMR 7271).

11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

11.2.1 Teaching

Participants: Hanna Abi Akl, Hajer Akid, Molka Dhouib, Catherine Faron, Fabien Gandon, Aline Menin, Guillaume Méroué, Franck Michel, Pierre Monnin, Andrea Tettamanzi, Célian Ringwald, Nicolas Robert, Marco Winckler.

- Hajer Akid:
 - Licence 2 Computer Science, Univ. Côte d’Azur: Data bases, 24h
- Molka Dhouib:
 - Master DSAI Univ. Côte d’Azur: Web of Data, 30h
- Catherine Faron:
 - Master 2/5A SI: Web of Data, 32 h, PNS, Univ. Côte d’Azur, France
 - Master 2/5A SI: Semantic Web, 32h, PNS, Univ. Côte d’Azur, France
 - Master Science, Data pipeline, 50h, DSTI, France
 - Master 2/5A SI: Coordination of the track on Artificial Intelligence and Knowledge Engineering, PNS, Univ. Côte d’Azur, France
- Fabien Gandon:
 - Master Science: Integrating Semantic Web technologies in Data Science developments, 72 h, M2, [Data ScienceTech Institute \(DSTI\)](#), France
 - Tutorial Inria Academy, “[CORESE and Semantic Web Standards](#)”, 6h x 9 sessions over the year
- Aline Menin:
 - Master 1, Methods and Tools for Technical and Scientific Writing, 18h ETD, Master DSAI, EUR DS4H, Univ. Côte d’Azur, France
 - Master 2, Data Visualization, 18h ETD, Master MBDS, EUR DS4H, Univ. Côte d’Azur, France
 - Master 2/5A, Data visualization, 13h ETD, PNS, Univ. Côte d’Azur, France
 - BUT 2, “Développement efficace”, “Qualité de développement”, 5h ETD, IUT, Univ. Côte d’Azur, France

- BUT 1, “Développement des Applications avec IHM”, 36h ETD, IUT, Univ. Côte d’Azur, France
- Master 2, Data Visualization, 18h ETD, Master Media et Humanité Numérique (MHN), EUR CREATES, Univ. Côte d’Azur, France
- Master 2, Data Processing, 18h ETD, Master Media et Humanité Numérique (MHN), EUR CREATES, Univ. Côte d’Azur, France
- Guillaume Mérroué:
 - Master 1/4A SI: NoSQL databases, 15h TD, PNS, Univ. Côte d’Azur, France
 - Licence 3/3A SI: Base de données relationnelles, 35 h TD, PNS, Univ. Côte d’Azur, France
 - Licence 3/3A SI: Initiation à la recherche scientifique, 10 h TD, PNS, Univ. Côte d’Azur, France
- Franck Michel:
 - Master 2/5A SI: Web of Data, 16 h, PNS, Univ. Côte d’Azur, France
- Pierre Monnin:
 - **Unite! Research School 2025**: Analogical reasoning through the lense of Knowledge Graphs, 1h, remote.
 - Master 2/5A SI: Machine Learning & Semantic Web, 2h CM, 3h TD, PNS, Univ. Côte d’Azur, France
 - Master 2 Applied Foreign Languages: Artificial Intelligence, professional applications, 5h CM, 5h TD, Univ. Côte d’Azur, France
 - Master 1 Adult Education: Introduction to Artificial Intelligence, 24h CM
 - Master 1/4A SI: NoSQL databases, 6h CM, 15h TD, PNS, Univ. Côte d’Azur, France
- Célian Ringwald:
 - Master 1/4A SI: NoSQL databases, 15h TD, PNS, Univ. Côte d’Azur, France
 - Licence 3/3 SI: ECUE Environnement informatique 12 h (TD), PNS, Univ. Côte d’Azur, France
- Nicolas Robert:
 - Master 1/4A SI: NoSQL databases, 15h TD, PNS, Univ. Côte d’Azur, France
 - Master 1/4A SI: Base de données relationnelles, 22 h (TD), PNS, Univ. Côte d’Azur, France
 - Licence 3/3A SI: Base de données relationnelles, 35 h TD, 1h course, PNS, Univ. Côte d’Azur, France
- Andrea Tettamanzi
 - Licence: Introduction à l’Intelligence Artificielle, 45 h ETD, L2, Univ. Côte d’Azur, France
 - Master: Logic for AI, 30 h ETD, M1, Univ. Côte d’Azur, France
 - Master: Web, 30 h ETD, M1, Univ. Côte d’Azur, France
 - Master: Algorithmes Évolutionnaires, 24.5 h ETD, M2, Univ. Côte d’Azur, France
 - Master: Modélisation del l’Incertitude, 24.5 h ETD, M2, Univ. Côte d’Azur, France
- Marco Winckler
 - Licence 3: Event-driven programming, 45 h ETD, PNS, Univ. Côte d’Azur, France
 - Master 2: Accessibility of Interactive Systems, 15 h ETD, PNS, Univ. Côte d’Azur, France
 - Master 2: Introduction to Scientific Research, 15 h ETD, PNS, Univ. Côte d’Azur, France
 - Master 2: Information Visualisation, 34 h ETD, PNS, Univ. Côte d’Azur, France
 - Master 2: Data Visualization, 15 h ETD, MBDS DS4H, Univ. Côte d’Azur, France.

- Master 2: Design of Interactive Systems, 34 ETD, PNS, Univ. Côte d’Azur, France
- Master 2: Evaluation of Interactive Systems, 34 ETD, PNS, Univ. Côte d’Azur, France
- Master 2: Multimodal Interaction Techniques, 15 ETD, PNS, Univ. Côte d’Azur, France
- Master 2: Coordination of the TER (Travaux de Fin d’Etude), PNS, Univ. Côte d’Azur, France
- Master 2: Coordination of the track on Human-Computer Interaction at the Informatics Department, PNS, Univ. Côte d’Azur, France

E-learning

- Mooc: Fabien Gandon, Olivier Corby & Catherine Faron, Web of Data and Semantic Web (FR), 7 weeks, [FUN](#), Inria, France Université Numérique, self-paced course 41002, Education for Adults, 24402 learners registered at the time of this report and 855 certificates/badges, [MOOC page](#).
- Mooc: Fabien Gandon, Olivier Corby & Catherine Faron, Introduction to a Web of Linked Data (EN), 4 weeks, [FUN](#), Inria, France Université Numérique, self-paced course 41013, Education for Adults, 5885 learners registered at the time of this report, [MOOC page](#).
- Mooc: Fabien Gandon, Olivier Corby & Catherine Faron, Web of Data (EN), 4 weeks, [Coursera](#), self-paced course Education for Adults, 7059 total learners at the time of this report, [MOOC page](#).

11.2.2 Supervision

PhD thesis

- Catherine Faron: co-supervision of the PhD thesis of Célian Ringwald, Nicolas Robert, Genesis Montenegro, Hanna Abi Akl.
- Fabien Gandon: co-supervision of the PhD thesis of Célian Ringwald, Guillaume Meroue, Ndeye-Emilie Mgengue, Genesis Montenegro, Hanna Abi Akl, Matthieu Feraud.
- Aline Menin: co-supervision of the PhD thesis of Clément Quere.
- Franck Michel: co-supervision of the PhD thesis of Célian Ringwald, Genesis Montenegro.
- Pierre Monnin: co-supervision of the PhD thesis of Nicolas Robert, Guillaume Meroue, Ndeye-Emilie Mgengue, Genesis Montenegro, Hanna Abi Akl.
- Marco Winckler: co-supervision of the PhD thesis of Clément Quere.

Internships

- Aline Menin: co-supervision of internships of Minh Huy Do (Master 1), Mohamed Fountir (Licence 2), Matteo Lachen (BUT2), Jeremy Moncada (Licence 3), Sajal Paudyal (Master 2).

11.2.3 Juries

- Catherine Faron:
 - Member of the PhD thesis defense jury for Thibaut Soulard, “When Facts Expire: Hybrid Approaches for Temporal Validation of Facts in Multiple and Heterogeneous Knowledge Graphs”, LISN, Université Paris-Saclay, France.
 - Member of the PhD thesis defense jury for Guilherme Santos Sousa, “Discovering expressive relations between knowledge graphs”, IRIT, Université de Toulouse, France.
 - Member of the jury COS PR 27, IRIT, Université de Toulouse, France.
 - Member of the jury COP (repyramidage) 27, IRIT, Université de Toulouse, France.
 - Member of the individual monitoring committee (CSI) of the PhD students Manon Ovide at Université de Tours and Ekaterina Sviridova at Université Côte d’Azur.

- Fabien Gandon:
 - Member of the jury COS PU CNAM, PR 27 « Web sémantique et Business Intelligence », Conservatoire National des Arts et Métier, CEDRIC Lab
 - Reviewer PhD William Charles « Des interprétations de documents et données à la connaissance historique : une approche ontologique appliquée au cas des territoires », IRIT, Univ. Toulouse
 - Member of CSI: Fanfu Wei (Eurecom), Pierre Epron (Univ. Paris Cite'), Greta Damo (Univ. Côte d'Azur), Cyprien Michel-Deletie (Univ. Côte d'Azur)
- Aline Menin
 - Member of the recruiting committee for Associated Professor (MCF 50), I3S, Univ. Côte d'Azur, France.
 - Member of the individual monitoring committee (CSI) of the PhD students Theo Szanto and Gregoire Picard.
 - Member of the recruiting committee for Temporary Assistant Professor and Researcher (ATER), Univ. Côte d'Azur, France.
- Pierre Monnin:
 - Reviewer of the PhD thesis of Raffaele Scaringi, "Empowering Deep Learning Models Through Contextual Knowledge", Università degli Studi di Bari Aldo Moro, Bari, Italy.
- Andrea Tettamanzi:
 - Member of the recruiting committee for Assistant Professor en CDD (RTT, Bando 2024-RTT-079) at Università Milano Bicocca, Italie.
 - Member of the PhD Committee of Alicia Bianchi at Université Côte d'Azur, Laboratoire ESPACE (géographie).
- Marco Winckler:
 - Member of the recruiting committee for Associated Professor (MCF 250801) at CNAM, France.
 - Member of the recruiting committee for Associated Professor (MCF 250643) at Université Paul Sabatier, France.
 - President of the recruiting committee for Associated Professor (MCF 250860) at Department MMI IUT CASTRES, Université Paul Sabatier, France.
 - Reviewer for the PhD of Maria Paula Corrêa Angelon. "Quality in use evaluation of smart environment applications by agent-based simulation". Université Polytechnique Hauts-de-France and INSA Hauts-de-France, Valenciennes, France.
 - Member of the jury of PhD Maylon Pires Macedo. "Towards a UX Data-Centric Approach: Providing Tools to Support Software Development". Federal University of São Carlos, UFSCar, Brazil.
 - Member of the individual monitoring committee (CSI First year) of the PhD students: Deborah Doré and Nicolas Robert, Université Côte d'Azur, France.
 - Member of the Research Master thesis jury of Fabio Neukirchen, Universidade Federal do Rio Grande do Sul, UFRGS, Porto Alegre, Brazil.

11.3 Popularization

11.3.1 Productions (articles, videos, podcasts, serious games, ...)

- Documentary and panel on TVMonaco: "IA : QUI SONT-ELLES ?", Fabien Gandon
- Blog Binaire, Le Monde, « Prêtez Attention : quand « prêter » est « données » » [post 1](#) and [post 2](#) Fabien Gandon, Franck Michel
- Podcast : [PAY ATTENTION: A Call to Regulate the Attention Market and Prevent Algorithmic Emotional Governance](#), Fabien Gandon, Franck Michel, Creative Process

11.3.2 Participation in Live events

- Fabien Gandon
 - Science Day "Fête de la Science" in Lyon : talk and panel October 9, 2025, Scietific Library, BU Sciences on the topic : "Quels sont les impacts de l'IA sur nos vies?"
 - Session Chiche, May 5, 2025 CIV Lycée international de Valbonne
 - Interviews with journalist students from Mastère à l'École du Journalisme de Nice (EDJ)
- Pierre Monnin, Ndeye-Emilie Mbengue
 - Presentation "Knowledge Graphs & Analogical Reasoning". Visit of students from EURECOM to the Centre Inria d'Université Côte d'Azur. November 5, 2025. Sophia-Antipolis, France
 - Presentation "Knowledge Graphs & Analogical Reasoning". Visit of students from ENS Lyon to the I3S Laboratory. November 5, 2025. Sophia-Antipolis, France

12 Scientific production

12.1 Major publications

- [1] D. Allemang, J. Hendler and F. Gandon. *Semantic Web for the Working Ontologist*. 3. ACM, June 2020. DOI: [10.1145/3382097](https://doi.org/10.1145/3382097). URL: <https://hal.inria.fr/hal-02939606>.
- [2] A. Ballout. 'Active learning for axiom discovery'. Université Côte d'Azur, 24th June 2024. URL: <https://theses.hal.science/tel-04680907>.
- [3] A. Ben othmane. 'CARS - A multi-agent framework to support the decision making in uncertain spatio-temporal real-world applications'. Université Côte d'Azur, 12th Oct. 2017. URL: <https://hal.archives-ouvertes.fr/tel-01680984>.
- [4] K. R. Bouzidi. 'Semantic web models to support the creation of technical regulatory documents in building industry'. Université Nice Sophia Antipolis, 11th Sept. 2013. URL: <https://tel.archives-ouvertes.fr/tel-00876366>.
- [5] L. Cadorel. 'Qualifying and quantifying uncertainty of geolocation information extracted from french real estated ads'. Université Côte d'Azur, 24th Jan. 2024. URL: <https://theses.hal.science/tel-04524701>.
- [6] L. Costabello. 'Context-aware access control and presentation of linked data'. Université Nice Sophia Antipolis, 29th Nov. 2013. URL: <https://tel.archives-ouvertes.fr/tel-00934617>.
- [7] P. F. Diallo. 'Sociocultural and temporal aspects in ontologies dedicated to virtual communities'. COMUE Université Côte d'Azur (2015 - 2019); Université de Saint-Louis (Sénégal), 16th Sept. 2016. URL: <https://tel.archives-ouvertes.fr/tel-01402394>.
- [8] A. E. A. Djebri. 'Uncertainty Management for Linked Data Reliability on the Semantic Web'. Université Côte D'Azur, 24th Feb. 2022. URL: <https://hal.archives-ouvertes.fr/tel-03679118>.
- [9] A. Ettorre. 'Towards an interpretable model of learners in a learning environment based on knowledge graphs'. Université Côte d'Azur, 28th Nov. 2022. URL: <https://theses.hal.science/tel-04060324>.
- [10] R. Felin. 'Evolutionary knowledge discovery from RDF data graphs'. Université Côte d'Azur, 22nd Nov. 2024. URL: <https://theses.hal.science/tel-04874737>.
- [11] F. Gandon. 'Distributed Artificial Intelligence And Knowledge Management: Ontologies And Multi-Agent Systems For A Corporate Semantic Web'. Université Nice Sophia Antipolis, 7th Nov. 2002. URL: <https://tel.archives-ouvertes.fr/tel-00378201>.
- [12] R. Gazzotti. 'Knowledge graphs based extension of patients' files to predict hospitalization'. Theses. Université Côte d'Azur, Apr. 2020. URL: <https://hal.archives-ouvertes.fr/tel-02959420>.

- [13] N. Halliwell. ‘Evaluating and improving explanation quality of graph neural network link prediction on knowledge graphs’. Université Côte d’Azur, 23rd Nov. 2022. URL: <https://theses.hal.science/tel-03907696>.
- [14] R. Hasan. ‘Predicting query performance and explaining results to assist Linked Data consumption’. Université Nice Sophia Antipolis, 4th Nov. 2014. URL: <https://tel.archives-ouvertes.fr/tel-01127124>.
- [15] M. Lefrançois. ‘Meaning-Text Theory lexical semantic knowledge representation : conceptualization, representation, and operationalization of lexicographic definitions’. Université Nice Sophia Antipolis, 24th June 2014. URL: <https://tel.archives-ouvertes.fr/tel-01071945>.
- [16] A. Macina. ‘SPARQL distributed query processing over linked data’. COMUE Université Côte d’Azur (2015 - 2019), 17th Dec. 2018. URL: <https://tel.archives-ouvertes.fr/tel-02340700>.
- [17] N. Marie. ‘Linked data based exploratory search’. Université Nice Sophia Antipolis, 12th Dec. 2014. URL: <https://tel.archives-ouvertes.fr/tel-01130622>.
- [18] Z. Meng. ‘Temporal and semantic analysis of richly typed social networks from user-generated content sites on the web’. Université Côte d’Azur, 7th Nov. 2016. URL: <https://hal.inria.fr/tel-01402612>.
- [19] F. Michel. ‘Integrating heterogeneous data sources in the Web of data’. Université Côte d’Azur, 3rd Mar. 2017. URL: <https://tel.archives-ouvertes.fr/tel-01508602>.
- [20] F. Michel, F. Gandon, V. Ah-Kane, A. Bobasheva, E. Cabrio, O. Corby, R. Gazzotti, A. Giboin, S. Marro, T. Mayer, M. Simon, S. Villata and M. Winckler. ‘Covid-on-the-Web: Knowledge Graph and Services to Advance COVID-19 Research’. In: ISWC 2020 - 19th International Semantic Web Conference. Athens / Virtual, Greece, 1st Nov. 2020. DOI: [10.1007/978-3-030-62466-8_19](https://doi.org/10.1007/978-3-030-62466-8_19). URL: <https://hal.archives-ouvertes.fr/hal-02939363>.
- [21] T. H. Nguyen. ‘Mining the semantic Web for OWL axioms’. Université Côte d’Azur, 2nd July 2021. URL: <https://tel.archives-ouvertes.fr/tel-03406784>.
- [22] C. Pasquier, C. Da Costa Pereira and A. G. B. Tettamanzi. ‘Extending a Fuzzy Polarity Propagation Method for Multi-Domain Sentiment Analysis with Word Embedding and POS Tagging’. In: *Frontiers in Artificial Intelligence and Applications*. ECAI 2020 - 24th European Conference on Artificial Intelligence. Vol. 325. Santiago de Compostela, Spain: IOS Press, 29th Aug. 2020, pp. 2140–2147. DOI: [10.3233/FAIA200338](https://doi.org/10.3233/FAIA200338). URL: <https://hal.inria.fr/hal-02936130>.
- [23] O. Seye. ‘Sharing and reusing rules for the Web of data’. Université Nice Sophia Antipolis; Université Gaston Berger de Saint Louis, 15th Dec. 2014. URL: <https://hal.inria.fr/tel-01096306>.
- [24] M. Tounsi Dhoub. ‘Knowledge engineering in the sourcing domain for the recommendation of providers’. Université Côte d’Azur, 26th Mar. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03336353>.
- [25] D. M. Tran. ‘Discovering multi-relational association rules from ontological knowledge bases to enrich ontologies’. Université Côte d’Azur; Université de Danang (Vietnam), 23rd July 2018. URL: <https://tel.archives-ouvertes.fr/tel-01926812>.

12.2 Publications of the year

International journals

- [26] M. Dhoub, A. Barbe, C. Faron, A. Zucker and M. Corneli. ‘Grappe de connaissance outillée au service des chercheurs en histoire de la zoologie antique et médiévale’. In: *Revue Ouverte d’Intelligence Artificielle* 6.1-2 (2025), pp. 85–106. DOI: [10.5802/roia.94](https://doi.org/10.5802/roia.94). URL: <https://hal.science/hal-05078267> (cit. on p. 25).
- [27] A. Menin, P. Buono and M. Winckler. ‘Muvin: A Visual Analytics Tool for Exploring Dynamic Collaboration Networks from Knowledge Graphs’. In: *International Journal of Human-Computer Interaction* (17th Nov. 2025). DOI: [10.1080/10447318.2025.2581255](https://doi.org/10.1080/10447318.2025.2581255). URL: <https://hal.science/hal-05302463> (cit. on p. 32).

- [28] A. Menin, C. Quere, J. Wagner, S. Chardonnel, P.-A. Davoine, W. Stuerzlinger, C. M. Dal Sasso Freitas, L. Nedel and M. Winckler. ‘What about thematic information? An analysis of the multidimensional visualization of individual mobility’. In: *Visual Informatics* 9.1 (Mar. 2025), pp. 99–115. DOI: [10.1016/j.visinf.2025.02.002](https://doi.org/10.1016/j.visinf.2025.02.002). URL: <https://hal.science/hal-04944676> (cit. on pp. 32, 33).
- [29] H.-Y. Wu, F. A. S. Robert, F. F. Gallo, K. Pirkovets, C. Quere, J. Delachambre, S. Ramanoël, A. Gros, M. Winckler, L. Sassatelli, M. Hayotte, A. Menin and P. Kornprobst. ‘Exploring, walking, and interacting in virtual reality with simulated low vision: a living contextual dataset’. In: *Scientific Data* 12.330 (25th Feb. 2025), p. 14. URL: <https://inria.hal.science/hal-04429351> (cit. on p. 33).
- [30] N. Yacoubi Ayadi, C. Faron, F. Michel, R. Bossy and A. Barbe. ‘Construction et exploitation d’un graphe de connaissances sur la littérature scientifique en sciences de la vie’. In: *Revue Ouverte d’Intelligence Artificielle* 6.1-2 (21st Nov. 2025), pp. 107–129. DOI: [10.5802/roia.95](https://doi.org/10.5802/roia.95). URL: <https://hal.science/hal-05391973> (cit. on p. 26).

International peer-reviewed conferences

- [31] M. Abd Nikooie Pour, E. Blomqvist, P. Giesteira Cotovio, A. Coulet, L. Ferraz, S. Hertling, S. Jain, E. Jiménez-Ruiz, F. Kraus, P. Lambrix, H. Li, Y. Li, X. Liu, P. Monnin, H. Paulheim, C. Pesquita, A. Sharma, P. Shvaiko, M. Silva, G. Sousa, C. Trojahn, J. Vataščinová, B. Yaman, O. Zamazal and L. Zhou. ‘Results of the Ontology Alignment Evaluation Initiative 2025’. In: OM 2025 - Ontology Matching 2025. Nara, Japan, 2nd Nov. 2025. URL: <https://inria.hal.science/hal-05447839> (cit. on p. 24).
- [32] H. Akid, M. W. Mabrouk, S. Arbaoui, A. Samet and B. Ammar. ‘BATT2GRAPH: A Hybrid CNN-LSTM and Temporal Graph-Based Approach for Lithium-Ion Battery SOH Prediction and Anomaly Detection’. In: *Lecture Notes in Computer Science*. iiWAS 2025 - 27th International Conference on Information Integration and Web Intelligence. Vol. LNCS-6330. Lecture Notes in Computer Science. Matsue, Japan: Springer Nature Switzerland; Springer Nature Switzerland, 3rd Dec. 2025, pp. 127–141. DOI: [10.1007/978-3-032-11976-6_9](https://doi.org/10.1007/978-3-032-11976-6_9). URL: <https://hal.science/hal-05449517> (cit. on p. 30).
- [33] H. A. Akl. ‘Investigating Language Model Capabilities to Represent and Process Formal Knowledge: A Preliminary Study to Assist Ontology Engineering’. In: *Rules and Reasoning, 9th International Joint Conference, RuleML+RR 2025*. the International Joint Conference on Rules and Reasoning (RuleML+RR) 2025. Instabul, Turkey, 21st Sept. 2025. URL: <https://inria.hal.science/hal-05248053> (cit. on p. 31).
- [34] M. Bekbergenova, L. Pradi, B. Navet, E. Tysinger, F. Michel, M. Feraud, Y. Taghzouti, Y. Zhou Chen, O. Kirchhoffer, F. Mehl, M. Legrand, T. Jiang, M. Pagni, S. Hassoun, J.-L. Wolfender, W. Bittremieux, F. Gandon and L.-F. Nothias. ‘MetaboT: AI-based agent for natural language-based interaction with metabolomics knowledge graphs’. In: 33rd annual international conference on Intelligent Systems for Molecular Biology (ISMB 2025) / 24th Annual Conference of the European Conference on Computational Biology (ECCB 2025). Liverpool, United Kingdom, 20th July 2025. URL: <https://hal.science/hal-05235355> (cit. on p. 33).
- [35] H. Belgroun, F. Gandon and F. Michel. ‘Perspectives on a Pigouvian Regulation of the Attention Market’. In: Digital Democracy: Bridging Theory and Practice. 2nd Conference of the European Network for Digital Democracy (EDDY-2025). Paris, France, 12th June 2025. URL: <https://hal.science/hal-05103176> (cit. on p. 34).
- [36] H. Belgroun, F. Michel and F. Gandon. ‘No Such Thing as Free Brain Time: For a Pigouvian Tax on Attention Capture’. In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. Eighth AAAI/ACM Conference on AI, Ethics, and Society (AIES 2025). Vol. 8. 1. Madrid, Spain, 15th Oct. 2025, pp. 367–378. DOI: [10.1609/aies.v8i1.36555](https://doi.org/10.1609/aies.v8i1.36555). URL: <https://hal.science/hal-05240862> (cit. on p. 34).

- [37] S. Bistarelli, V. David, P. Monnin, F. Santini and C. Taticchi. ‘Fast Computing of Dung Semantics in Acyclic Probabilistic Argumentation Frameworks’. In: The 39th Annual AAAI Conference on Artificial Intelligence (AAAI 2025). Thirty-Ninth AAAI Conference on Artificial Intelligence, AAAI 2025. Philadelphia, United States, 25th Feb. 2025. URL: <https://inria.hal.science/hal-04876258> (cit. on p. 30).
- [38] S. Bouveret, H. Gilbert, J. Lang and G. Méroué. ‘Constrained Serial Dictatorships Can Be Fair’. In: IJCAI 2025 - Thirty-Fourth International Joint Conference on Artificial Intelligence. Montreal, Canada: International Joint Conferences on Artificial Intelligence Organization, 20th May 2025, pp. 3762–3770. DOI: [10.24963/ijcai.2025/418](https://doi.org/10.24963/ijcai.2025/418). URL: <https://hal.science/hal-05393240>.
- [39] M. Buffa, M. Winckler, Q. Escobar, S. Demont, A. Hofr and A. Mir-Sadjadi. ‘Interactive Audio Sculpting: Plugin Customization and UI Affordances in Immersive Environments’. In: IS 2025 - IEEE International Symposium on the Internet of Sounds. L’Aquila, Italy, 29th Oct. 2025. URL: <https://inria.hal.science/hal-05388540> (cit. on pp. 14, 33).
- [40] M. Buffa, M. Winckler, Q. Escobar, S. Demont, A. Hofr and A. Mir-Sadjadi. ‘Using Web Audio Plugins in 3D environments’. In: <https://zenodo.org/communities/wac/>. WAC 2025 - Web Audio Conference. Vol. <https://zenodo.org/communities/wac/>. Paris, France: Zenodo, 2025. DOI: [10.5281/zenodo.17641873](https://doi.org/10.5281/zenodo.17641873). URL: <https://inria.hal.science/hal-05388555> (cit. on p. 33).
- [41] C. da Costa Pereira, C. Pasquier and A. G. B. Tettamanzi. ‘Propagation-Based Domain-Transferable Gradual Sentiment Analysis’. In: *SciTePress digital library*. ICAART 2025 - 17th International Conference on Agents and Artificial Intelligence. Vol. 3. Proceedings of the 17th International Conference on Agents and Artificial Intelligence - Volume 3: ICAART. Porto, Portugal: SCITEPRESS - Science and Technology Publications, 2025, pp. 520–527. DOI: [10.5220/0013158200003890](https://doi.org/10.5220/0013158200003890). URL: <https://hal.science/hal-05010804> (cit. on p. 28).
- [42] N.-E. Mbengue, M. Couceiro, Y. Lepage and P. Monnin. ‘Which Pairs to Choose? Exploring Analogical Competency for Knowledge Graph Pruning’. In: IARML 2025 - Workshop on the Interactions between Analogical Reasoning and Machine Learning at IJCAI 2025. Montreal, Canada, 16th Aug. 2025. URL: <https://inria.hal.science/hal-05162211> (cit. on p. 30).
- [43] A. Menin and C. Faron. ‘DIVA: An Ontology-based Approach to Model User Activity within Visualization Systems’. In: SAC 2026 - 41st ACM/SIGAPP Symposium on Applied Computing. Thessaloniki, Greece, 23rd Mar. 2026. DOI: [10.1145/3748522.3779893](https://doi.org/10.1145/3748522.3779893). URL: <https://inria.hal.science/hal-05403198> (cit. on p. 25).
- [44] N. B. Ocampo, E. Cabrio and S. Villata. ‘From Hidden to Harmful: Connecting Implicit and Explicit Hate Through Implied Statements’. In: WI-IAT 2025 - 24th IEEE/WIC International Conference on Web Intelligence and Intelligent Agent Technology. London, United Kingdom, 15th Nov. 2025. URL: <https://hal.science/hal-05453245>.
- [45] L. Pradi, T. Jiang, M. Feraud, M. Bekbergenova, Y. Taghzouti and L.-F. Nothias. ‘An AI Pipeline for Scientific Literacy and Discovery: a Demonstration of Perspicacité-AI Integration with Knowledge Graphs’. In: Joint Proceedings of Industry, Doctoral Consortium, Posters and Demos of the 24th International Semantic Web Conference (ISWC-C 2025). Vol. Vol-4085. Nara, Japan, 31st Oct. 2025, pp. 462–467. URL: <https://hal.science/hal-05290005> (cit. on p. 34).
- [46] C. Quere, A. Menin, E. Bernart, C. Maria Dal Sasso Freitas, L. Nedel, P. Moura, H.-Y. Wu and M. Winckler. ‘WristNotes: Detachable Menu for Annotations in Immersive Environments’. In: 20th IFIP TC13 International Conference on Human-Computer Interaction (INTERACT 2025). Belo Horizonte, Minas Gerais, Brazil, 8th Sept. 2025. URL: <https://hal.science/hal-05155843> (cit. on p. 33).
- [47] C. Quere, A. Menin, H.-Y. Wu, P. Moura and M. Winckler. ‘INSPIRE : An INcremental, SPatio-temporal and Immersive approach for Representing and Exploring mobility data’. In: *ACM digital library*. CHIItaly 2025 - 16th Biannual Conference of the Italian SIGCHI Chapter. Salerno, Italy, 6th Oct. 2025. DOI: [10.1145/3750069.3750444](https://doi.org/10.1145/3750069.3750444). URL: <https://hal.science/hal-05178219> (cit. on p. 33).

- [48] C. Ringwald, F. Gandon, C. Faron, F. Michel and H. A. Akl. ‘Kastor: Fine-tuned Small Language Models for Shape-based Active Relation Extraction’. In: *Extended Semantic Web Conference 2025*. Portoroz, France, 1st June 2025. DOI: [10.1007/978-3-031-94575-5_6](https://doi.org/10.1007/978-3-031-94575-5_6). URL: <https://hal.science/hal-05078493> (cit. on p. 29).
- [49] C. Ringwald, F. Gandon, C. Faron, F. Michel and H. Akl. ‘Overcoming the Generalization Limits of SLM Finetuning for Shape-Based Extraction of Datatype and Object Properties’. In: *ACM digital library. KCAP 2025 - Thirteenth International Conference on Knowledge Capture. K-CAP ’25: Proceedings of the 13th Knowledge Capture Conference 2025*. Dayton, OH, USA,, United States: ACM, 15th Oct. 2025, pp. 9–17. DOI: [10.1145/3731443.3771342](https://doi.org/10.1145/3731443.3771342). URL: <https://hal.science/hal-05285428> (cit. on p. 29).
- [50] F. A. S. Robert, H.-Y. Wu, L. Sassatelli and M. Winckler. ‘The Triangle of Misunderstanding in Interactive Virtual Narratives: Gulfs Between System, Designers and Players’. In: *CHI 2025 - CHI Conference on Human Factors in Computing Systems*. Yokohama, Japan, 25th Apr. 2025, pp. 1–16. DOI: [10.1145/3706598.3713700](https://doi.org/10.1145/3706598.3713700). URL: <https://inria.hal.science/hal-05241510> (cit. on p. 32).
- [51] N. Robert, F. Gandon and M. Lefrançois. ‘OLIVAW: ACIMOV’s GitHub robot assisting agile collaborative ontology development’. In: *IEEE Xplore. 24th IEEE /WIC International Conference on Web Intelligence and Intelligent Agent Technology*. London, United Kingdom, 15th Nov. 2025. URL: <https://hal.science/hal-05291789> (cit. on p. 26).
- [52] N. Robert, P. Monnin and C. Faron. ‘Semantically Enriched Datasets for Link Prediction: DB100k+, NELL-995+ and YAGO3-10+’. In: *1st International Workshop on Advanced Neuro-Symbolic Applications Co-located with ECAI 2025*. Bologna, Italy, 26th Oct. 2025. URL: <https://hal.science/hal-05291884> (cit. on pp. 24, 29).
- [53] Y. Taghzouti, F. Michel, T. Jiang, L.-F. Nothias and F. Gandon. ‘User Interface and Agent Interface for Online Generation of Knowledge Graph’s Competency Questions and Question-Query Training Sets’. In: *Proceedings of the Second International Workshop on Retrieval-Augmented Generation Enabled by Knowledge Graphs (RAGE-KG 2025)co-located with the 24th International Semantic Web Conference (ISWC 2025)*. RAGE-KG 2025: The Second International Workshop on Retrieval-Augmented Generation Enabled by Knowledge Graphs, co-located with ISWC 2025. Vol. 4079. Nara, Japan, 27th Oct. 2025, pp. 153–159. URL: <https://hal.science/hal-05289962> (cit. on p. 34).
- [54] Y. Taghzouti, F. Michel, T. Jiang, L.-F. Nothias and F. Gandon. ‘Q²Forge: Minting Competency Questions and SPARQL Queries for Question-Answering Over Knowledge Graphs’. In: *Knowledge Capture Conference 2025 (K-CAP ’25)*, December 10–12, 2025, Dayton, OH, USA. Dayton, OH, United States, 10th Dec. 2025. DOI: [10.1145/3731443.3771350](https://doi.org/10.1145/3731443.3771350). URL: <https://hal.science/hal-05070442> (cit. on p. 34).

National peer-reviewed Conferences

- [55] S. Taoufiq, M. Tounsi Dhouib, A. Daubresse, C. Faron, M. Hayotte and S. Scoffier-Mériaux. ‘Conception et implémentation d’une ontologie des facteurs psychosociaux liés aux attitudes et comportements alimentaires dysfonctionnels chez les sportifs’. In: *Revue des Nouvelles Technologies de l’Information*. Conférence Francophone sur l’Extraction et la Gestion des Connaissances. Anglet, France, 26th Jan. 2026. URL: <https://hal.science/hal-05392013> (cit. on p. 25).

Scientific book chapters

- [56] A. Ciortea, M. Castellucci, K. Schultz, J. Lemée, D. Vachtsevanou, F. Gandon, S. Mayer, V. Berger, S. Burattini, A. Ricci and O. Boissier. ‘Yggdrasil: An Artifact-Based Framework for Hypermedia Multi-Agent Systems’. In: *Agents and Multi-Agent Systems Development*. Springer Nature Switzerland, 8th Jan. 2026, pp. 181–207. DOI: [10.1007/978-3-032-01082-7_7](https://doi.org/10.1007/978-3-032-01082-7_7). URL: <https://inria.hal.science/hal-05460210> (cit. on p. 31).

Edition (books, proceedings, special issue of a journal)

- [57] *Proceedings of the Second International Workshop on Hypermedia Multi-Agent Systems (HyperAgents 2025) co-located with 28th European Conference on Artificial Intelligence (ECAI 2025)*. HyperAgents 2025 : Hypermedia Multi-Agent Systems 2025. Vol. Vol-4084. Bologna, Italy: CEUR-WS.org, 30th Oct. 2025. URL: <https://hal.science/hal-05356585> (cit. on p. 45).

Reports & preprints

- [58] C. Leturc, F. Gandon and M. Lefrançois. *Possible Hyperworlds: representing Kripke models in RDF*. 10th Dec. 2025. URL: <https://hal.science/hal-05329535>.
- [59] P. Maillot, C. Faron, F. Gandon, F. Michel and P. Monnin. *Eat your own KR: a KR-based approach to index Semantic Web Endpoints and Knowledge Graphs*. INRIA, 2nd July 2025. URL: <https://hal.science/hal-05140375> (cit. on p. 28).
- [60] A. Ollagnier and A. Menin. *DataLens: Enhancing Dataset Discovery via Network Topologies*. 30th July 2025. URL: <https://hal.science/hal-05189348> (cit. on p. 33).
- [61] C. Ringwald, F. Gandon, C. Faron, F. Michel and H. A. Akl. *A systematic review of relation extraction task since the emergence of Transformers*. 21st Nov. 2025. URL: <https://hal.science/hal-05349792> (cit. on p. 28).
- [62] Y. Taghzouti, F. Michel, T. Jiang, L.-F. Nothias and F. Gandon. *Q²Forge : Générer des Collections de Questions de Compétences et Requêtes SPARQL pour Interroger des Graphes de Connaissances en Langue Naturelle*. 4th Dec. 2025. URL: <https://hal.science/hal-05399933> (cit. on p. 34).

Other scientific publications

- [63] M. Bekbergenova, L. Pradi, B. Navet, Y. Taghzouti, E. Tysinger, J.-L. Wolfender, W. Bittremieux, F. Gandon and L.-F. Nothias. ‘MetaboT: AI-based agent for interaction with metabolomics knowledge graphs’. In: EDSFA 2nd year PhD Workshop. Sophia -Antipolis, France, 20th May 2025. URL: <https://hal.science/hal-05235422> (cit. on p. 33).
- [64] M. Bekbergenova, L. Pradi, E. Tysinger, B. Navet, O. Kirchhoffer, J.-L. Wolfender, F. Mehl, M. Pagni, W. Bittremieux, F. Gandon and L.-F. Nothias. ‘MetaboT: A conversational AI agent for accessible Mass Spectrometry Metabolomics Data Mining’. In: Benelux Metabolomics Days 2025. Leiden, Netherlands, 27th Aug. 2025. URL: <https://hal.science/hal-05235376> (cit. on p. 33).
- [65] M. Bekbergenova, L. Pradi, E. Tysinger, B. Navet, O. Kirchhoffer, J.-L. Wolfender, F. Mehl, M. Pagni, W. Bittremieux, F. Gandon and L.-F. Nothias. ‘MetaboT: A Conversational AI-Agent for Accessible Mass Spectrometry Metabolomics Data Mining’. In: The 21st Annual Conference of the Metabolomics Society. Prague, Czech Republic, 22nd June 2025. URL: <https://hal.science/hal-05235334> (cit. on p. 33).
- [66] M. Buffa, M. Winckler, Q. Escobar, S. Demont, A. Hofr and A. Mir-Sadjadi. ‘Creating Collaborative and immersive Modular Audio Installations with WAM Jam Party’. In: IS 2025 - IEEE International Symposium on the Internet of Sounds. L’aquila, Italy, 29th Oct. 2025. URL: <https://inria.hal.science/hal-05388549> (cit. on p. 33).
- [67] M. Buffa, M. Winckler, Q. Escobar, S. Demont, A. Hofr and A. Mir-Sadjadi. ‘WAM Jam Party: a collaborative real-time music creation experience in VR’. In: WAC 2025 - Web Audio Conference. Paris Auditorium du Petit Palais, France: Zenodo, 2025. DOI: [10.5281/zenodo.17761181](https://doi.org/10.5281/zenodo.17761181). URL: <https://inria.hal.science/hal-05388770> (cit. on p. 33).
- [68] M. Buffa, M. Winckler and A. Mir-Sadjadi. ‘Embodied Virtual Instruments in Web-Based Multi-User VR: A Case Study with a 3D Drum Kit and Web Audio Modules’. In: WAC 2025 - 9th Web Audio Conference. Paris, France, 19th Nov. 2025. URL: <https://hal.science/hal-05363346> (cit. on p. 33).

- [69] Y. Taghzouti, F. Michel, T. Jiang, L.-F. Nothias and F. Gandon. ‘Building Questions and Queries Datasets for Knowledge Graphs: a Demo of Q²Forge’. In: Joint Proceedings of Industry, Doctoral Consortium, Posters and Demos of the 24th International Semantic Web Conference (ISWC-C 2025). Vol. Vol-4085. Nara, Japan, 31st Oct. 2025, pp. 408–413. URL: <https://hal.science/hal-05289986> (cit. on pp. 14, 34).
- [70] Y. Taghzouti, F. Michel, T. Jiang, L.-F. Nothias and F. Gandon. ‘Building Questions and Queries Datasets for Knowledge Graphs: a Demo of Q²Forge Pipeline’. In: ISWC 2025: The 24th International Semantic Web Conference. CEUR-4085. Nara, Japan, 2nd Nov. 2025. URL: <https://hal.science/hal-05332768> (cit. on p. 34).
- [71] Y. Taghzouti, L.-F. Nothias, F. Michel and F. Gandon. ‘Gen²KGBot: Generic Generative Knowledge Graph Robot: a conversational agent to access (scientific) data and knowledge in natural language’. In: AI Action Summit. Paris, France, 6th Feb. 2025. URL: <https://hal.science/hal-05333244> (cit. on p. 34).

12.3 Cited publications

- [72] N. Y. Ayadi, C. Faron, F. Michel, F. Gandon and O. Corby. ‘Computing and Visualizing Agro-Meteorological Parameters based on an Observational Weather Knowledge Graph’. In: *Companion Proceedings of the ACM Web Conference 2023, WWW 2023, Austin, TX, USA, 30 April 2023 - 4 May 2023*. ACM, 2023, pp. 242–245. DOI: [10.1145/3543873.3587357](https://doi.org/10.1145/3543873.3587357). URL: <https://doi.org/10.1145/3543873.3587357> (cit. on p. 12).
- [73] O. Boissier, A. Ciorcea, F. Gandon, S. Mayer and A. Ricci. ‘Hypermedia Multi-Agent Systems’. In: *Agents on the Web*. Vol. 23081. Wadern, Germany: Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Feb. 2023, pp. 117–118. URL: <https://hal-emse.ccsd.cnrs.fr/emse-04240524> (cit. on p. 12).
- [74] [SW Rel.] R. Cérés, O. Corby, E. Demairy and F. Gandon, *Corese* version 4.5.0, July 2023. Inria ; CNRS ; Université Côte d’Azur. LIC: CECILL-C. DOI: [10.5281/zenodo.10376241](https://zenodo.org/record/10376241), HAL: [hal-04170333v1](https://hal.science/hal-04170333v1), URL: <https://hal.science/hal-04170333>, vcs: <https://github.com/Wimmics/corese>, SWHID: [swh:1:dir:7ec9a0bccba1254457985de26d524d0d5c4a9c5](https://www.swinsight.org/record/7ec9a0bccba1254457985de26d524d0d5c4a9c5) (cit. on p. 26).
- [75] P.-A. Champin and R. Wenning. ‘Knowledge representation and reasoning in personal knowledge graphs’. en. In: *Personal Knowledge Graphs (PKGs): Methodology, tools and applications*. IET Digital Library, Nov. 2023, pp. 55–81. DOI: [10.1049/PBPC063E_ch3](https://doi.org/10.1049/PBPC063E_ch3). URL: https://digital-library.theiet.org/content/books/10.1049/pbpc063e_ch3 (visited on 16/03/2024) (cit. on p. 10).
- [76] O. Corby, R. Dieng-Kuntz and C. Faron-Zucker. ‘Querying the Semantic Web with Corese Search Engine’. In: *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI’2004, including Prestigious Applicants of Intelligent Systems, PAIS 2004, Valencia, Spain, August 22-27, 2004*. Ed. by R. L. de Mántaras and L. Saitta. IOS Press, 2004, pp. 705–709 (cit. on p. 11).
- [77] O. Corby, R. Dieng-Kuntz, C. Faron-Zucker and F. Gandon. ‘Searching the Semantic Web: Approximate Query Processing Based on Ontologies’. In: *IEEE Intell. Syst.* 21.1 (2006), pp. 20–27. DOI: [10.1109/MIS.2006.16](https://doi.org/10.1109/MIS.2006.16). URL: <https://doi.org/10.1109/MIS.2006.16> (cit. on p. 11).
- [78] O. Corby, C. Faron, F. Gandon, D. Graux and F. Michel. ‘Beyond Classical SERVICE Clause in Federated SPARQL Queries: Leveraging the Full Potential of URI Parameters’. In: *Proceedings of the 17th International Conference on Web Information Systems and Technologies, WEBIST 2021, October 26-28, 2021*. Ed. by F. J. D. Mayo, M. Marchiori and J. Filipe. SCITEPRESS, 2021, pp. 65–76. DOI: [10.5220/0010660300003058](https://doi.org/10.5220/0010660300003058). URL: <https://doi.org/10.5220/0010660300003058> (cit. on p. 11).

- [79] O. Corby and C. Faron-Zucker. ‘A Transformation Language for RDF Based on SPARQL’. In: *Web Information Systems and Technologies - 11th International Conference, WEBIST 2015, Lisbon, Portugal, May 20-22, 2015, Revised Selected Papers*. Ed. by V. Monfort, K. Krempels, T. A. Majchrzak and Z. Turk. Vol. 246. Lecture Notes in Business Information Processing. Springer, 2015, pp. 318–340. DOI: [10.1007/978-3-319-30996-5_16](https://doi.org/10.1007/978-3-319-30996-5_16). URL: https://doi.org/10.1007/978-3-319-30996-5%5C_16 (cit. on p. 11).
- [80] O. Corby and C. Faron-Zucker. ‘Implementation of SPARQL Query Language Based on Graph Homomorphism’. In: *Conceptual Structures: Knowledge Architectures for Smart Applications, 15th International Conference on Conceptual Structures, ICCS 2007, Sheffield, UK, July 22-27, 2007, Proceedings*. Ed. by U. Priss, S. Polovina and R. Hill. Vol. 4604. Lecture Notes in Computer Science. Springer, 2007, pp. 472–475. DOI: [10.1007/978-3-540-73681-3_37](https://doi.org/10.1007/978-3-540-73681-3_37). URL: https://doi.org/10.1007/978-3-540-73681-3%5C_37 (cit. on p. 11).
- [81] O. Corby, C. Faron-Zucker and F. Gandon. ‘A Generic RDF Transformation Software and Its Application to an Online Translation Service for Common Languages of Linked Data’. In: *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II*. Ed. by M. Arenas, Ó. Corcho, E. Simperl, M. Strohmaier, M. d’Aquin, K. Srinivas, P. Groth, M. Dumontier, J. Heflin, K. Thirunarayan and S. Staab. Vol. 9367. Lecture Notes in Computer Science. Springer, 2015, pp. 150–165. DOI: [10.1007/978-3-319-25010-6_9](https://doi.org/10.1007/978-3-319-25010-6_9). URL: https://doi.org/10.1007/978-3-319-25010-6%5C_9 (cit. on p. 11).
- [82] C. da Costa Pereira, D. Dubois, H. Prade and A. G. B. Tettamanzi. ‘Parsimonious Representation of Knowledge Uncertainty using Metadata about Validity and Completeness’. In: *ICAART 2022 - 14th International Conference on Agents and Artificial Intelligence*. Vol. 2. Proceedings of the 14th International Conference on Agents and Artificial Intelligence - (Volume 2). Online Streaming, Portugal: SCITEPRESS : Science and Technology Publications, Feb. 2022, pp. 441–449. DOI: [10.5220/0010874000003116](https://hal.science/hal-03719603). URL: <https://hal.science/hal-03719603> (cit. on p. 10).
- [83] C. d’Amato, L. Mahon, P. Monnin and G. Stamou. ‘Machine Learning and Knowledge Graphs: Existing Gaps and Future Research Challenges’. In: *Transactions on Graph Data and Knowledge 1.1* (Dec. 2023), pp. 1–35. DOI: [10.4230/TGDK.1.1.8](https://inria.hal.science/hal-04353543). URL: <https://inria.hal.science/hal-04353543> (cit. on p. 10).
- [84] A. Delteil, C. Faron-Zucker and R. Dieng. ‘Learning Ontologies from RDF annotations’. In: *IJCAI’2001 Workshop on Ontology Learning, Proceedings of the Second Workshop on Ontology Learning OL’2001, Seattle, USA, August 4, 2001 (Held in conjunction with the 17th International Conference on Artificial Intelligence IJCAI’2001)*. Vol. 38. CEUR Workshop Proceedings. CEUR-WS.org, 2001. URL: https://ceur-ws.org/Vol-38/delteil%5C_ol.pdf (cit. on p. 11).
- [85] M. T. Dhoubi, C. Faron and A. G. B. Tettamanzi. ‘Measuring Clusters of Labels in an Embedding Space to Refine Relations in Ontology Alignment’. In: *J. Data Semant.* 10.3-4 (2021), pp. 399–408. DOI: [10.1007/s13740-021-00137-8](https://doi.org/10.1007/s13740-021-00137-8). URL: <https://doi.org/10.1007/s13740-021-00137-8> (cit. on p. 11).
- [86] A. E. A. Djebri, A. G. B. Tettamanzi and F. Gandon. ‘Publishing Uncertainty on the Semantic Web: Blurring the LOD Bubbles’. In: *ICCS 2019 - 24th International Conference on Conceptual Structures*. Marburg, Germany, July 2019, pp. 42–56. DOI: [10.1007/978-3-030-23182-8_4](https://inria.hal.science/hal-02167174). URL: <https://inria.hal.science/hal-02167174> (cit. on p. 10).
- [87] A. E. A. Djebri, A. G. B. Tettamanzi and F. Gandon. ‘Task-Oriented Uncertainty Evaluation for Linked Data Based on Graph Interlinks’. In: *EKAW 2020 - 22nd International Conference on Knowledge Engineering and Knowledge Management*. Bozen-Bolzano, Italy, Sept. 2020. DOI: [10.1007/978-3-030-61244-3_15](https://hal.science/hal-02933190). URL: <https://hal.science/hal-02933190> (cit. on p. 10).
- [88] R. Felin, C. Faron and A. G. B. Tettamanzi. ‘A Framework to Include and Exploit Probabilistic Information in SHACL Validation Reports’. In: *The Semantic Web - 20th International Conference, ESWC 2023, Hersonissos, Crete, Greece, May 28 - June 1, 2023, Proceedings*. Ed. by C. Pesquita, E. Jiménez-Ruiz, J. P. McCusker, D. Faria, M. Dragoni, A. Dimou, R. Troncy and S. Hertling. Vol. 13870. Lecture Notes in Computer Science. Springer, 2023, pp. 91–104. DOI: [10.1007/978-3-030-61244-3_15](https://doi.org/10.1007/978-3-030-61244-3_15)

- 3-031-33455-9_6. URL: https://doi.org/10.1007/978-3-031-33455-9%5C_6 (cit. on p. 11).
- [89] R. Felin, P. Monnin and C. F. A. G. B. Tettamanzi. ‘An Algorithm Based on Grammatical Evolution for Discovering SHACL Constraints’. In: *Genetic Programming - 27th European Conference, EuroGP 2024, Held as Part of EvoStar 2024, Aberystwyth, Wales, April 3-5, 2024, Proceedings*. Lecture Notes in Computer Science. 2024 (cit. on p. 11).
- [90] S. Firmenich, G. A. Bosetti, G. Rossi, M. Winckler and T. Barbieri. ‘Abstracting and Structuring Web Contents for Supporting Personal Web Experiences’. In: *Web Engineering - 16th International Conference, ICWE 2016, Lugano, Switzerland, June 6-9, 2016. Proceedings*. Vol. 9671. Lecture Notes in Computer Science. Springer, 2016, pp. 77–95. doi: [10.1007/978-3-319-38791-8_5](https://doi.org/10.1007/978-3-319-38791-8_5). URL: https://doi.org/10.1007/978-3-319-38791-8%5C_5 (cit. on p. 13).
- [91] F. Gandon. ‘Merry hMAS and Happy New Web: A Wish for Standardizing an AI-Friendly Web Architecture for Hypermedia Multi-Agent Systems’. In: *Dagstuhl-Seminar 21072 : Autonomous Agents on the Web*. Dagstuhl Schloss. Dagstuhl, Germany, Feb. 2021, p. 3. URL: <https://inria.hal.science/hal-03960471> (cit. on p. 12).
- [92] R. Gazzotti and F. Gandon. ‘When owl: sameAs is the Same: Experimenting Online Resolution of Identity with SPARQL Queries to Linked Open Data Sources’. In: *Proceedings of the 17th International Conference on Web Information Systems and Technologies, WEBIST 2021, October 26-28, 2021*. Ed. by F. J. D. Mayo, M. Marchiori and J. Filipe. SCITEPRESS, 2021, pp. 41–52. doi: [10.5220/0010654400003058](https://doi.org/10.5220/0010654400003058). URL: <https://doi.org/10.5220/0010654400003058> (cit. on p. 11).
- [93] F.-Z. Hannou, V. Charpenay, M. Lefrançois, C. Roussey, A. Zimmermann and F. Gandon. ‘The ACIMOV Methodology: Agile and Continuous Integration for Modular Ontologies and Vocabularies’. In: *MK 2023 - 2nd Workshop on Modular Knowledge associated with FOIS 2023 - the 13th International Conference on Formal Ontology in Information Systems*. Sherbrooke, Canada, July 2023. URL: <https://laas.hal.science/hal-04187236> (cit. on p. 10).
- [94] F.-Z. H. Hannou, V. Charpenay, M. Lefrançois, C. Roussey, A. Zimmermann and F. Gandon. ‘La méthodologie ACIMOV pour l’intégration agile et continue des modules ontologiques’. In: *35. Journées francophones d’Ingénierie des Connaissances (IC 2024) @ Plate-Forme Intelligence Artificielle (PFIA 2024)*. La Rochelle, France, July 2024, pp. 127–128. URL: <https://hal-emse.ccsd.cnrs.fr/emse-04636109> (cit. on p. 26).
- [95] H. Huang and F. Gandon. ‘Learning URI Selection Criteria to Improve the Crawling of Linked Open Data (Extended Abstract)’. In: *IJCAI 2020 - 29th International Joint Conference on Artificial Intelligence*. Yokohama, Japan, Jan. 2021. URL: <https://inria.hal.science/hal-03064912> (cit. on p. 11).
- [96] N. Hubert, P. Monnin, A. Brun and D. Monticolo. ‘Treat Different Negatives Differently: Enriching Loss Functions with Domain and Range Constraints for Link Prediction’. In: *Lecture Notes in Computer Science. Semantic Web - 21st International Conference, ESWC 2024, Hersonissos, Greece, May 2024*. URL: <https://inria.hal.science/hal-04491230> (cit. on p. 29).
- [97] N. Hubert, P. Monnin, A. Brun and D. Monticolo. ‘Treat Different Negatives Differently: Enriching Loss Functions with Domain and Range Constraints for Link Prediction’. In: *The Semantic Web - 21st International Conference, ESWC 2024, Hersonissos, Crete, Greece, May 26 - 30, 2024, Proceedings*. 2024. URL: <https://doi.org/10.48550/arXiv.2303.00286> (cit. on p. 11).
- [98] L. Jarnac, M. Couceiro and P. Monnin. ‘Relevant Entity Selection: Knowledge Graph Bootstrapping via Zero-Shot Analogical Pruning’. In: *CIKM ’23: The 32nd ACM International Conference on Information and Knowledge Management*. CIKM ’23: The 32nd ACM International Conference on Information and Knowledge Management. Birmingham (UK), United Kingdom: ACM, Oct. 2023, pp. 934–944. doi: [10.1145/3583780.3615030](https://doi.org/10.1145/3583780.3615030). URL: <https://inria.hal.science/hal-04254979> (cit. on p. 30).

- [99] P. Maillot, J. Andersen, S. Cazalens, C. Faron, F. Gandon, P. Lamarre and F. Michel. ‘An Open Platform for Quality Measures in a Linked Data Index’. In: *WWW '24: The ACM Web Conference 2024*. Singapore, Singapore: ACM, May 2024, pp. 1087–1090. doi: [10.1145/3589335.3651443](https://doi.org/10.1145/3589335.3651443). URL: <https://inria.hal.science/hal-04575211> (cit. on p. 28).
- [100] P. Maillot, O. Corby, C. Faron, F. Gandon and F. Michel. ‘IndeGx: A Model and a Framework for Indexing RDF Knowledge Graphs with SPARQL-based Test Suites’. In: *Journal of Web Semantics* (Jan. 2023). doi: [10.1016/j.websem.2023.100775](https://doi.org/10.1016/j.websem.2023.100775). URL: <https://hal.science/hal-03946680> (cit. on pp. 11, 28).
- [101] P. Maillot, O. Corby, C. Faron, F. Gandon and F. Michel. ‘Metadatamatic: A Web application to Create a Dataset Description’. In: *WWW '23 - The ACM Web Conference 2023*. WWW '23 Companion: Companion Proceedings of the ACM Web Conference 2023. Austin TX, United States: ACM, Apr. 2023, pp. 123–126. doi: [10.1145/3543873.3587328](https://doi.org/10.1145/3543873.3587328). URL: <https://hal.science/hal-04171196> (cit. on p. 10).
- [102] A. Menin, M. N. Do, C. Dal Sasso Freitas, O. Corby, C. Faron, A. Giboin and M. Winckler. ‘Using Chained Views and Follow-up Queries to Assist the Visual Exploration of the Web of Big Linked Data’. In: *International Journal of Human-Computer Interaction* (Aug. 2022). doi: [10.1080/10447318.2022.2112529](https://doi.org/10.1080/10447318.2022.2112529). URL: <https://hal.science/hal-03518845> (cit. on p. 13).
- [103] A. Menin, F. Michel, F. Gandon, R. Gazzotti, E. Cabrio, O. Corby, A. Giboin, S. Marro, T. Mayer, S. Villata and M. Winckler. ‘Covid-on-the-Web: Exploring the COVID-19 Scientific Literature through Visualization of Linked Data from Entity and Argument Mining’. In: *Quantitative Science Studies* (Nov. 2021). doi: [10.1162/qss_a_00164](https://doi.org/10.1162/qss_a_00164). URL: <https://hal.science/hal-03404580> (cit. on p. 12).
- [104] F. Michel, C. Faron-Zucker, O. Corby and F. Gandon. ‘Enabling Automatic Discovery and Querying of Web APIs at Web Scale using Linked Data Standards’. In: *Companion of The 2019 World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*. Ed. by S. Amer-Yahia, M. Mahdian, A. Goel, G. Houben, K. Lerman, J. J. McAuley, R. Baeza-Yates and L. Zia. ACM, 2019, pp. 883–892. doi: [10.1145/3308560.3317073](https://doi.org/10.1145/3308560.3317073). URL: <https://doi.org/10.1145/3308560.3317073> (cit. on p. 11).
- [105] F. Michel, C. Faron-Zucker and J. Montagnat. ‘A Generic Mapping-based Query Translation from SPARQL to Various Target Database Query Languages’. In: *Proceedings of the 12th International Conference on Web Information Systems and Technologies, WEBIST 2016, Volume 2, Rome, Italy, April 23-25, 2016*. Ed. by T. A. Majchrzak, P. Traverso, V. Monfort and K. Krempels. SciTePress, 2016, pp. 147–158. doi: [10.5220/0005905401470158](https://doi.org/10.5220/0005905401470158). URL: <https://doi.org/10.5220/0005905401470158> (cit. on p. 11).
- [106] F. Michel, C. Faron-Zucker and J. Montagnat. ‘Bridging the Semantic Web and NoSQL Worlds: Generic SPARQL Query Translation and Application to MongoDB’. In: *Trans. Large Scale Data Knowl. Centered Syst.* 40 (2019), pp. 125–165. doi: [10.1007/978-3-662-58664-8_5](https://doi.org/10.1007/978-3-662-58664-8_5). URL: https://doi.org/10.1007/978-3-662-58664-8_5 (cit. on p. 11).
- [107] F. Michel and F. Gandon. ‘Pay Attention: a Call to Regulate the Attention Market and Prevent Algorithmic Emotional Governance’. In: *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. Vol. 7 (2024): Proceedings of the Seventh AAAI/ACM Conference on AI, Ethics, and Society (AIES-24) 7. AAAI/ACM. San Jose (CA), United States: AAAI Press, Washington, DC, USA, Oct. 2024. doi: [10.1609/aies.v7i1.31696](https://doi.org/10.1609/aies.v7i1.31696). URL: <https://hal.science/hal-04758276> (cit. on p. 34).
- [108] F. Michel, O. Gargominy, S. Tercerie and C. Faron Zucker. ‘A Model to Represent Nomenclatural and Taxonomic Information as Linked Data. Application to the French Taxonomic Register, TAXREF’. In: *ISWC 2017 Workshop on Semantics for Biodiversity (S4Biodiv 2017)*. Vol. CEUR Vol. 1933. Vienna, Austria, Oct. 2017, pp. 1–12. URL: <https://hal.science/hal-01617708> (cit. on p. 23).
- [109] T. H. Nguyen. ‘Mining the semantic Web for OWL axioms. (Fouille du Web sémantique à la recherche d’axiomes OWL)’. PhD thesis. University of Côte d’Azur, Nice, France, 2021. URL: <https://tel.archives-ouvertes.fr/tel-03406784> (cit. on p. 11).

- [110] C. Ringwald. ‘Learning Pattern-Based Extractors from Natural Language and Knowledge Graphs: Applying Large Language Models to Wikipedia and Linked Open Data’. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. No. 21: IAAI-24, EAAI-24, AAAI-24 Student Abstracts, Undergraduate Consortium and Demonstrations 21. Vancouver, Canada, Feb. 2024, pp. 23411–23412. DOI: [10.1609/aaai.v38i21.30406](https://doi.org/10.1609/aaai.v38i21.30406). URL: <https://hal.science/hal-04526050> (cit. on p. 28).
- [111] C. Ringwald, F. Gandon, C. Faron, F. Michel and H. A. Akl. ‘12 shades of RDF: Impact of Syntaxes on Data Extraction with Language Models’. In: *Lecture Notes in Computer Science*. Vol. 15344. Lecture Notes in Computer Science. Accepted at ESWC 2024. Hersonissos, Greece: Springer Nature Switzerland, May 2024, pp. 81–91. DOI: [10.1007/978-3-031-78952-6_8](https://doi.org/10.1007/978-3-031-78952-6_8). URL: <https://hal.science/hal-04581124> (cit. on pp. 14, 29).
- [112] C. Ringwald, F. Gandon, C. Faron, F. Michel and H. A. Akl. ‘Well-Written Knowledge Graphs: Most Effective RDF Syntaxes for Triple Linearization in End-to-End Extraction of Relations from Texts (Student Abstract)’. In: *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*. AAAI Press, 2024, pp. 23631–23632. DOI: [10.1609/AAAI.V38I21.30502](https://doi.org/10.1609/AAAI.V38I21.30502). URL: <https://doi.org/10.1609/aaai.v38i21.30502> (cit. on pp. 10, 28).
- [113] S. Sarkadi, A. G. B. Tettamanzi and F. Gandon. ‘Interoperable AI: Evolutionary Race Towards Sustainable Knowledge Sharing’. In: *IEEE Internet Computing*. IEEE Internet Computing 26.6 (Oct. 2022). Article and Supplementary Material, p. 8. DOI: [10.1109/MIC.2022.3214378](https://doi.org/10.1109/MIC.2022.3214378). URL: <https://inria.hal.science/hal-03807461> (cit. on p. 11).
- [114] T. R. Silva, J. Hak and M. Winckler. ‘A Formal Ontology for Describing Interactive Behaviors and Supporting Automated Testing on User Interfaces’. In: *Int. J. Semantic Comput.* 11.4 (2017), pp. 513–540. DOI: [10.1142/S1793351X17400219](https://doi.org/10.1142/S1793351X17400219). URL: <https://doi.org/10.1142/S1793351X17400219> (cit. on p. 13).
- [115] A. G. B. Tettamanzi, C. Faron-Zucker and F. Gandon. ‘Possibilistic testing of OWL axioms against RDF data’. In: *Int. J. Approx. Reason.* 91 (2017), pp. 114–130. DOI: [10.1016/j.ijar.2017.08.012](https://doi.org/10.1016/j.ijar.2017.08.012). URL: <https://doi.org/10.1016/j.ijar.2017.08.012> (cit. on p. 11).
- [116] M. Tikat, A. Menin, M. Buffa and M. Winckler. ‘Engineering Annotations to Support Analytical Provenance in Visual Exploration Processes’. In: *Web Engineering - 22nd International Conference, ICWE 2022, Bari, Italy, July 5-8, 2022, Proceedings*. Ed. by T. D. Noia, I. Ko, M. Schedl and C. Ardito. Vol. 13362. Lecture Notes in Computer Science. Springer, 2022, pp. 204–218. DOI: [10.1007/978-3-031-09917-5_14](https://doi.org/10.1007/978-3-031-09917-5_14). URL: https://doi.org/10.1007/978-3-031-09917-5_14 (cit. on p. 13).
- [117] E. Tysinger, M. Pagni, O. Kirchhoffer, F. Mehl, F. Gandon, J.-L. Wolfender and L.-F. Nothias. ‘An Artificial Intelligence Agent for Navigating Knowledge Graph Experimental Metabolomics Data’. In: *2023 Swiss Metabolomics Society Annual Meeting*. Frontiers in Metabolomics. Swiss Metabolomics Society Zurich and ETH Zurich. Zurich, Switzerland, Sept. 2023. URL: <https://hal.science/hal-04381448> (cit. on pp. 13, 34).
- [118] N. Yacoubi Ayadi, S. Bernard, R. Bossy, M. Courtin, B. G. Happi Happi, P. Larmande, F. Michel, C. Nedellec, C. Roussey and C. Faron. ‘A Unified approach to publish semantic annotations of agricultural documents as knowledge graphs’. In: *Smart Agricultural Technology* 8 (Jan. 2024), p. 43. DOI: [10.1016/j.atech.2024.100484](https://doi.org/10.1016/j.atech.2024.100484). URL: <https://hal.science/hal-04495022> (cit. on p. 23).
- [119] A. Yurchyshyna, C. Faron-Zucker, N. L. Thanh and A. Zarli. ‘Knowledge capitalisation and organisation for conformance checking model in construction’. In: *Int. J. Knowl. Eng. Soft Data Paradigms* 2.1 (2010), pp. 15–32. DOI: [10.1504/IJKESDP.2010.030464](https://doi.org/10.1504/IJKESDP.2010.030464). URL: <https://doi.org/10.1504/IJKESDP.2010.030464> (cit. on p. 11).