Fraunhofer FIT

Fraunhofer Institute for Applied Information Technology FIT

Beyond words: Grounding AI Agents in Enterprise Knowledge

Dr. Diego Collarana

Knowledge Centric AI
Fraunhofer FIT | Sankt Augustin | Germany

Fraunhofer FIT – Institut für angewandte Informationstechnik

Research Driven Approach for Innovation



© Markus Jürgens





- Publications
- PhD Thesis
- Lectures

Bundesministerium der Finanzen

Bundesministerium für Wirtschaft und Energie

Bundesministerium für Bildung und Forschung

Bundesministerium für Verkehr und digitale Infrastruktur

DFG

Fraunhofer FIT

Generative AI Research & Innovation Programs at Fraunhofer FIT

Overview

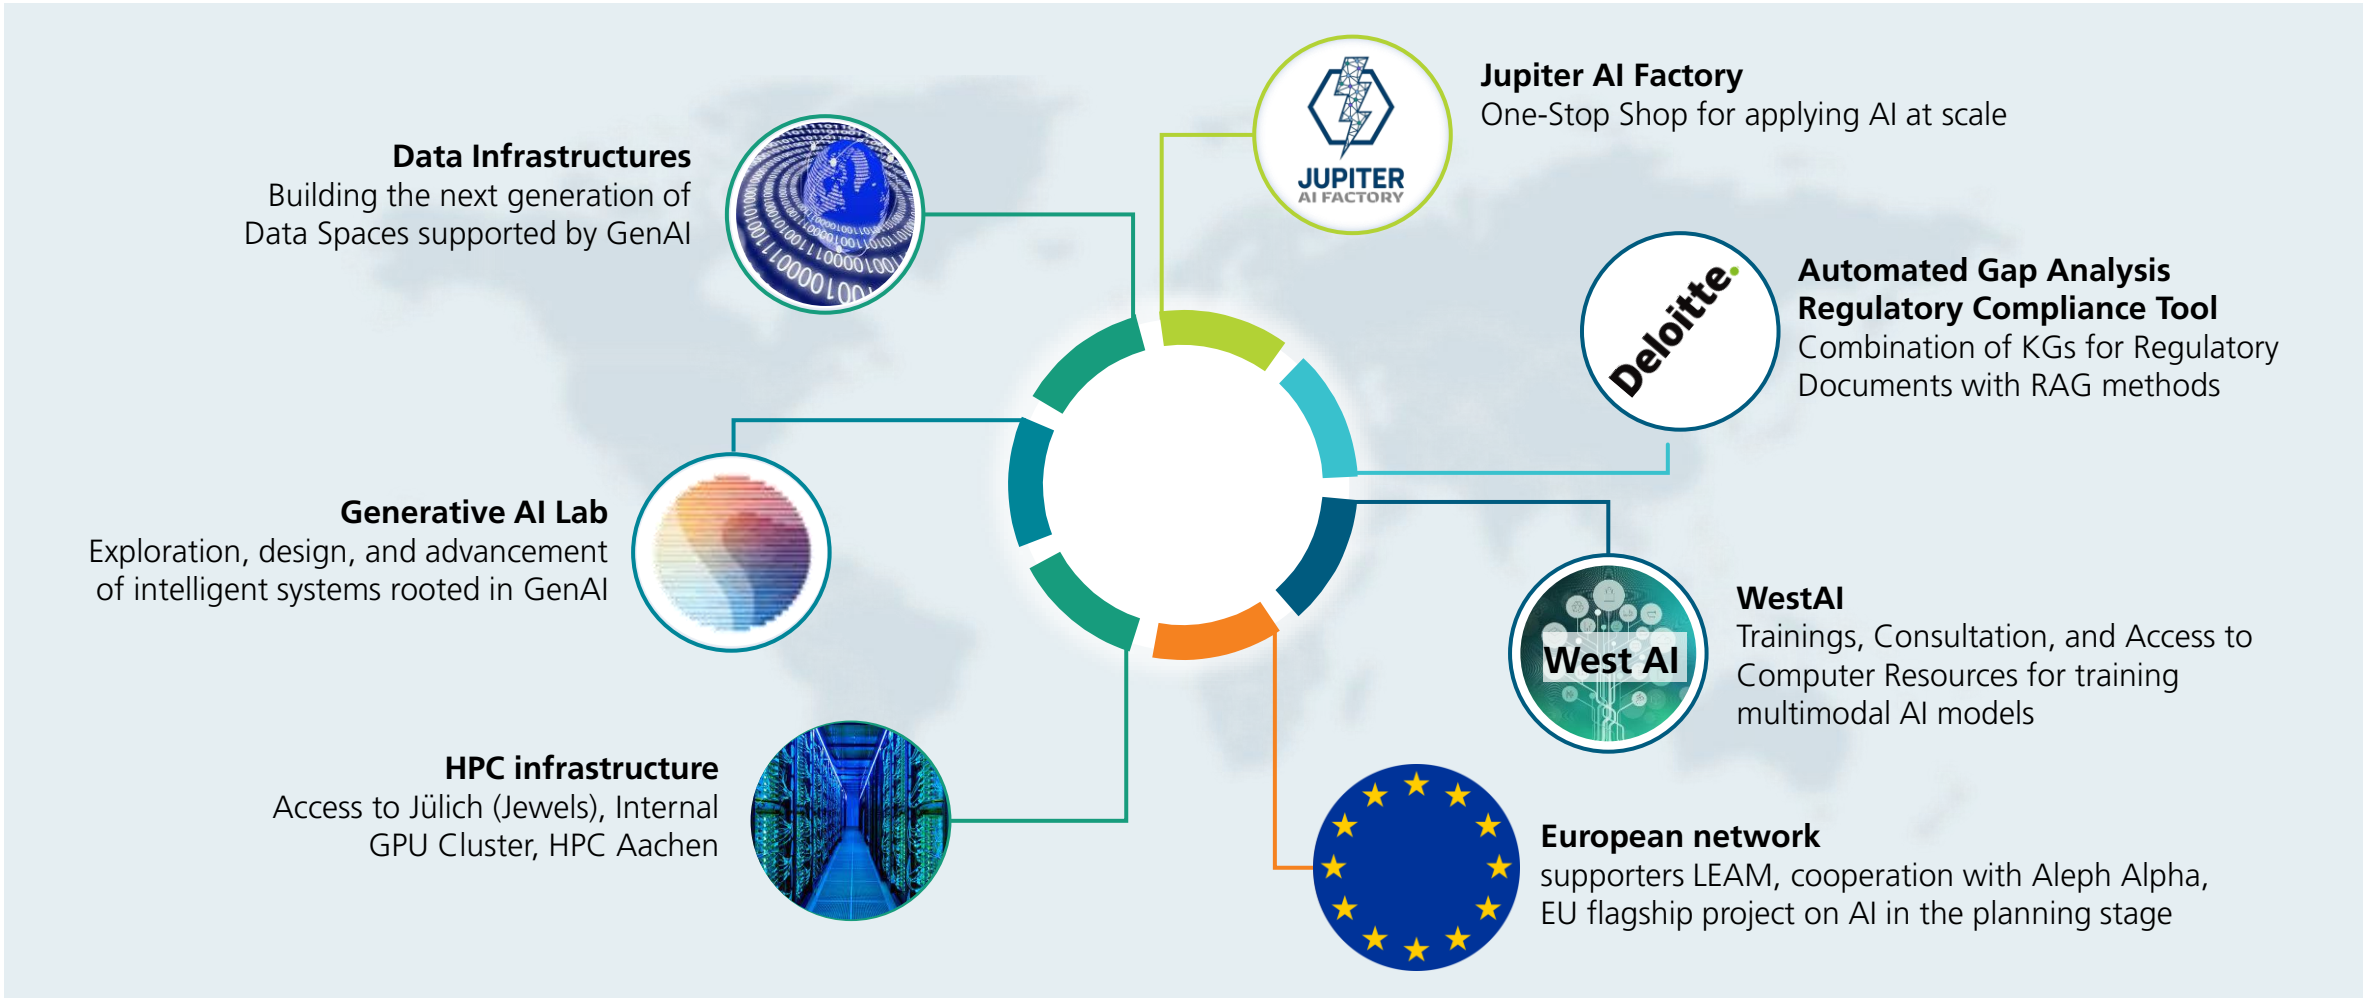

Fraunhofer FIT

Jupiter AI Factory

One-Stop Shop for applying AI at scale

- Enable world-class research on **domain-specific** AI solutions
- Provide **AI knowledge and technology transfer** and support to regional, national, and European stakeholders
- Develop deep **expertise** and highly **skilled talents** for industry and public services
- Provide cloud-based scalable **inference resources** via the dedicated JARVIS module
- Enable access to the **largest data resources at JSC** and to the Common **European Data Spaces**

https://www.jaif.ai/

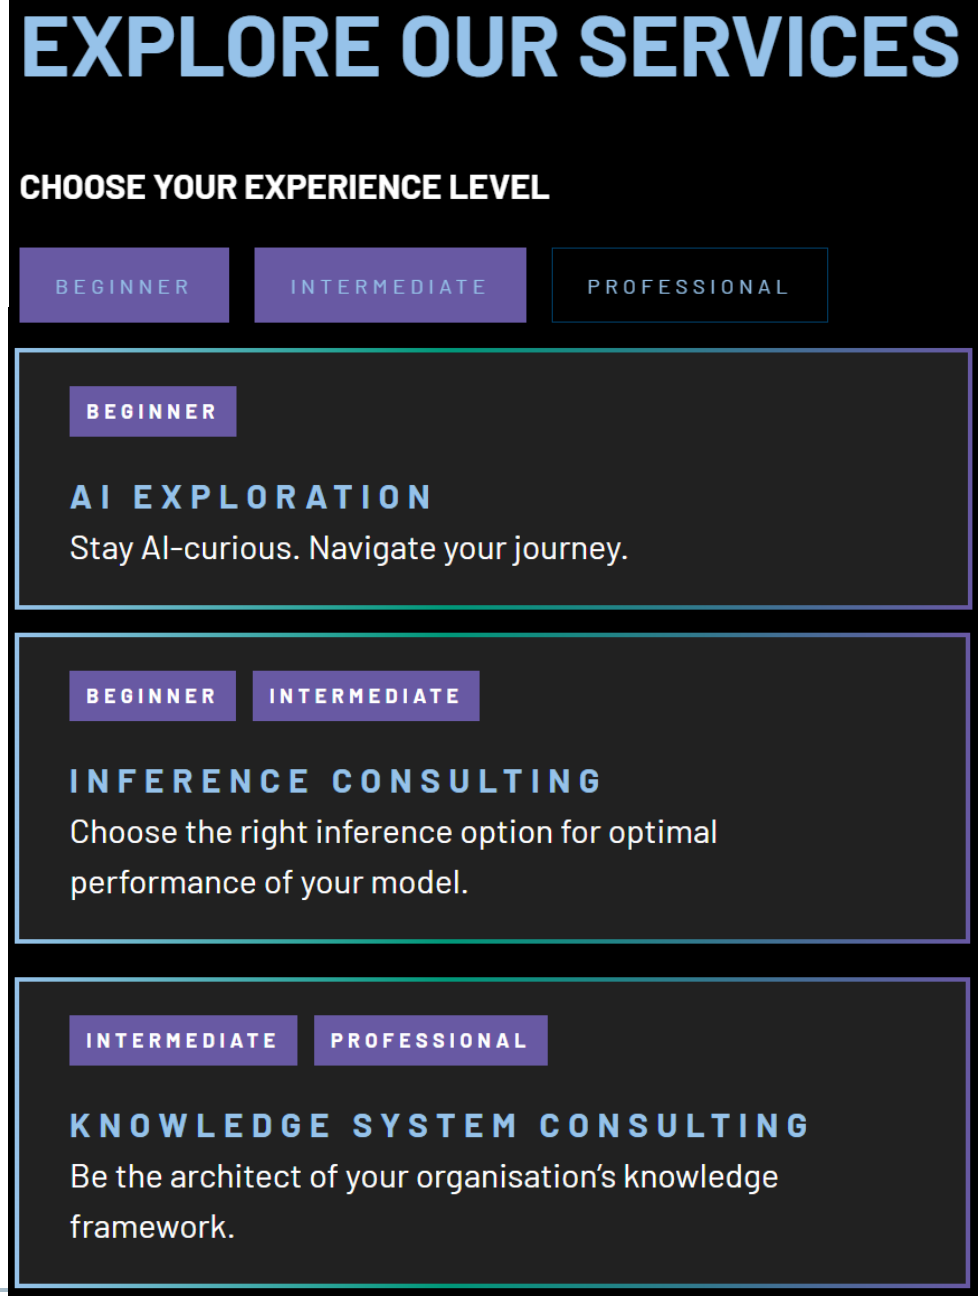

Fraunhofer FIT

Agentic AI is at the heartbeat of Industry. Everyone is looking how to build an Agentic solution to automate processes, increase productivity and reduce costs. However,…

 © Fraunhofer FIT

Fraunhofer FIT

Agentic AI is at the hearth bit of Industry. Everyone is looking how to build an Agentic solution to automate processes, increase productivity and reduce costs. However,…

Q: Find approved purchase orders (POs) for the SE workforce that use preferred suppliers

MIT Study: 95% of AI projects fail

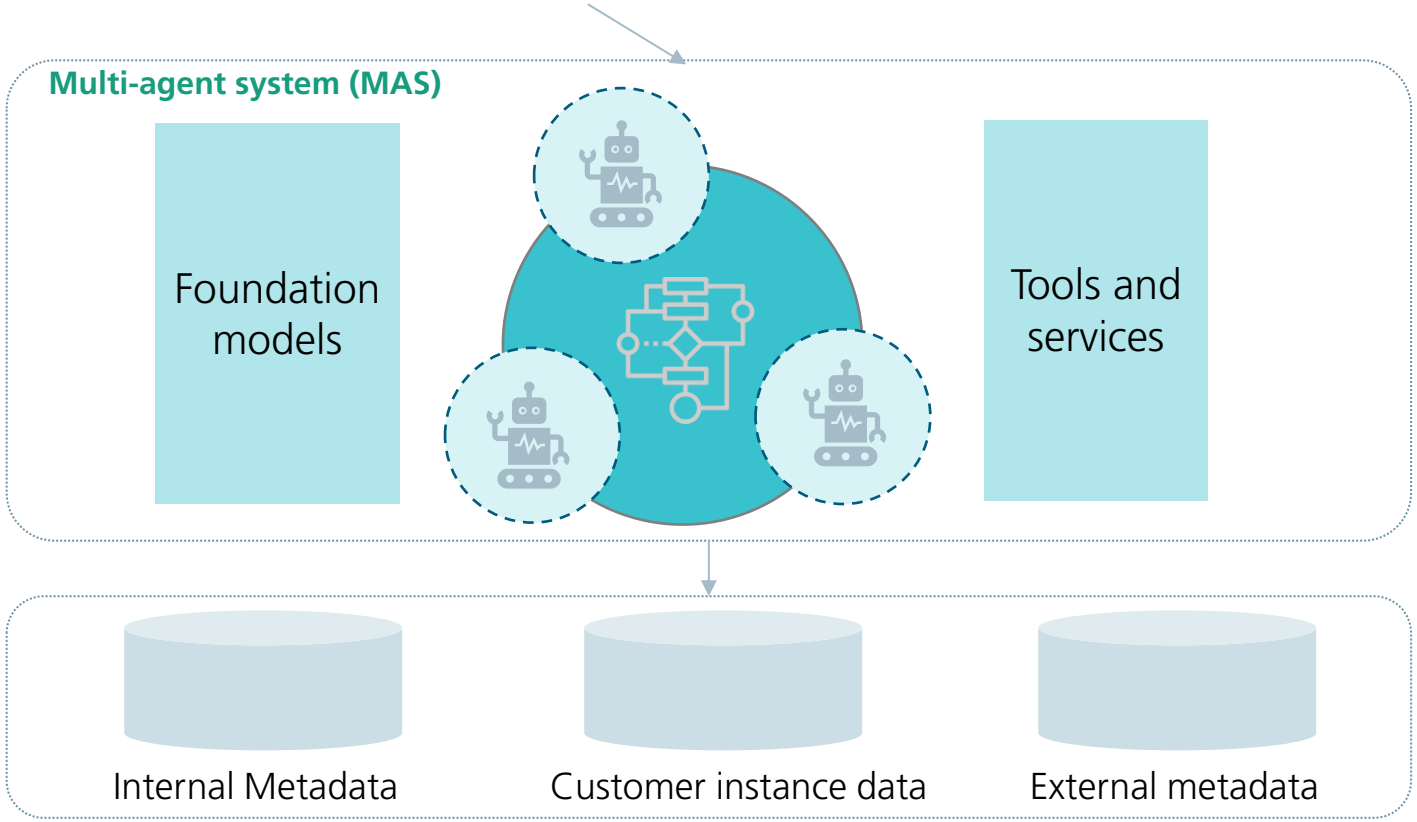

Fraunhofer FIT

Agentic AI is at the hearth bit of Industry. Everyone is looking how to build an Agentic solution to automate processes, increase productivity and reduce costs. However,…

The challenge is not “reasoning” but grounding

Q: Find approved POs for SE workforce that use preferred suppliers

AI-Assistant needs to: (1) Discover the right sources by matching the terms; (2) Understand by mapping terms to a shared world model; and (3) Act efficiently by multi-sourcing querying and integrating knowledge from cross-systems

Slide inspired by the keynote at IJCAI NORA Workshop 2026, by Felix Sasaki

Neuro-Symbolic Backbone for Accessing and Using Enterprise Data

From static repositories towards autonomous actionable knowledge a hybrid approach

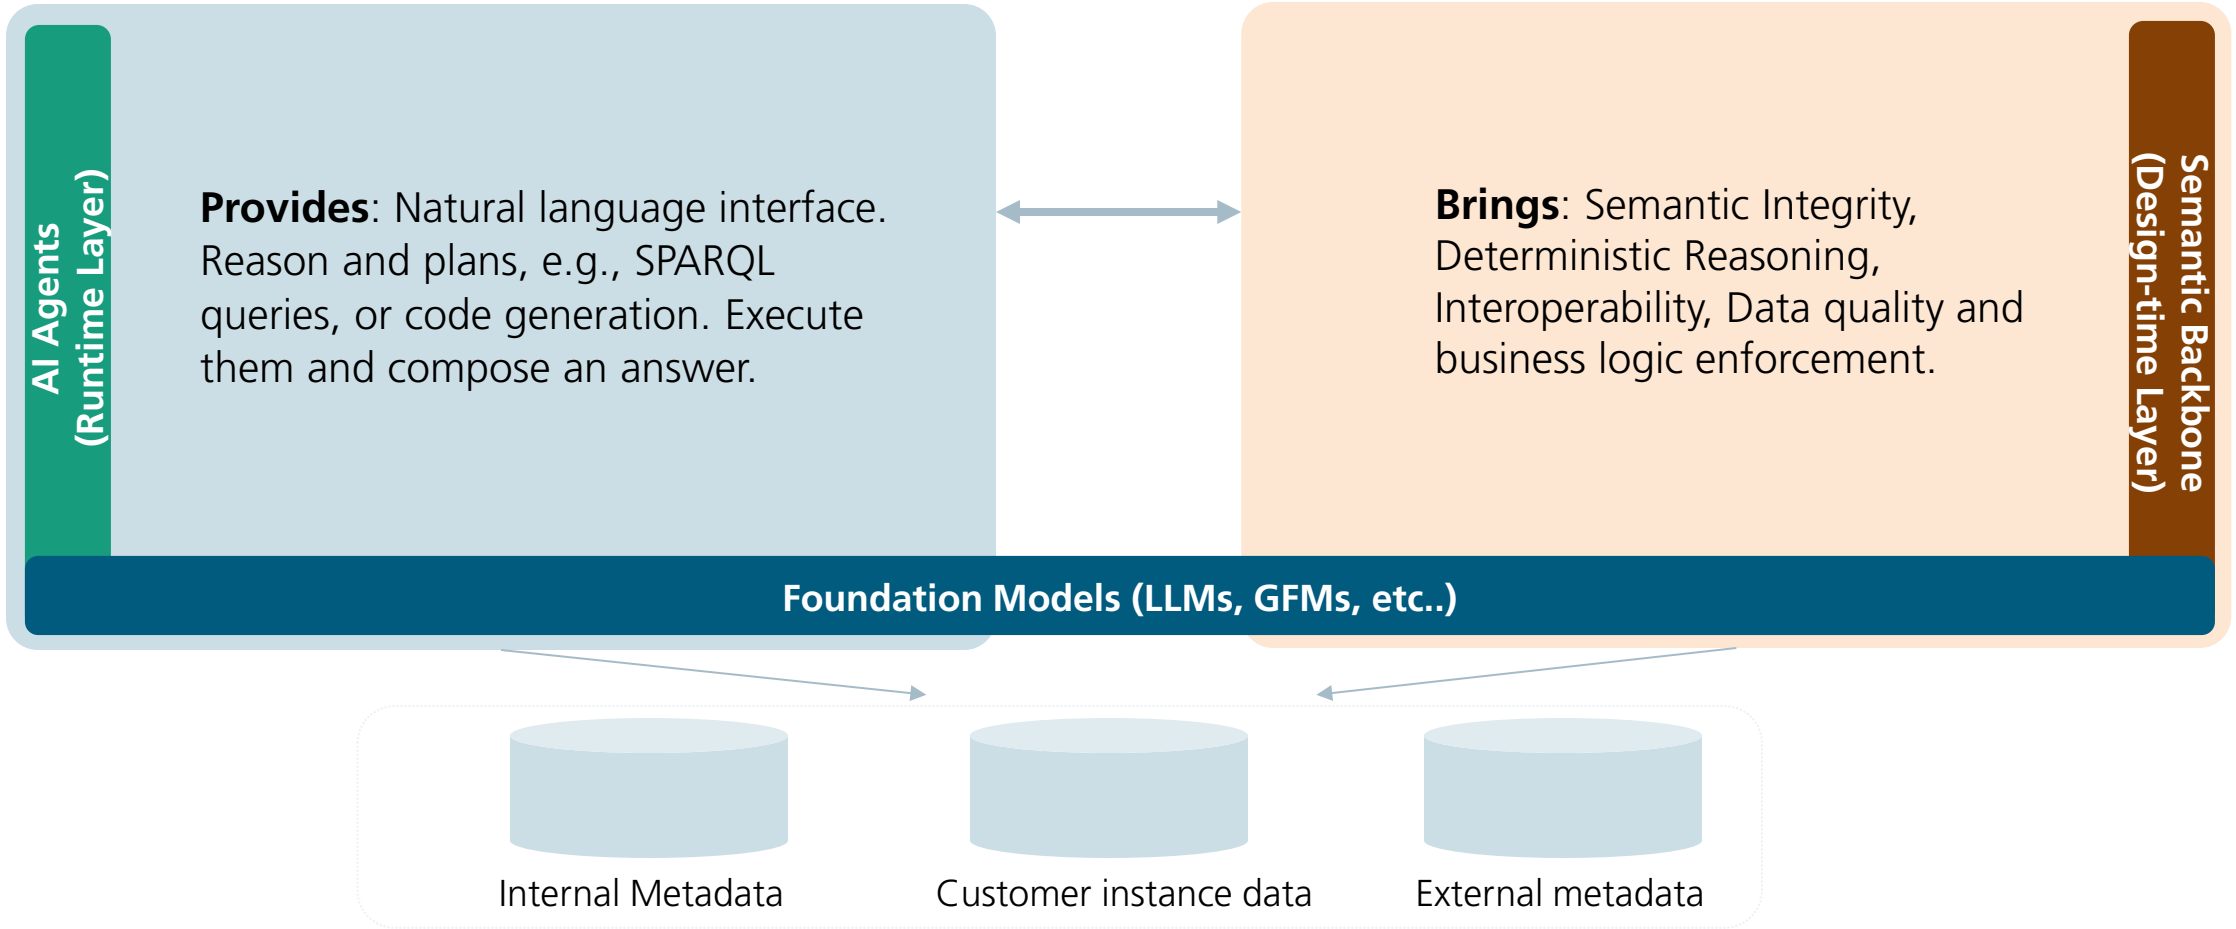

Neuro-Symbolic Backbone for Accessing and Using Enterprise Data

From static repositories towards autonomous actionable knowledge a hybrid approach

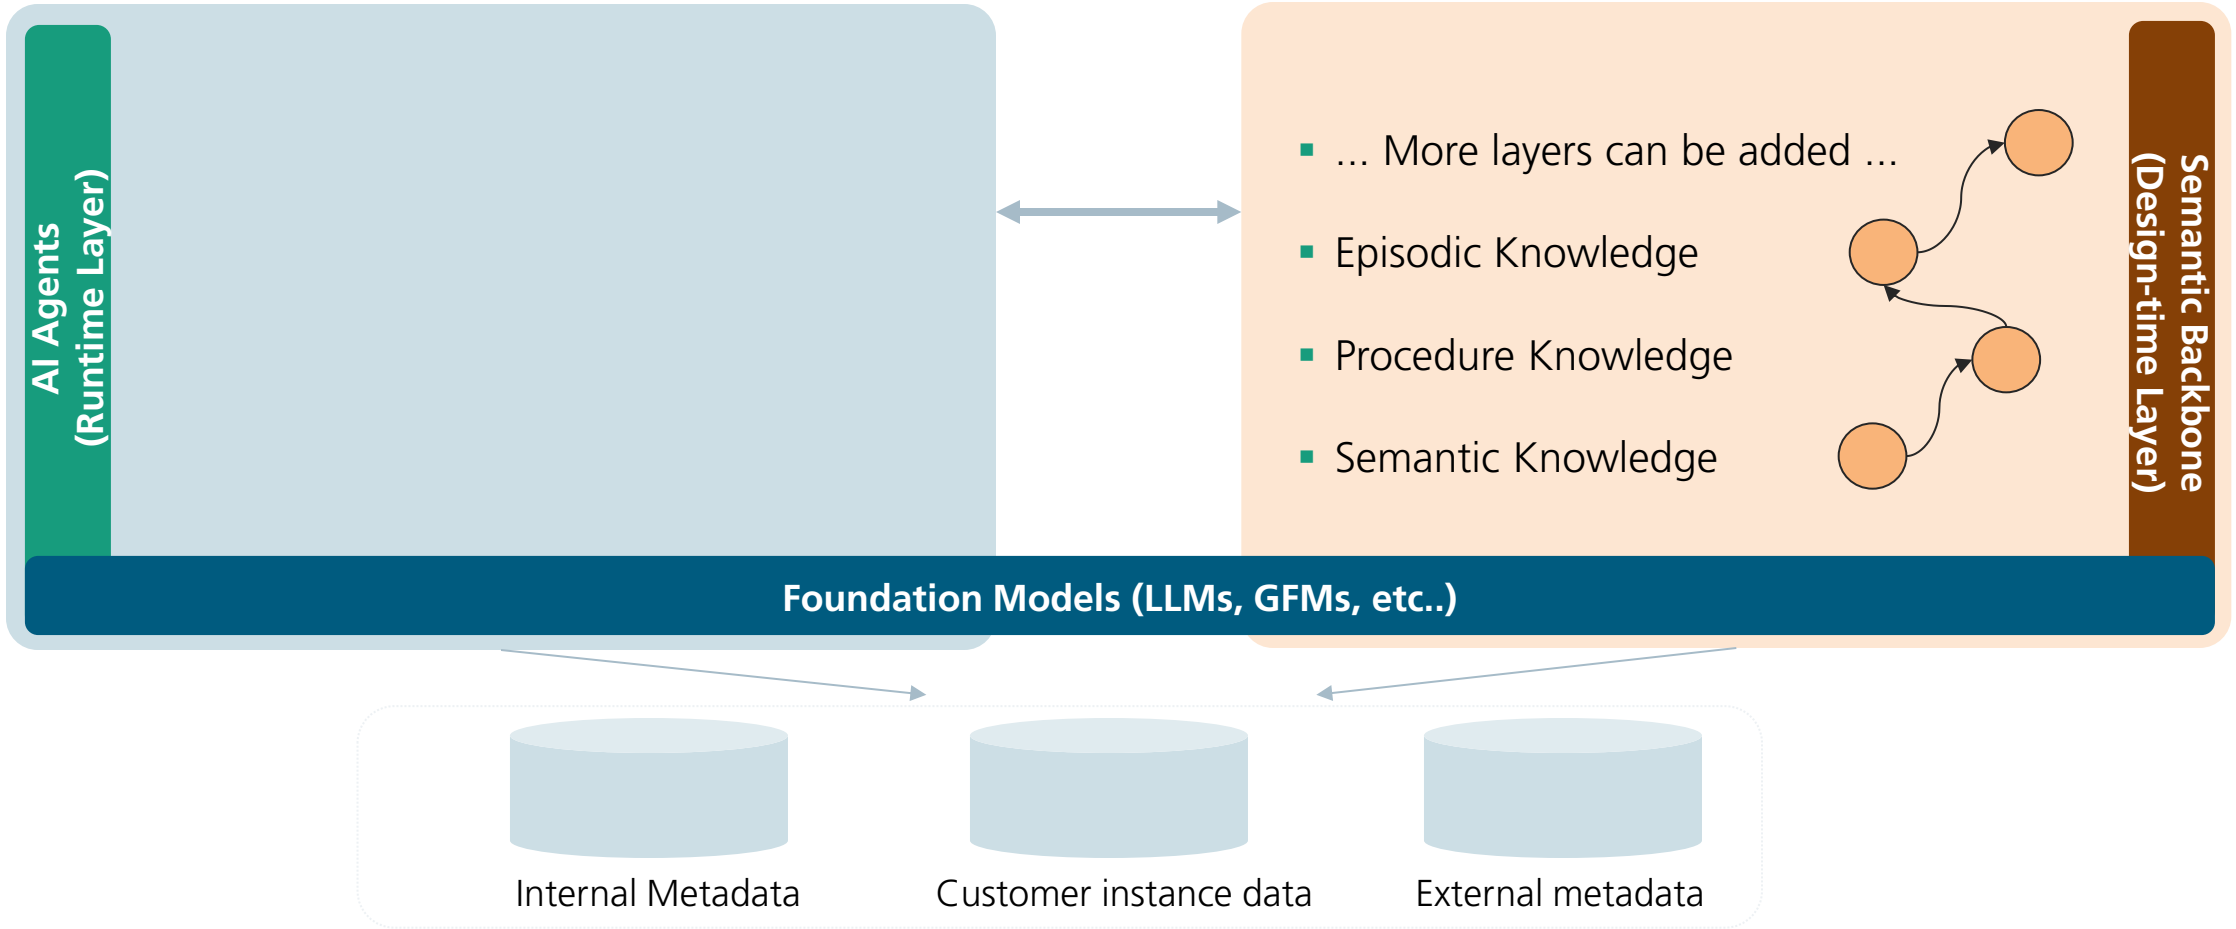

Neuro-Symbolic Backbone for Accessing and Using Enterprise Data

From static repositories towards autonomous actionable knowledge a hybrid approach

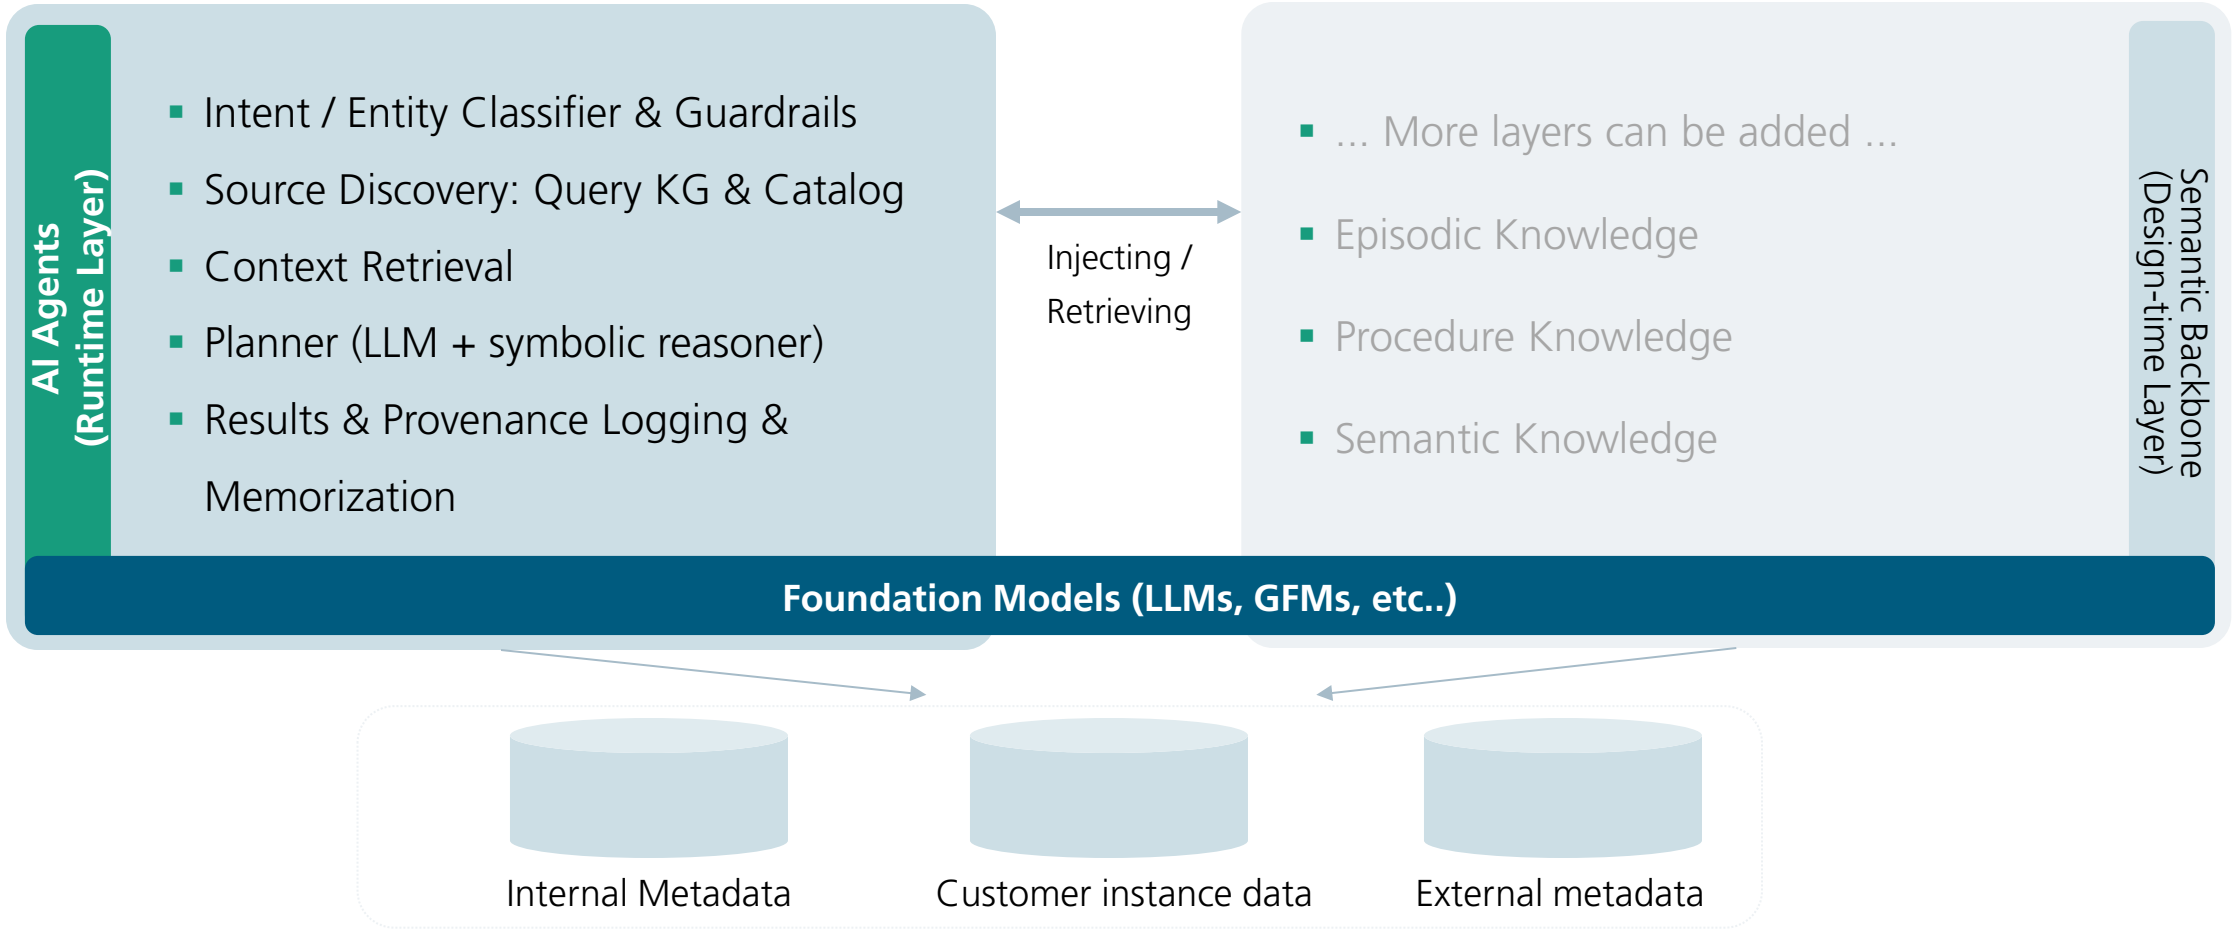

Neuro-Symbolic Backbone for Accessing and Using Enterprise Data

From static repositories towards autonomous actionable knowledge a hybrid approach

Brings: Semantic Integrity, i.e., meaning across silos and precise retrieval.
Deterministic Reasoning, i.e., data quality and business logic enforcement.
Interoperable knowledge for reuse across silos and value chains.

Research Gaps:

1. Automation Gap
   - Translate Questions to SPARQL
   - Translate Text (Policies) into RDF following an Ontology
   - Universal Retriever
2. Universal Temporal KG embeddings
   - Graph Foundation Models → THGFM: Dual-Branch Temporal Heterogeneous Graph Fusion Model at ISWC'26

Semantic Backbone (Design-time Layer)

- SHACL Constrains
- Embedding Index (KG embeddings, textual)
- Knowledge Graph (RDF Stores)
- Instance Matching (Entity Resolution)
- Ontology (OWL Classes/Properties)
- Schema Alignment & Mappings
- Triple Extraction (schema/column names, docs)



Neuro-Symbolic Backbone for Accessing and Using Enterprise Data

From static repositories towards autonomous actionable knowledge a hybrid approach

AI Agents (Runtime Layer)

- Intent / Entity Classifier & Guardrails
- Source Discovery: Query KG & Catalog
- Context Retrieval (table samples, docs, ontologies)
- Planner (LLM + symbolic reasoner)
- Execution Engine (SQL/SPARQL endpoints, APIs, ...)
- Results & Provenance Logging & Memorization
- Answer Synthesis & Analysis & Explanation (LLM)

Provides: Seamless communication with user, i.e., natural language. Finds relevant sources and generate plans, e.g., executable SPARQL queries, or code generation. Execute the plan and compose an answer.

Relevant Research Gaps:

1. Long-term memory
   - Ontology defining how to store knowledge
2. Detect undesirable behavior, e.g., hallucinated answers
   - Steering vectors calculated out of LLMs residual stream



Fraunhofer FIT

Uses Cases and Applications

Research Driven Approach for Innovation

1. Automated Generation and Validation of ODRL Policies

Neurosymbolic Policy Reasoning for ODRL

RQ: How can AI agents collaboratively interpret natural-language policies, map them into ODRL compliance schemas, and detect logical or regulatory contradictions across domains?

Natural-language policy

Acme Corporation permits Research Lab to read the Climate Dataset solely for academic research before 1 January 2027 at 00:00 UTC. Research Lab is prohibited from distributing the dataset.

ODRL representation in RDF (Turtle)

```
@prefix odrl: <http://www.w3.org/ns/odrl/2/> .
@prefix xsd:  <http://www.w3.org/2001/XMLSchema#> .
@prefix ex:   <https://example.org/> .

ex:ClimateDataAgreement
   a odrl:Agreement ;

   odrl:permission [
      odrl:assigner ex:AcmeCorporation ;
      odrl:assignee ex:ResearchLab ;
      odrl:target ex:ClimateDataset ;
      odrl:action odrl:read ;

      odrl:constraint [
         odrl:leftOperand odrl:purpose ;
         odrl:operator odrl:eq ;
         odrl:rightOperand ex:AcademicResearch
      ] ;
      odrl:constraint [
         odrl:leftOperand odrl:dateTime ;
         odrl:operator odrl:lt ;
         odrl:rightOperand
            "2027-01-01T00:00:00Z"^^xsd:dateTime
      ]
   ] ;  odrl:prohibition [
      odrl:assigner ex:AcmeCorporation ;
      odrl:assignee ex:ResearchLab ;
      odrl:target ex:ClimateDataset ;
      odrl:action odrl:distribute
   ] .
```

Publications:

[1] What Does ODRL Mean? A Cross-Level Ontological Grounding of Permissions, Prohibitions, and Duties in UFO-L
[2] Axis Decomposition for ODRL: Resolving Dimensional Ambiguity in Policy Constraints through Interval Semantics
[3] From Instructions to ODRL Usage Policies: An Ontology Guided Approach
[4] Denotational Semantics for ODRL: Knowledge-Based Constraint Conflict Detection

From Instructions to ODRL Usage Policies: An Ontology Guided Approach

Mustafa et al. @ VDBL on LLMs + KGs

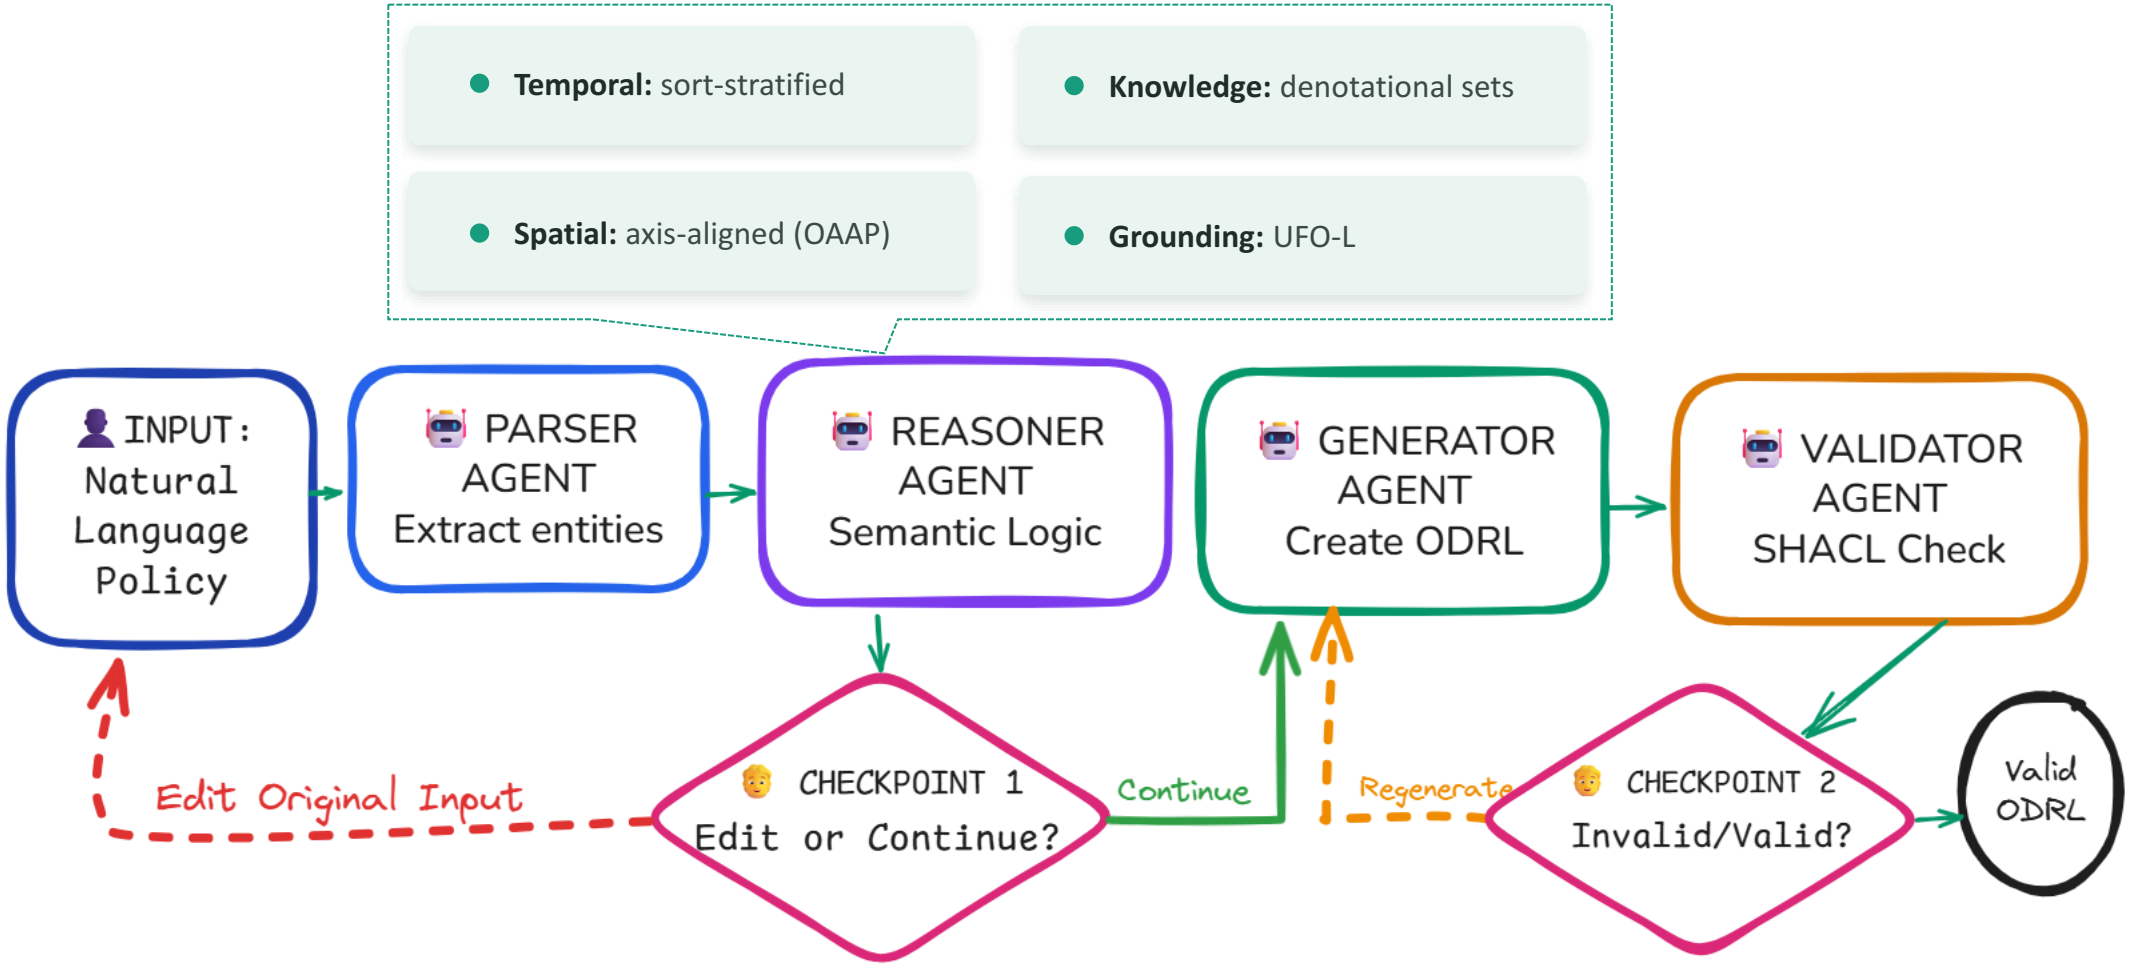

Demo

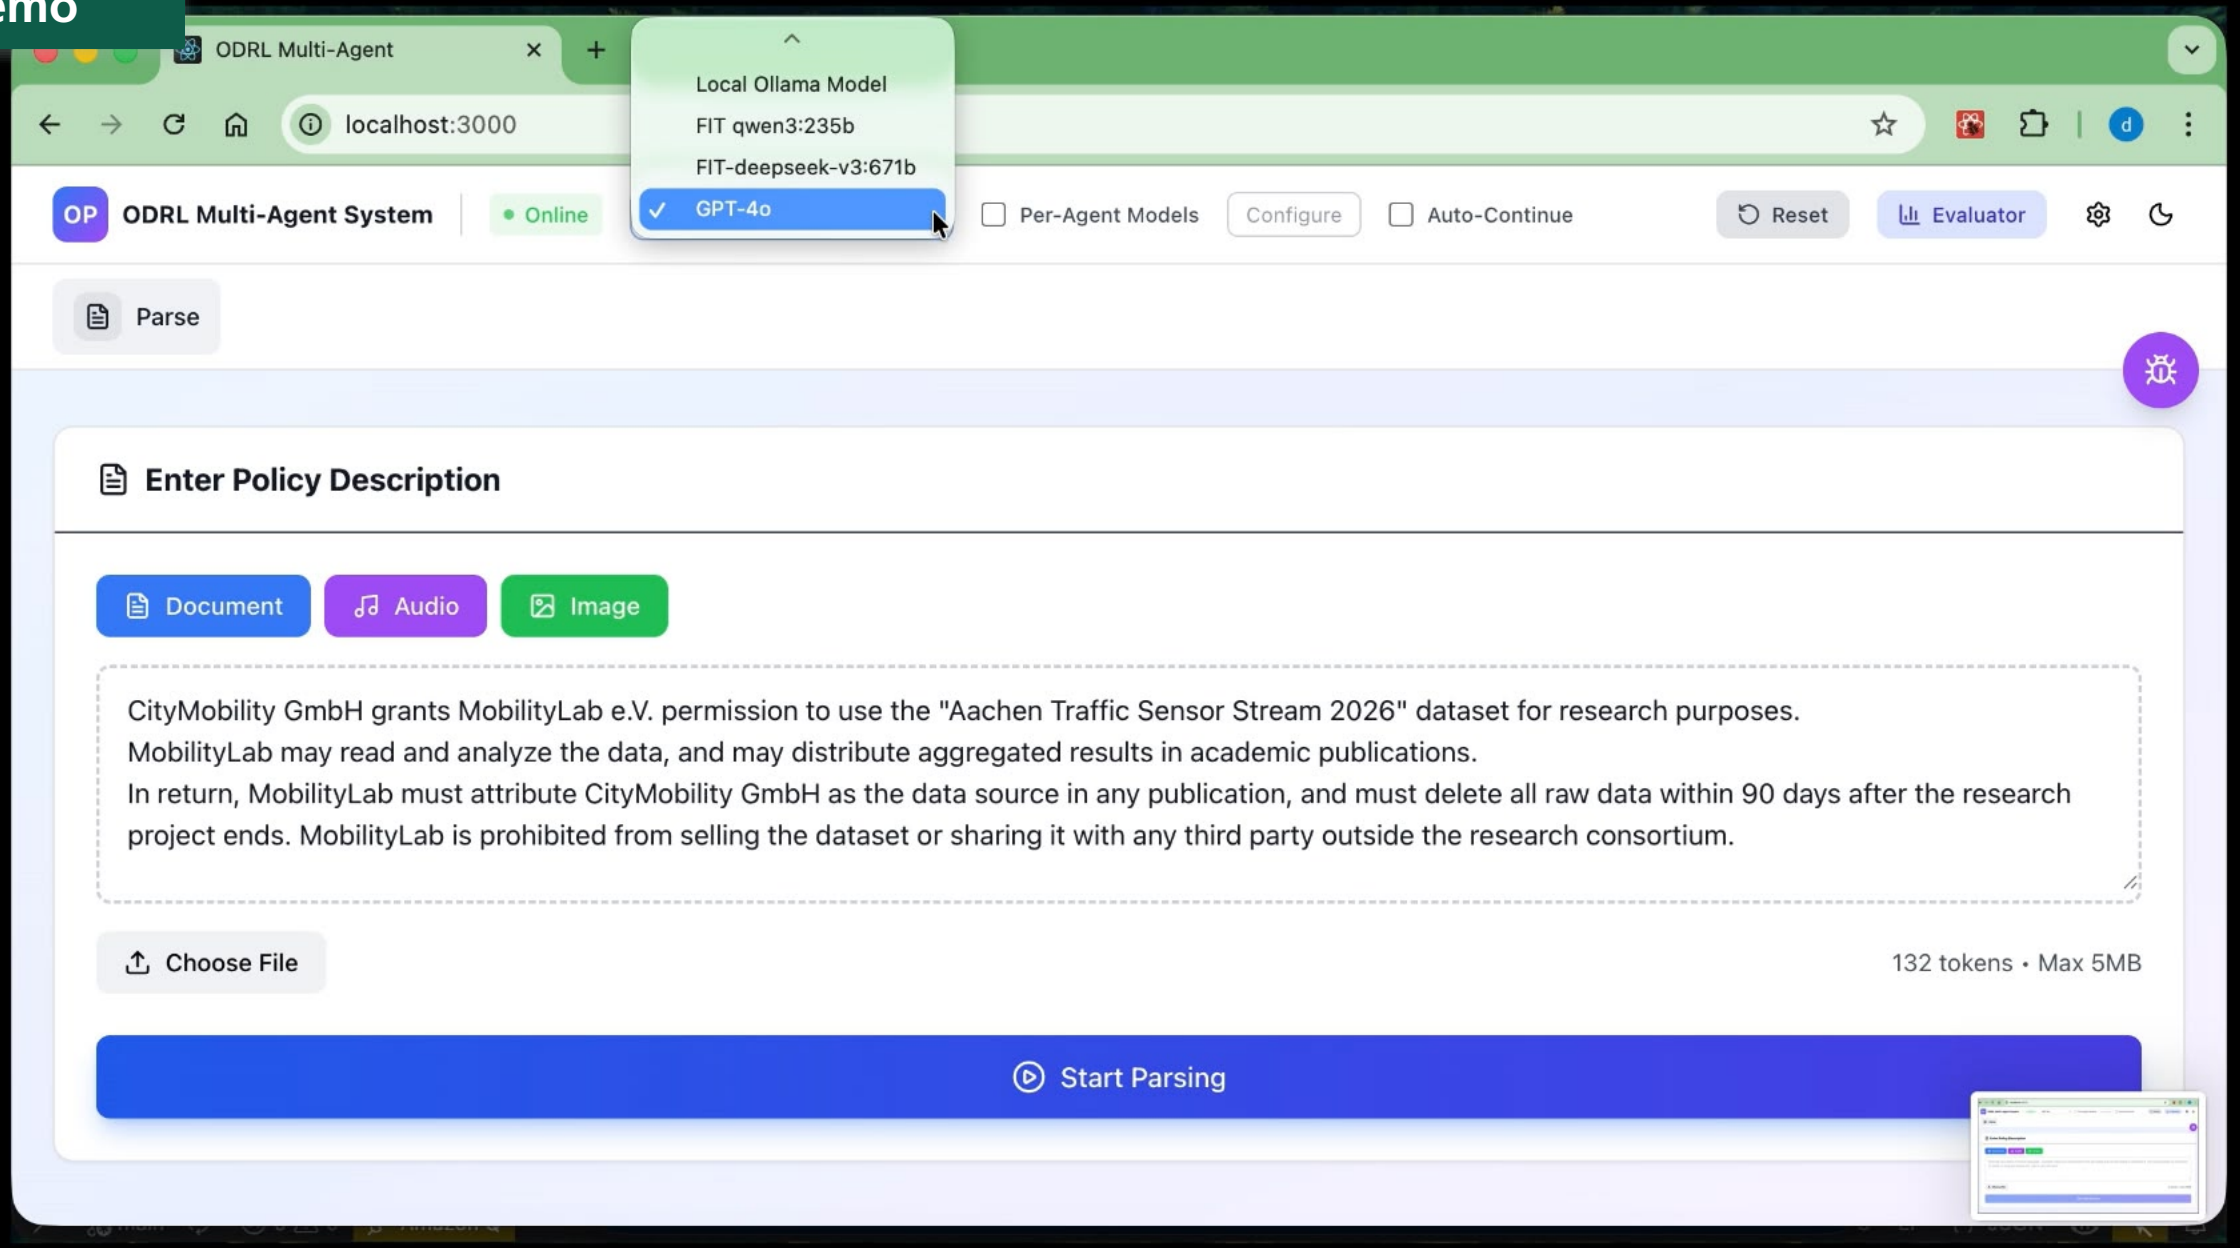
ODRL Multi-Agent
localhost:3000
Local Ollama Model
FIT qwen3:235b
FIT-deepseek-v3:671b
GPT-4o
OP
ODRL Multi-Agent System
Online
Per-Agent Models
Configure
Auto-Continue
Reset
Evaluator
Parse
Enter Policy Description
Document
Audio
Image
CityMobility GmbH grants MobilityLab e.V. permission to use the "Aachen Traffic Sensor Stream 2026" dataset for research purposes.
MobilityLab may read and analyze the data, and may distribute aggregated results in academic publications.
In return, MobilityLab must attribute CityMobility GmbH as the data source in any publication, and must delete all raw data within 90 days after the research project ends. MobilityLab is prohibited from selling the dataset or sharing it with any third party outside the research consortium.
Choose File
132 tokens • Max 5MB
Start Parsing

Long-term Memory Gap

2. Data Quality and Analytics Services for Data Spaces

RQ: How can AI agents retain and reuse insights from previous advanced analytics tasks, generated code, and identified localized data quality issues to improve the accuracy, efficiency, and reliability of future analytics tasks?

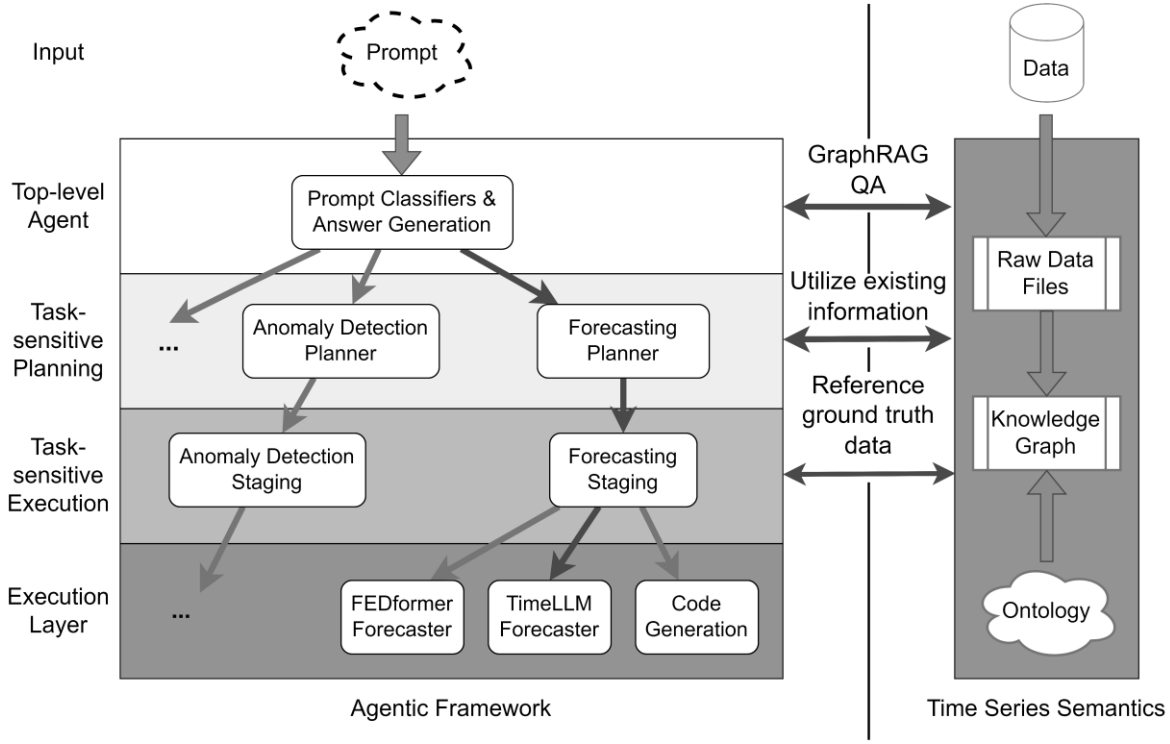

Publications:

[1] SemTS: Ontology and Vocabularies for the Semantic Categorization of Time Series Knowledge
[2] Code2Onto: Multi-Agent System for Code-Driven Ontology Population
[3] Semantic Intelligence: Graph RAG-Driven Agents for Time Series Analytics

Fraunhofer FIT

Long-term Memory Gap

Ontology and Vocabularies for Semantic Time Series Knowledge

Graß et. al. @ ESWC 2026, resource track

Gap: Existing ontologies focus on dataset definitions, catalogs, and methodical information. Not for persisting insights at the end of data pipelines.

Contribution: An Ontology representation of reusable knowledge out of data pipelines (focus on Time Series but extensible).

- Data knowledge: Insights derived from data
- Scenario Knowledge: Domain expertise / Validated knowledge
- Method Knowledge: Best practice settings (e.g., Models, Hyperparameters)

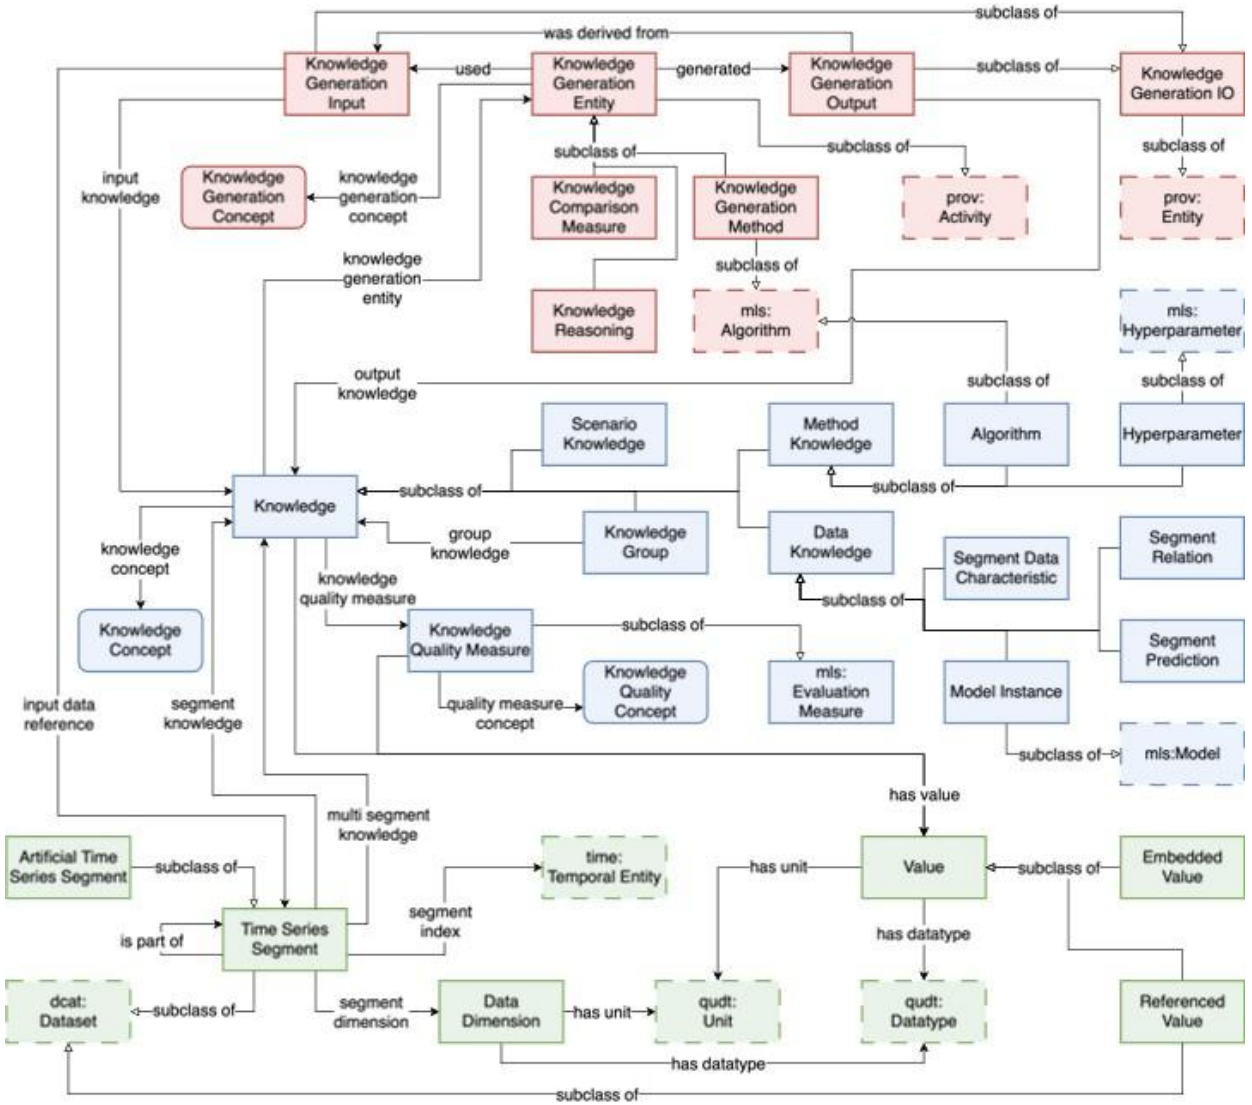

Fraunhofer FIT

Demo

SemTS Agent

Analysis | Knowledge Graph | Session: c6031e7e...

Ask a Question

Examples

New Conversation 0 messages in conversation

Sh

Submit | Clear Input | New Chat

Tip: Ask follow-up questions like 'show me option B' or 'explain that in more detail'

Response

New conversation started! Ask your first question.

Visualization

Visualization will appear here after analysis...

Conversation History

Clear

No conversation history yet.

Execution Trace

Expand All

Trace will appear here...

Knowledge Used

Retrieved knowledge will appear here...

Retriever Gap

3. AI-Powered EU Regulatory Compliance Assistant

RQ: How to integrate fragmented regulatory sources to provide up-to-date, context-aware interpretations of EU regulations for identify potential compliance risks delivering transparent, evidence-based justifications for its recommendations?



Publications:
[1] Graph RAG Best Practices: A Modular Evaluation Framework
[2] GRAFT - Graph Retrieval Augmented Generation Fine-Tuning Approach
[3] ABEL: Artificial Buddy for Effective Learning. IJCKG 2025: 329-338
[4] Graph RAG in the Wild: Insights and BestPractices from Real-World Applications

Fraunhofer FIT

3. AI-Powered EU Regulatory Compliance Assistant

GraphRAG Best Practices: A Modular Evaluation Framework…

Gap: Most practitioner pipeline design choices in GraphRAG are ad hoc. Which chunking? Retriever? Traversal depth? reranker? Prompt serialization? etc.

Contribution: We broke the GraphRAG pipeline into its three core stages — indexing, retrieval, and generation, develop an evaluation framework and benchmarked:

- 17 embedding models
- 6 retriever types across multiple traversal depths
- 9 reranking models
- 4 graph serialization formats

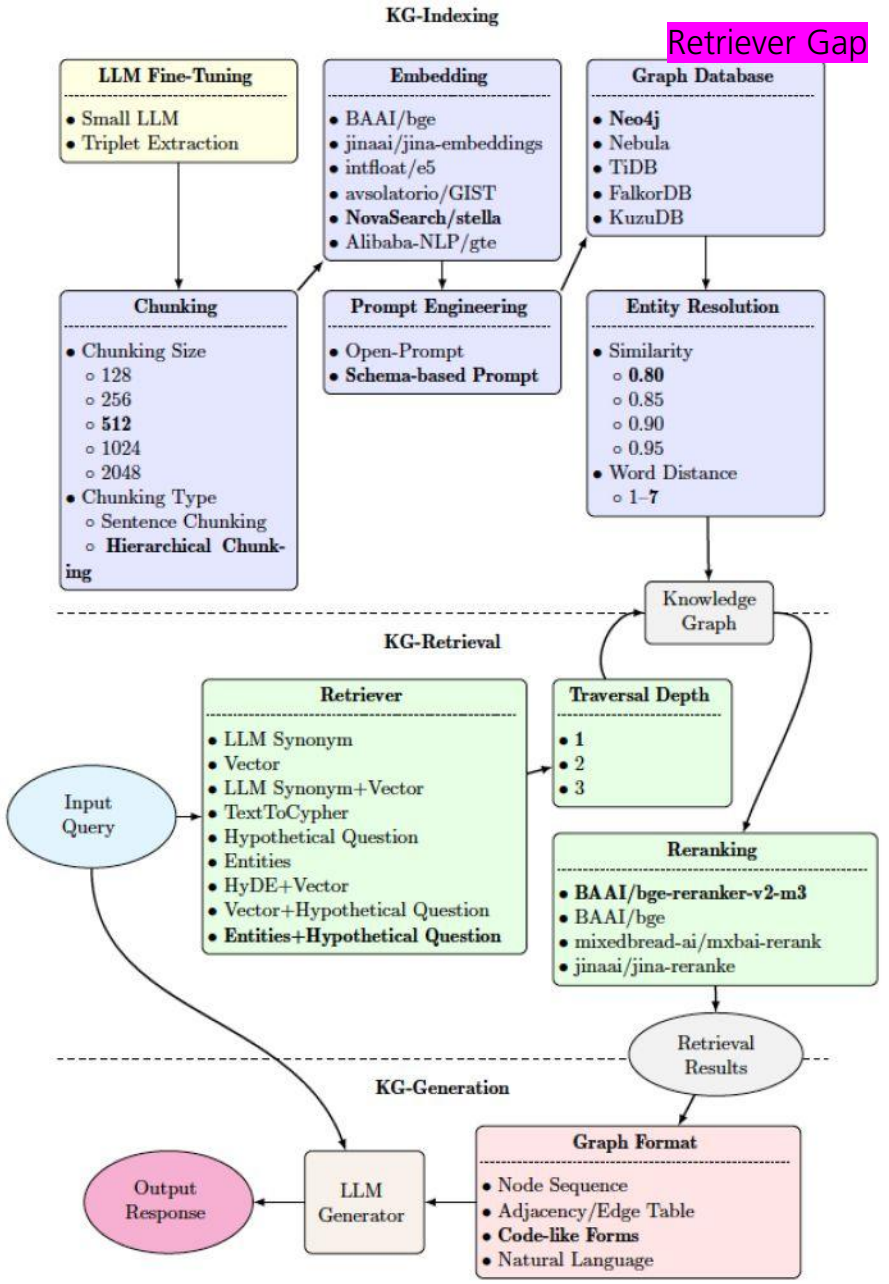

Fraunhofer FIT

Demo



Main conclusion and takeaways

Why LLM Agents Need a Semantic Layer

Grounding AI Agents in enterprise knowledge requires neuro-symbolic approaches

1. **Shared meaning, not model-dependent interpretation**
   Knowledge graphs provide an explicit semantic contract for business entities, processes, and rules—grounding agents in agreed enterprise definitions.

2. **Connected knowledge, memory, and experience**
   The semantic layer links enterprise knowledge with agents past insights, actions, and outcomes, enabling contextual reasoning and reuse of experience.

3. **Traceable and accountable actions**
   Services built on the semantic layer can validate proposed actions, trace decisions to their knowledge sources, and support auditing against enterprise policies.

4. **Model independence and technological sovereignty**
   Keeping enterprise knowledge separate from the LLM reduces vendor dependence and supports switching models or deploying local models—strengthening organizational and national control over data and AI infrastructure.

Vielen Dank für Ihre Aufmerksamkeit

Fraunhofer FIT

Fraunhofer Institute for Applied Information Technology FIT

Diego Collarana
Head of Knowledge Centric AI at Fraunhofer FIT