OLYMPIA AI DATA INFRASTRUCTURE SUMMIT 2026

Day 1 Recap

Olympia AI Data Infrastructure Summit 2026 · Heraklion, Wednesday 4 November 2026

Yannis Tzitzikas

FORTH-ICS and University of Crete

Day 1 at a glance

KEYNOTE

Joseph Sifakis

Trustworthy autonomous agents



Turing Award 2007, and a son of Heraklion

SESSION 1

Foundations of knowledge and agentic AI

Jeff Z. Pan

From parameter scaling to agent scaling

Martin Doerr

Knowledge about the past

Bo Zhao

State management for agentic AI

SESSION 2

Semantic knowledge engineering for AI agents

Diego Collarana

Grounding agents in enterprise knowledge

Nikolaos Bassiliades

Language models meet ontology engineering

Claudia d'Amato

Neural-symbolic link prediction

SESSION 3

From semantic extraction to agentic systems

Arijit Khan

Graph-native LLMs and the AI harness

Grigoris Tsagkatakis

Agentic frameworks for space data

Georgios Paliouras

The BioASQ challenge

THREE THREADS

1. **Agents need more than a model**
2. **Semantics as the backbone of knowledge**
3. **From extraction to operational agents**

Agents need more than a model

Scale, state, memory and verification around the model

Joseph Sifakis
Keynote: Trustworthy Autonomous Agents

- Autonomy is defined by functionality, not by a technique
- Trust needs rigorous design and validation, beyond today's AI. Symbolic AI is needed

Jeff Z. Pan
From Parameter Scaling to Agent Scaling

- From scaling model parameters to scaling cooperating agents
- Knowledge and memory as the glue between agents

Bo Zhao
State Management for Agentic AI Systems

- Evolving context and verification results create a "state wall"
- State abstractions for scheduling, training and experience data

Arijit Khan
Graph-native LLMs and Agentic Systems

- The harness (retrieval, tools, memory) largely decides reliability
- KG-RAG for matching, alignment and fact checking

Semantics as the backbone of knowledge

Neural and symbolic methods are complementary, and the arrows go both ways

Martin Doerr
Challenges of Knowledge about the Past

- Knowledge about the past rests on evidence and interpretation
- Event-centric models with provenance (CIDOC CRM) make it integrable. LLMs make it possible to create Very Large Knowledge Bases

Diego Collarana
Grounding AI Agents in Enterprise Knowledge

- KGs give agents a world model: semantics, IDs, provenance, governance
- Agents can make costly KG curation continuous

Nikolaos Bassiliades
Language Models and Semantic Knowledge

- LLMs draft and match ontologies; reasoning validates and explains
- Ontological constraints steer KG completion

Claudia d'Amato
Neural-Symbolic Link Prediction

- Embeddings scale but ignore schema; benchmarks drop it
- Background knowledge and a protocol to evaluate explanations

From semantic extraction to operational agents

Georgios Paliouras

The BioASQ challenge

10+
years

180+
teams

41
countries

- Shared benchmarks for semantic indexing, question answering and extraction
- The US National Library of Medicine moved to fully automated MeSH indexing
- LLM-based systems are strong, yet challenges remain open

Grigoris Tsagkatakis

Agentic Frameworks for Space Data Management

1. **Semantic access**
   to large, heterogeneous space and Earth-observation data
2. **Agent memory**
   to keep context across long analysis tasks
3. **Operational services**
   scaling from prototypes to operations

Real domains, from biomedicine to space data, are where semantic access, agent memory and evaluation are put to the test.

Five cross-cutting messages

Trust is engineered, not assumed

Rigorous design, verification and evaluation around the model.

Sifakis · Doerr, Zhao · Khan

State and memory are infrastructure

Context, experience and reasoning traces need managed storage

Pan · Zhao · Tsagkatakis

Neural and symbolic, in both directions

LLMs build knowledge; ontologies guide, validate and explain

Sifakis, Pan, Doerr, d'Amato · Bassiliades · Collarana

Evaluation is a research problem itself

Shared benchmarks, explanation protocols, locating misalignment

Paliouras · d'Amato · Khan

Knowledge lifecycles become continuous

Provenance-aware, agent-assisted curation keeps knowledge current

Doerr · Collarana · Bassiliades

Questions to take into Day 2

- What is worth **remembering** for an enterprise agent, and at what cost?
- **Who builds and maintain**s the enterprise ontology: people, agents, or both?
- Which **guarantees** can we realistically offer for agents built on LLMs?
- Which shared **benchmarks** would move the field, and could we build them together?

Tomorrow

09:00 Keynotes: Diego Calvanese, Yannis Ioannidis

10:40 Adaptive and decentralised intelligence

13:30 Operational AI and structural intelligence

15:20 Panel: Ontology and Agent Memory

16:50 Recap, joint statement and closing

Thank you for a great first day! See you at Petousis Taverna at 18:45.

Petousis Restaurant: 8-9 min walk

