OLYMPIA AI DATA INFRASTRUCTURE SUMMIT 2026

PANEL

Ontology and Agent Memory

Turning Research into Product Capabilities

Moderator: Yannis Tzitzikas

FORTH-ICS and University of Crete

Our panelists

Yannis Ioannidis
University of Athens

Data infrastructure and data management at scale

Diego Calvanese
Free University of Bozen-Bolzano

Reasoning, data access and integration

Dimitris Plexousakis
University of Crete and FORTH-ICS

Knowledge representation, evolution and reasoning

Arijit Khan
Aalborg University

Data systems and scalable knowledge processing

Claudia d'Amato
University of Bari

Learning, inference and knowledge graphs

How we work: a brief answer, then others **challenge or build on it**. In question 4, everyone names **one capability and one breakthrough**.

The central question

How can ontologies and structured knowledge become a practical foundation for AI agents that remember, reason, and act reliably?

THREE COMMUNITIES

- **Ontology and knowledge-representation researchers**
- **AI-agent researchers**
- **Industrial R&D leaders**

BEYOND "KNOWLEDGE GRAPHS IMPROVE AI"

1. ***What do ontologies contribute to agent memory?***
2. ***What can be implemented today?***
3. ***What is still missing for dependable products?***

AI agents: where ontology meets memory

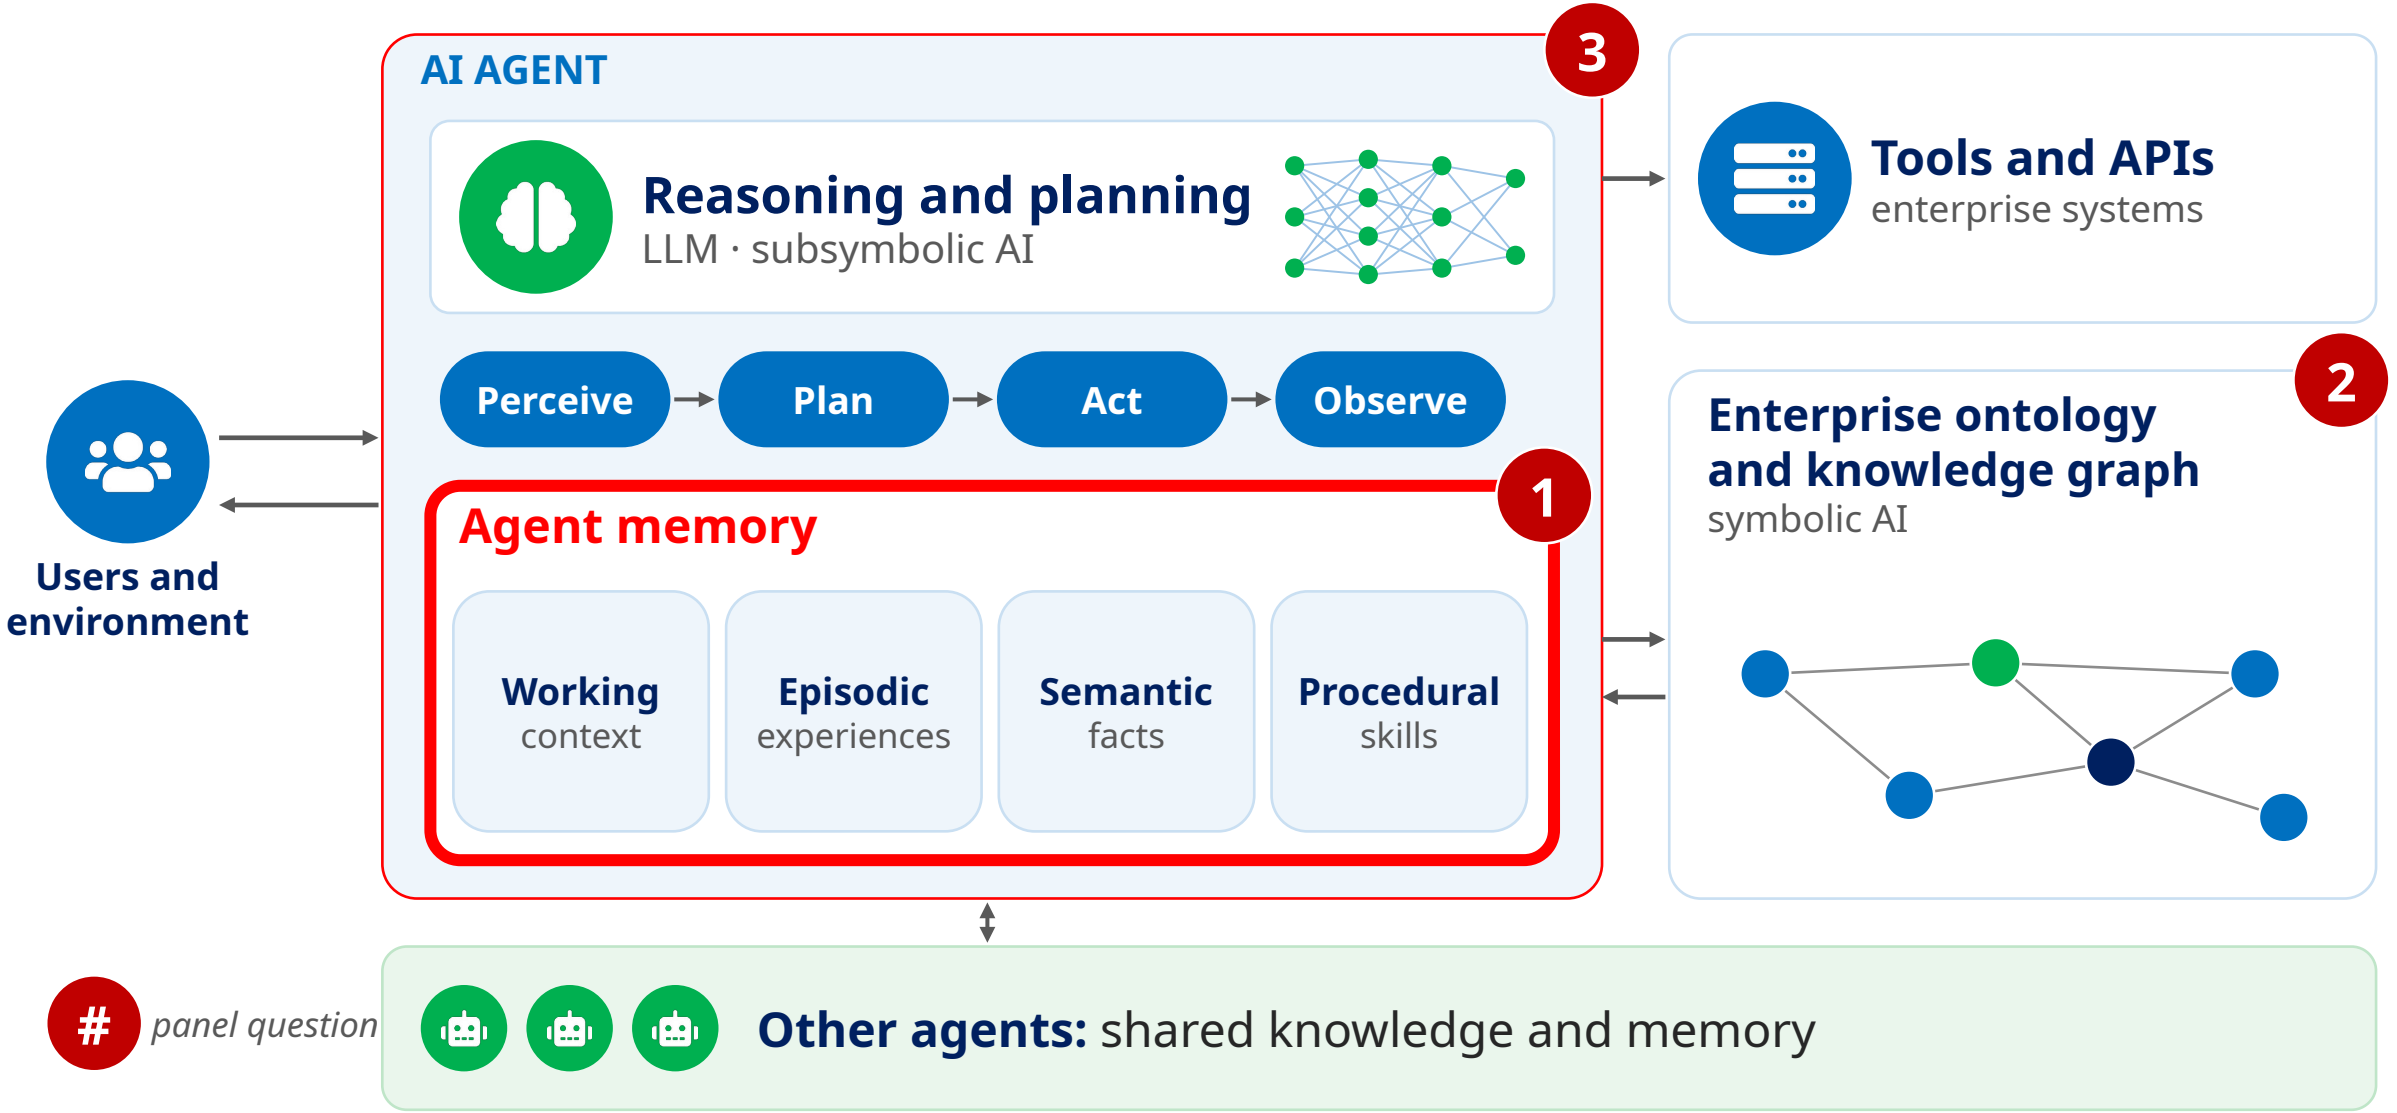

Common vocabulary (1/2): memory

1 **What an agent remembers**

$M_t = \langle W, E, S, P \rangle \qquad S \subseteq G$

$\text{keep}(x) \;\Leftrightarrow\; \text{value}(x) - \text{cost}(x) > \theta$

Facts, entities, events, preferences and decisions live in S, grounded in G.

2 **Consistent and trustworthy memory**

$M_{t+1} = M_t \oplus \Delta \quad \text{accepted iff } O \cup S_{t+1} \text{ consistent}$

$\text{prov}(x) = \langle \text{source, time} \rangle \qquad \text{retrieve if } \text{conf}(x) < \tau$

Updates respect the ontology; low confidence triggers retrieval (the agent doesn't blindly trust everything it remembers)

Notation

W, E, S, P: working, episodic, semantic, procedural memory
G graph
O ontology
Δ update

Common vocabulary (2/2): architecture

3 **Architecture**

$$a_t = \pi(o_t, \text{recall}(M_t, o_t))$$

$$\text{ans}(q) = \text{cert}(q, \langle O, \text{mappings}, D \rangle) \text{ or } \text{cert}(q, \text{warehouse}(O,D))$$

the agent (a) decides what to do based on what it sees now plus what it remembers that is relevant, (b) the agent asks questions in business terms, and the system translates them through the ontology and mappings into queries over the existing databases, returning only answers that are logically guaranteed.

4 **From research to product**

$$\text{adopt} \iff \text{gain(quality)} > \text{cost(build + maintain)}$$

in 2 years: one capability, one breakthrough

Teams adopt ontology-based memory when the gain exceeds the cost.

Notation: π policy · o observation · a action · cert certain answers · D enterprise data

Time plan: 15:20–16:50



15:20–15:25	Opening	*Moderator*
15:25–15:35	Position statements (2 min each)	*All panelists*
15:35–15:50	Q1 · What should an agent remember?	*d'Amato, Khan*
15:50–16:05	Q2 · Consistent, evolving, trustworthy memory	*Plexousakis, Ioannidis*
16:05–16:20	Q3 · The right architecture	*Calvanese, Khan*
16:20–16:40	**Q4 · A product within two years**	*All panelists*
16:40–16:48	Questions from the audience	*Moderator*
16:48–16:50	Closing summary	*Moderator*

Position statements (2’ each)

- ❑Yannis Ioannidis
- ❑Diego Calvanese
- ❑Dimitris Plexousakis
- ❑Arijit Khan
- ❑Claudia d’Amato

Panel Questions

01 · What should an agent remember?

What should an AI agent actually remember?

1 15:35

- Agents keep summaries, retrieve documents, use knowledge graphs: what is missing?
- What should memory represent explicitly?
- Where do ontologies make a fundamental difference compared with IR indexes, and vector stores?

WHAT COULD MEMORY REPRESENT?

- **Facts**
- **Entities and relations**
- **Events**
- **Preferences**
- **Goals**
- **Actions and decisions**

Unstructured

chunks, vectors

vs

Ontology-based

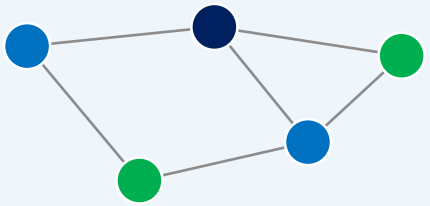

typed, explicit

OPENS **Claudia d'Amato · Arijit Khan**

Talks: *Doerr, Collarana, Jui*

02 · Consistent, evolving, trustworthy memory

How do we keep agent memory ***consistent, evolving, and trustworthy?***

2 · 15:50

- ***How should agents update memory, handle conflicts, separate facts from assumptions?***
- ***When should an agent retrieve evidence rather than trust what it remembers?***
- ***How can a data system manage memories for hundreds of thousands of agents?***

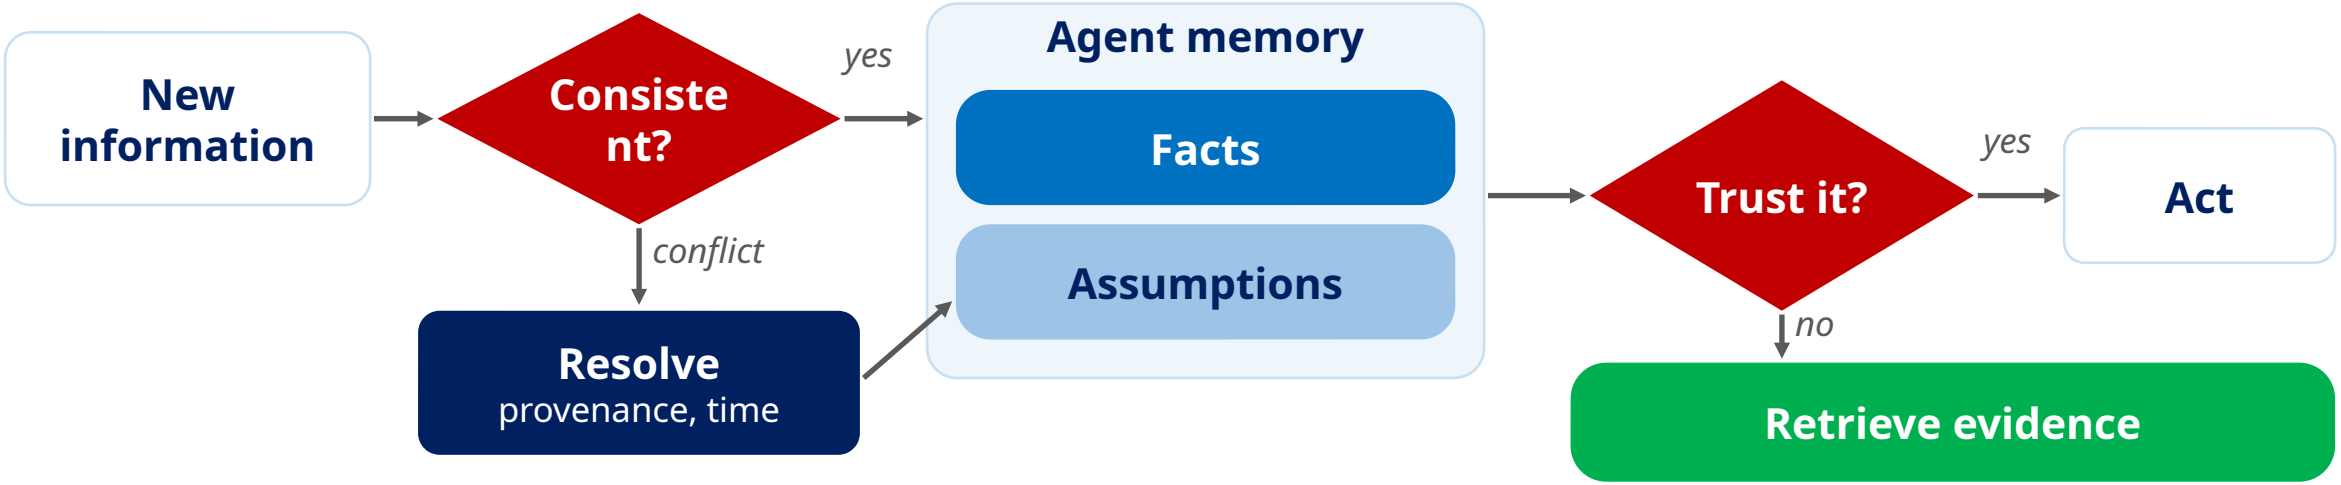

OPENS **Dimitris Plexousakis · Yannis Ioannidis**

Talks: *d'Amato, Bassiliades, Chalkiadakis, Jui*

03 · The right **architecture**

Combining LLMs, ontologies and agent memory in a real product

3 **16:05**

- *Knowledge graph, ontology-based data layer, vector store, relational DB, or a mix?*
- *What belongs in the foundation model, and what in explicit representations?*
- *How should we store and manage this data for the right trade-offs?*

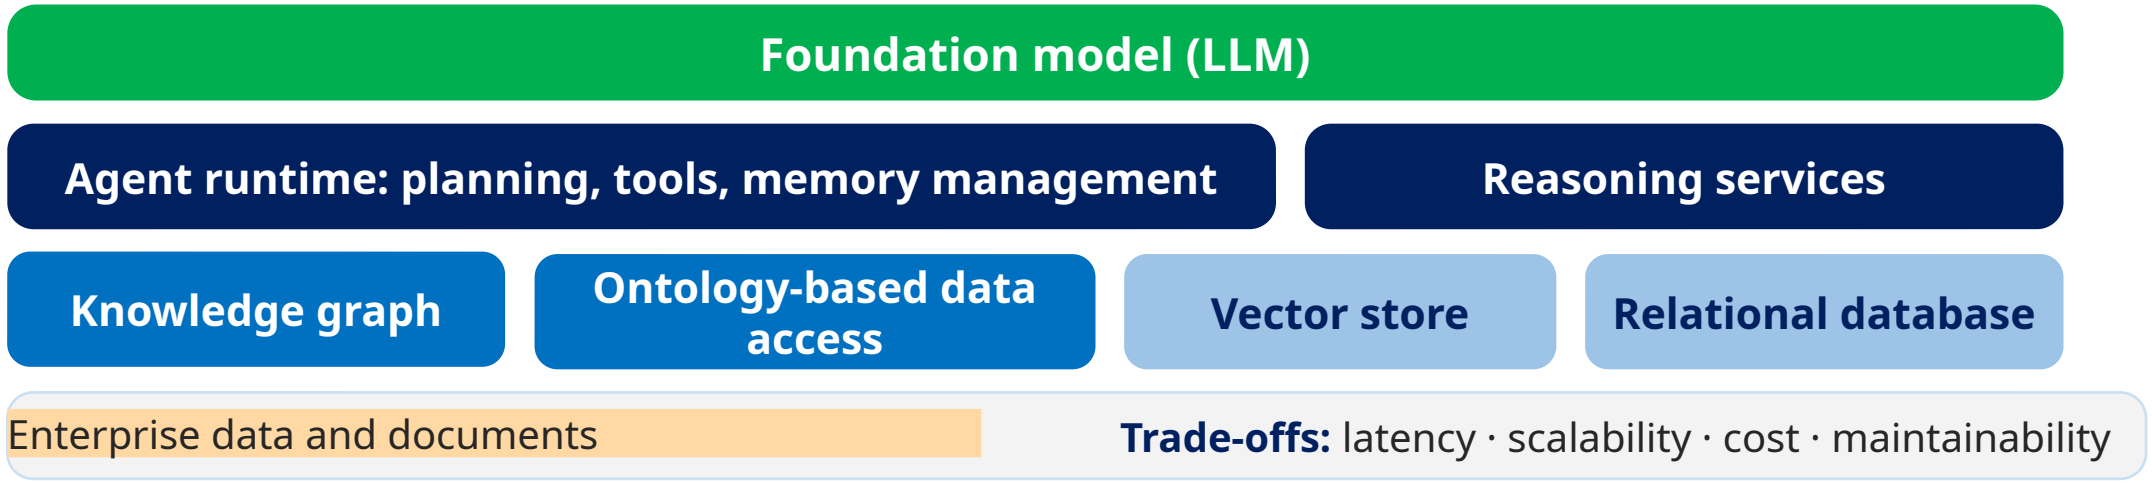

OPENS **Diego Calvanese · Arijit Khan**

Talks: *Calvanese, Khan, Magoutis, Zhao*

04 · A product within two years

What would it take to turn these ideas into a product capability?

4 16:20

- ***Which concrete use case should we prioritize?***
- ***What is the biggest remaining obstacle?***
- ***What measurable improvement would convince a product team?***

Use case
which one first?

2 years →

Product capability

Which measurable improvement convinces a product team?

BIGGEST REMAINING OBSTACLE?

- **Knowledge acquisition**
- **Ontology construction**
- **Integration with data**
- **Reasoning performance**
- **Evaluation**
- **Privacy**
- **Cost of memory**

ALL **Everyone: one capability, one breakthrough**

Talks: *Pan, Tsoumakas, Paliouras*

One capability, one breakthrough

Within two years: which capability should become a product, and what breakthrough makes it practical?

Panelist	Capability to productize	Breakthrough needed
Yannis Ioannidis		
Diego Calvanese		
Dimitris Plexousakis		
Arijit Khan		
Claudia d'Amato		

These answers feed the joint statement. Many thanks to our panelists!

Questions from the audience

Summit recap and closing
(including points for a draft joint statement)

Day 1 recap

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.

FORTH INSTITUTE OF COMPUTER SCIENCE

Day 1- 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

Day 2 recap

Day 2 recap: at a glance

KEYNOTES

Data and knowledge infrastructure

Diego Calvanese

Yannis Ioannidis

SESSION 4

Adaptive and decentralised intelligence

Panos Vassiliadis

Ioannis Pitas

Giorgos Chalkiadakis

SESSION 5

Operational AI and structural intelligence

Grigoris Tsoumakas

Kostas Magoutis

Jui Shang-Ling

Two keynotes and six talks, chaired by Theodore Patkos and Giorgos Flouris

Keynotes: infrastructure for the agentic era

Diego Calvanese

Knowledge Graphs from an Industrial Perspective

- Virtual knowledge graphs: the ontology reaches data where it lives
- From research to industry: Ontopic

Yannis Ioannidis

The Future of Data Infrastructure in the Agentic AI Era

- Agents become first-class users of data infrastructure
- AI for DataManagment
- Datamanagement for AI

Meaning moves into the data layer: infrastructure must integrate, remember and govern.

Adaptive and decentralised intelligence

Panos Vassiliadis

From analytic intentions to data stories

- Users state an intention, not a query
- Answers become data stories

Ioannis Pitas

Decentralized machine learning

- DNN agents teach and learn from each other
- No catastrophic forgetting

Giorgos Chalkiadakis

Safe uncertainty-driven exploration

- Uncertainty decides when to explore
- Safe self-evolution from experience

Intelligence adapts: to the user's intent, to other agents, and to its own uncertainty.

Operational AI and structural intelligence

Grigoris Tsoumakas

From semantic extraction to autonomous agents

- Efficient extraction, custom RAG
- Continual learning, data science agents

Kostas Magoutis

Autoscaling of AI services

- Memory is central to serving LLMs
- Scaling means moving state across nodes

Jui Shang-Ling

From data infrastructure to structural intelligence

- Ontology as backbone, memory as evolving knowledge
- Explicit levels of assurance

Operational AI needs structure: elastic infrastructure, continual learning, explicit assurance.

Day 2 in four messages

Semantics bridge data and decisions — *Calvanese · Jui · Vassiliadis*

Agents are the new users of data infrastructure — *Ioannidis · Magoutis · Tsoumakas*

Learning becomes collaborative and continual — *Pitas · Tsoumakas · Chalkiadakis*

Uncertainty and assurance must be explicit — *Chalkiadakis · Jui*

Thank you to all Day 2 speakers!

Day 1 and Day 2 messages

Trust is engineered, not assumed

Rigorous design, verification and evaluation around the model.

State and memory are infrastructure

Context, experience and reasoning traces need managed storage

Neural and symbolic, in both directions

LLMs build knowledge; ontologies guide, validate and explain

Evaluation is a research problem itself

Shared benchmarks, explanation protocols, locating misalignment

Knowledge lifecycles become continuous

Provenance-aware, agent-assisted curation keeps knowledge current

Semantics bridge data and decisions

Agents are the new users of data infrastructure

Learning becomes collaborative and continual

Uncertainty and assurance must be explicit

OLYMPIA AI DATA INFRASTRUCTURE SUMMIT 2026

RECAP, JOINT STATEMENT AND CLOSING · 16:50–17:10

Towards a joint statement

Yannis Tzitzikas

FORTH-ICS and University of Crete

Olympia AI Data Infrastructure Summit 2026 · Thursday 5 November 2026

Our goals for the last 20 minutes

Synergies

Which ideas and groups could work together?

Next steps

What we could do after Olympia, and who could leads it?

Joint statement

A short shared text we can all sign



Outcome: *an agreed draft of the Olympia joint statement and a short list of next steps, each with a coordinator.*

Five threads across the two days

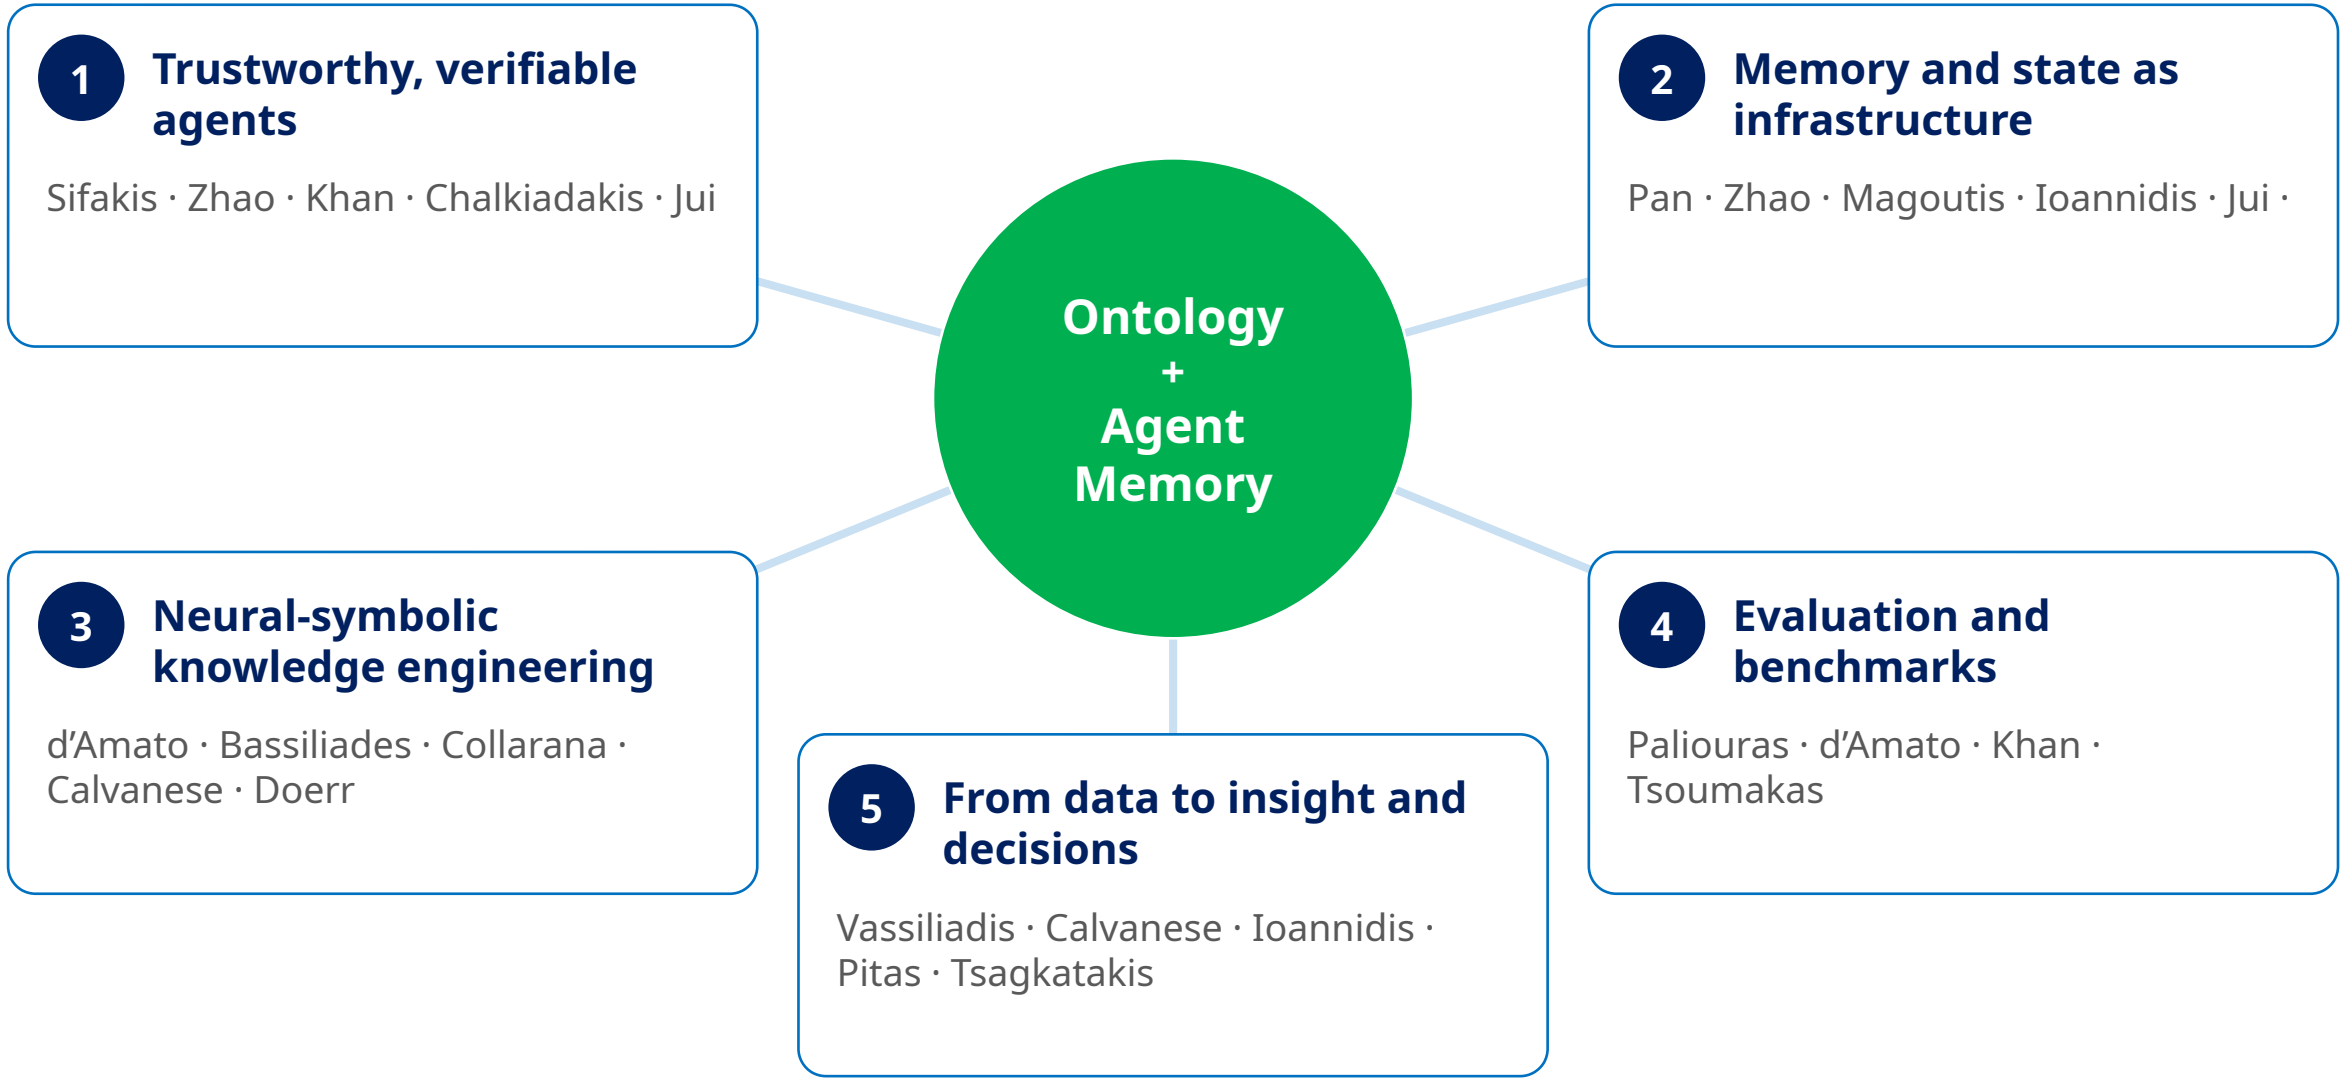

Interesting cross-overs among speakers

Pan × Zhao
Scaling agents with managed state and memory

Plexousakis × Jui
Belief revision and verifiable reasoning for evolving memory

Sifakis × Chalkiadakis
Trustworthy agents that evolve safely from experience

d'Amato × Paliouras
A shared benchmark for knowledge graph explanations

Calvanese × Doerr × Collarana
Patterns for enterprise ontologies and virtual knowledge graphs

Bassiliades × Tsoumakas
LLM-assisted ontology engineering and knowledge extraction

Ioannidis × Khan × Magoutis
Data infrastructure serving many agents at scale

Tsagkatakis × Pitas × Vassiliadis
From Earth-observation data to learned insight and data stories

Which other pairings do you see?

Potential next steps

A

Asilomar-style report

Build on the panel's common vocabulary; a short white paper, probably arxiv

B

A shared benchmark

A BioASQ-style challenge on agent memory and ontology quality

C

Joint proposals

Horizon Europe and national calls, academia and industry together

D

Open testbeds

Realistic enterprise scenarios and datasets with industry partners

E

Working groups

One small group per topic, each with a named coordinator

Potential next steps

A

Asilomar-style report

Build on the panel's common vocabulary; a short white paper, probably arxiv

Yannis Tzitzikas and Nikos Ntarmos could prepare the 1st version and then share to you for feedback and enrichments.

We could publish it at arxiv.
Raise your hand if you agree

Draft statement: Yannis v0.1 (before the summit)

- Agentic memory is a decision-making problem over information, not merely a storage problem. The system must decide ***what*** to preserve, ***how*** to represent it, when to ***update*** it, what **evidence** to retrieve, and **when not to answer**.
- These choices are coupled: a poor extraction can corrupt a graph, an over-compressed summary can erase a temporal exception, and a highly relevant retrieval can still be insufficient evidence for an action.
- Ontologies and knowledge graphs could contribute substantially where explicit semantics matter: stable identity, typed relations, temporal validity, provenance, constraint checking, multi-hop retrieval, and auditable updates.
- They contribute less directly to salience estimation, open-ended language interpretation, discovery of useful abstractions, and calibrated risk decisions. Thus symbolic AI is best treated as a *reliability and structure layer* inside a **hybrid memory architecture**, rather than as a complete memory solution.
- The key research question is: ***how to combine learned flexibility with explicit, testable constraints while measuring utility, cost, and risk over long-horizon interactions.***

Yannis thoughts during the summit

- How to build an autonomous and trustworthy agent
  - Equip it with an SL-Language model, and feed it with resources (data and docs) and tools.
    - Of course the better data is, the better
      - Therefore o-based integration(materialized or virtual) is great!
  - Based on the frequency of incoming requests the agent decides that it is worth building and maintaining a symbolic representation (for explainability, computational economy , etc). It creates it (if there is not any one available), and uses external knowledge to verify it (and aligning it with the terminology and meaning of the rest world).
    - It leverages symbolic ai techniques internally for exploiting it.
  - Based on its operation, it updates its memory (including the symbolic part)
  - Externally it exhibits contract-based operations and SLAs
    - These include trust
  - Agent memory can be seen as a sophisticated cache

Collaborative Statement

https://docs.google.com/document/d/17LMFnDjAR_mAslYqZKYVGUJjZrO1UB7xu0KwWIbJmuM/edit?tab=t.0

OLYMPIA AI DATA INFRASTRUCTURE SUMMIT 2026

DAY 2 · FINAL SESSION

Closing

Yannis Tzitzikas

FORTH-ICS and University of Crete

Olympia AI Data Infrastructure Summit 2026 · Thursday 5 November 2026

Thank you to our speakers and panelists

KEYNOTES	Joseph Sifakis	Diego Calvanese	Yannis Ioannidis		
DAY 1	Bo Zhao	Martin Doerr	Claudia d'Amato	Nikolaos Bassiliades	
	Diego Collarana	Georgios Paliouras	Grigoris Tsoumakas	Arijit Khan	Grigoris Tsagkatakis
DAY 2	Panos Vassiliadis	Kostas Magoutis	Ioannis Pitas	Giorgos Chalkiadakis	Jui Shang-Ling
PANEL	Yannis Ioannidis	Dimitris Plexousakis	Diego Calvanese	Arijit Khan	Claudia d'Amato

Many thanks for making Olympia 2026 happen

Our partner: Huawei

See Yeen Tan

Shen Yan

Nikos Ntarmos

Yongjian Lyu

Our host: FORTH

Lida Harami

Maria Moutsaki

Michalis Mountantonakis

Venue and support

Megaron Hotel, Heraklion

Synthesis Group

...and thank you all for two days of ideas and dialogue!

OLYMPIA AI DATA INFRASTRUCTURE SUMMIT 2026

Thank you, and safe travels!

“Data shapes memory, Knowledge drives intelligence.”

Stay in touch: olympia2026@ics.forth.gr

Olympia AI Data Infrastructure Summit 2026 · Heraklion, Crete · 4–5 November 2026