

The 12th International Artificial Intelligence Symposium – AIS 2026

*AIS: A Residential Symposium held in Conjunction with the
International Artificial Intelligence Summer School (IAISS)*

***Riva del Sole Resort & SPA – Castiglione della Pescaia (Grosseto) Tuscany, Italy
21 – 24 September 2026***

**<https://ais2026.icas.events>
ais@icas.cc**

Schedule Ver. 0.6 (5 pages) – September 15

	Mon, 21 Sept	Tue, 22 Sept	Wed, 23 Sept	Thu, 24 Sept
09:00 – 09:50	G. De Giacomo	R. Bogacz	Oral Session 2	G. Paliouras
09:50 – 10:40	G. De Giacomo	R. Bogacz		M. Eckstein
10:40 – 11:20	Coffee break	Coffee break	Coffee break	Coffee break
11:20 – 12:10	G. De Giacomo	S. L. Risi	Oral Session 3	M. Eckstein
12:10 – 13:00	J. Hayes	S. L. Risi		M. Eckstein
13:00 – 15:00	Lunch	Lunch	Lunch	Lunch
15:00 – 15:50	J. Hayes	S. L. Risi	Oral Session 4	Oral Session 6
15:50 – 16:40	C. Feichtenhofer	Sumi Helal		
16:40 – 17:20	Coffee break	Coffee break	M. Le Morvan	
17:20 – 18:10	C. Feichtenhofer	Oral Session 1	M. Le Morvan	
18:10 – 19:00	C. Feichtenhofer		M. Le Morvan	
19:00 – 20:00	Free Time		Oral Session 5	
20:00	Dinner	Social Dinner*	Dinner	Dinner

AIS 2026 & IAISS 2026

As in previous editions, the conference will be co-located with the *International Artificial Intelligence Summer School (IAISS 2026)*, fostering a unique and stimulating environment for interdisciplinary exchange.

This co-location is intended to strengthen the interaction between researchers, practitioners, and students, promoting the cross-fertilization of ideas and the development of new collaborations within the artificial intelligence community.

All conference participants will have full access to IAISS 2026 activities, including IAISS 2026 keynote lectures delivered by leading experts, thereby enriching the overall scientific experience and providing additional opportunities for learning, discussion, and networking.

IAISS/AIS 2026 Lecturers/Keynote Speakers

Rafal Bogacz, *University of Oxford, UK*

Giuseppe De Giacomo, *University of Oxford, UK*

Maria K. Eckstein, *Google DeepMind, London, UK*

Christoph Feichtenhofer, *Meta Superintelligence Labs, USA*

Jamie Hayes, *Google DeepMind, UK*

Sumi Helal, *University of Bologna, Italy & University of Florida, USA*

Marine Le Morvan, *INRIA Saclay, France*

Sebastian Lins Risi, *Sakana AI, Japan & IT University of Copenhagen, Denmark*

George Paliouras, *National Centre for Scientific Research "Demokritos", Athens, Greece*

IAISS/AIS 2026 Lectures

Rafal Bogacz, *University of Oxford, UK*

Lecture 1/2 “Predictive coding: training neural networks with local rules”

Lecture 2/2 “Predictive coding: advantages over backpropagation”

Giuseppe De Giacomo, *University of Oxford, UK*

Lecture 1/3 “Foundations of Guardrail AI via Temporal Synthesis”

Lecture 2/3 “Foundations of Guardrail AI via Temporal Synthesis”

Lecture 3/3 “Foundations of Guardrail AI via Temporal Synthesis”

Maria K. Eckstein, *Google DeepMind, London, UK*

Lecture 1/3 “Computational Cognitive Modeling Using Reinforcement Learning”

Lecture 2/3 “Advancing Cognitive Modeling by Combining Hand-Crafted and Neural-Network Models”

Lecture 3/3 “Advancing Cognitive Modeling using Procedurally-Generated Task Spaces”

Christoph Feichtenhofer, *Meta Superintelligence Labs, USA*

Lectures 1 and 2 “On the impact of data for vision-language learning”

Lecture 3/3 “Building frontier visual segmentation models”

Jamie Hayes, *Google DeepMind, UK*

Lectures 1 and 2 “Agentic Security (how to attack and defend against prompt injection)”

Sumi Helal, *University of Bologna, Italy & University of Florida, USA*

Lecture “Generative Smart Spaces – A Novel Physical-AI Approach to Programming Future Internet-of-Things”

Marine Le Morvan, *INRIA Saclay, France*

Lecture 1/3 “The rise of Tabular Foundation Models”

Lecture 2/3 “TabICL: on open Tabular Foundation Model”

Lecture 3/3 “Next frontiers for TFM”

Sebastian Lins Risi, *Sakana AI, Japan & IT University of Copenhagen, Denmark*

Lecture 1/3 “Neuroevolution: The basics”

Lecture 2/3 “Neuroevolution: Diversity, Novelty, and Open-Ended Discovery”

Lecture 3/3 “Neuroevolution: Synergies with Generative AI”

George Paliouras, *National Centre for Scientific Research "Demokritos", Athens, Greece*

Lecture “Neuro-Symbolic Reasoning and Learning for Complex Event Recognition/Forecasting”

Useful Information

1. **Registration:** *Sept 20 @ 16:30 – 20:00* (in the foyer of the conference centre near The Crown conference room). Inside the Riva del Sole it is mandatory to wear the badge.
2. Conference Room: **The Crown** (in the conference centre of the Riva del Sole Resort & SPA, num 7 on the resort map).
3. **Very Important:** *during lessons and Q&A sessions, absolute silence is recommended, you must not make noise, you must not disturb the Lecturer during the lesson or while she/he is answering questions posed by the audience.*
4. **Group Photo:** *Sept 21 @ 13:00 Meeting point: terrace adjacent to the restaurant.*
5. **Resort Map** and **Castiglione della Pescaia Map**.
6. Obviously, we recommend bringing your work laptop.
7. Please test your presentation during coffee breaks/lunches. Please remember *to bring your own laptop/tablet*. As a backup, you could also bring a USB stick. There will be HDMI access. It's always advisable to bring your own adapters. VERY IMPORTANT: before the presentation, during a coffee break or lunch break, test your presentation by connecting your laptop/tablet to the video projector and scrolling through/viewing all the slides!
8. **For social dinner on the Sept 22, we advise everyone **to dress up** (jacket and tie, suit, etc.).*
9. **Certificates** will be given out on the last day of the course. To receive the certificate you must have at least 85% of class attendance (we have CNNs/Transformers to compute and infer attendance ;-).
10. One last thing does not forget bathing suits, sunscreen, and stay hydrated.

AIS Leonardo Fibonacci Best Paper Award: <https://ais2026.icas.events/best-paper-award/>

Previous Editions: <https://ais2026.icas.events/past-editions/>

Past Keynote Speakers: <https://ais2026.icas.events/past-keynote-speakers/>

Oral Session 1, September 22, 17:20 – 20:00, The Crown Conference Room

- 17:20 – 17:35 Xiao Jin, Conner Parker and Rui Feng, "Structure-Informed Multi-Head Attention for Correlated Predictors"
- 17:35 – 17:50 Evann Rabeau, Jérôme Iste, Ismael Freire and Mehdi Khamassi, "Model-free multi-agent reinforcement learning simulation of the influence of empathy on the sustainability of agents' consumption in a common-pool resource dilemma"
- 17:50 – 18:05 Tian Li, Huizhi Liang, Lucy Robinson and Varun Ojha, "Predicting Cognitive Behavioural Therapy Outcomes based on Clinical Assessments and Referral Letters"
- 18:05 – 18:20 Valerio Bellandi and Stefano Siccardi, "Predicting Energy Consumption in Buildings and Mobility with ML and Graphs"
- 18:20 – 18:35 Ilinca Tiriblecea and Gabriel Turinici, "Inferring Missing Trajectory Data with Temporal Convolutional Networks"
- 18:35 – 18:50 Qingliang Meng, Luogeng Xiong, Wei Liang, Huizhi Liang and Tian Li, "Unified Timbre-Style Representation Learning for Controllable Speech Synthesis"
- 18:50 – 19:05 Simon Hakenes and Tobias Glasmachers, "Macro-Action Topological Navigation under Noisy Localization using Reinforcement Learning"
- 19:05 – 19:20 Jianfei Xu, Huizhi Liang, Judith Harrison, John-Paul Taylor, Jaume Bacardit and Varun Ojha, "HypTI: Hyperbolic Geometry-Aware Multimodal Fusion Framework for Dementia Diagnosis Based on Tabular-Image Data"
- 19:20 – 19:35 Tõnn Talpsepp, Kristjan Liivamägi and Tarvo Vaarmets, "What Does a Transformer Learn from an Investor's Trade Sequence and Educational Profile?"
- 19:35 – 19:50 Mark Callan and Meghana Kshirsagar, "Grounded Multi-Agent Reasoning for Financial Due Diligence: A Deterministic-First Architecture with Claim- Level Verification"

Oral Session 2, September 23, 9:00 – 10:40, The Da Vinci Conference Room (remote presentations)

- 09:00 – 09:15 Giuseppe Trimigno and Gianfranco Lombardo, "Prompt-Induced Schema Awareness for Question Answering over Knowledge Graphs"
- 09:15 – 09:30 Gabriel Turinici, "Softmax gradient policy for variance minimization and risk-averse multi armed bandits"
- 09:30 – 09:45 Muhammad Hashim, Roberto Rugani, and Marco Raugi, "Techno-Economic and Microscopic Simulation of DC Fast Charging for Electric Vehicles in Renewable Energy Communities"
- 09:45 – 10:00 Nadejda Bocheva, Miroslava Stefanova and Milena Mihaylova, "Drift-Diffusion Modeling of Orientation Effects in Contour-Defined Triangle Detection under Time Pressure"
- 10:00 – 10:15 Sarah Youssef, Hussein Assaf, Esam Abdel-Raheem and Ziad Kobti, "Clinical Prior-Guided Multimodal Learning Framework for Brain Tumor Segmentation"
- 10:15 – 10:30 Jose Barrera-Garcia, Victor Macho, Broderick Crawford, Álvaro Olmedo, Felipe Cisternas-Caneo, Yoslandy Lazo, Ricardo Soto, José Maria Jurado Carmona, Giovanni Giachetti and José Manuel Gómez-Pulido, "When Features Leak: A Leakage Audit and Imbalance Benchmark for Manchester Triage"
- 10:30 – 10:45 Nicola Floris, Gianmarco Cherchi, Gianni Fenu, Marco Manolo Manca, Roberto Saia, Riccardo Scateni and Alessandro Sebastian Podda, "A Data-Centric, Comparative Approach to Real-Time Pedestrian Orientation Estimation for Autonomous Driving"

Oral Session 3, September 23, 11:20 – 13:00, The Da Vinci Conference Room (remote presentations)

- 11:20 – 11:35 Asma Atamna, Tom Maus, Fabian Kievelitz and Tobias Glasmachers, "Cumulative Learning Rate Adaptation: Revisiting Path-Based Schedules for SGD and Adam"
- 11:35 – 11:50 Marco Pilloni, Alessandro Giuliani, Alessandro Sebastian Podda, Alessia Pisu, Livio Pompianu and Salvatore Carta, "Semantic Normalization of Coercive Knowledge Graph Triples with Large Language Models"

11:50 – 12:05 Gauri Vaidya, Gauri Mehta, Yuanyuan Chen, Alison O'Connor and Meghana Kshirsagar, "A Static Digital Twin Framework with Predictive Analytics for Healthcare Operational Intelligence"

12:05 – 12:20 Minhquang Nguyen, Vanha Tran, Thiloan Bui, Truongminh Ngo and Thanhthu Nguyen, "Domain-Adapted Embedding Fine-tuning based ESG Question-Answering System Using GraphRAG and Large Language Models: A Case Study for the Vietnamese Banking Sector"

12:20 – 12:35 Melchiorre Cosseddu, Salvatore Carta, Alessandro Sebastian Podda, Livio Pompianu, Marco Manolo Manca and Alessandro Giuliani, "Mapping Business Ecosystems with LLMs: An Automated Pipeline for Knowledge Graph Construction from Corporate Websites"

12:35 – 12:50 José Luis Carrillo-Medina, Ivone Arias-Almeida, Pablo Morales-Paredes and Kevin Chuquitarco-Velasco, "A Two-Stage Framework for Facial Identification and Torso-Centered Person Re-Identification Using Body Silhouette and Clothing Color in Non-Overlapping Cameras"

12:50 – 13:05 Cheng Arn David Goh, "Signed Decomposition of the Regression Structure via Linear JEPA Training: An Algorithm, Empirical Validation, and a Machine-Checked Proof"

Oral Session 4, September 23, 15:00 – 16:30, The Da Vinci Conference Room (remote presentations)

15:00 – 15:15 Shasha Yu, Fiona Carroll and Barry L. Bentley, "Execution-Refusal Profiling of Tool-Augmented LLM Agents: Behavioural Characterization Across Autonomy Scaffolds"

15:15 – 15:30 Elena Bellodi and Francesco Resca, "Smart City Analytics for Sustainable Tourism: an AI-Driven Framework for Flow Forecasting in a Historic Borough"

15:30 – 15:45 Elena Bellodi and Francesco Resca, "Forecasting and Interpreting Urban Air Pollution Under Weather, Traffic and Airport Influence"

15:45 – 16:00 Emmanuel Osei-Brefo, Huizhy Liang, Carlos Amador, Gustavo Leon Cazares, Matthew Davis and Varun Ojha, "GenAnalysis: A Contract-Driven Multi-Agent Framework for Auditable Data Analysis and Interpretation"

16:00 – 16:15 Maryam Mosleh and Ellis Solaiman, "Faithful Role-Aware Explanations for Adaptive Learning: Translating Structured XAI Evidence with LLMs"

16:15 – 16:30 Shu Li, Huizhi Liang, Wanqing Zhao, Mohamad Kassem, Ankush Prashar and Varun Ojha, "HypCR: Hyperbolic Convolution Recurrent Model with Mixture-of-Experts for Multivariate Time-Series Forecasting"

Oral Session 5, September 23, 19:00 – 20:00, The Crown Conference Room

19:00 – 19:15 Andrea de Ruvo, Pierluigi Castelli, Nicolas Radomski and Adriano Di Pasquale, "PhenoCal: Natural-Control Label Calibration for Cross-Cohort Machine Learning"

19:15 – 19:30 Stacey Truex and Margaret Malan, "Miles: A System for Parameter Analysis in Practical, Private Federated Learning"

19:30 – 19:45 Yannick Kees, Elena Hoemann, Frank Köster and Sven Hallerbach, "Holistic Risk Assessment for AI Systems"

19:45 – 20:00 Francesco Morano and Giovanni Denaro, "Phishing4training: an approach to train users on highly personalised phishing"

Oral Session 6, September 24, 15:00 – 19:15, The Da Vinci Conference Room

15:00 – 15:15 Ioannis Demetriou and Eva Perivolari, "Explainable and Monotonic Machine Learning for Hotel Booking Cancellation Prediction"

15:15 – 15:30 Derrick Chan and Dustin van der Haar, "Deep Learning Approaches to Restore Rainy Images in Adverse Weather Conditions"

15:30 – 15:45 Ioannis Sfyarakis, Adam Booth, Rajiv Ranjan and Ellis Solaiman, "EdgeCross: A Distributed AI-Assisted Framework for DeFi Smart Contract Auditing"

15:45 – 16:00 Absalom Ezugwu, "A Unified World Model for Predictive Representation Learning and Latent-Space Decision-Making"

16:00 – 16:15 Michal Skalican, Patrik Kozlik, Filip Zatroch, Maximilian Strecansky, Stefan Dlugolinsky, Martin Bobak, Ladislav Hluchy and Igor Skrlec, "Explainable AI-Driven Pipeline Composition for Cloud-Edge Data Platforms"

16:15 – 16:30 Tahitoa Leygue, Astrid Sabourin, Christian Bolzmacher, Sylvain Bouchigny, Margarita Anastassova and Quoc-Cuong Pham, "CADENCE: Congruence-Aware Word-Level Fusion for Speech Emotion Recognition"

16:30 – 16:45 Agate Jarmakovica and Arnis Kirshners, "Transformer-Based Semantic Variable Layer for AI-Assisted Quality Assessment of Tabular Healthcare Data"

Oral Presentation

16:45 – 16:55 Benny Stein and Abhinav Pareek, "A Trace-Based Monitoring Approach for Trustworthy Agentic AI"

Remote Presentations

16:55 – 17:10 Nikolay Georgiev, "Systemic Divergence in LLM-Based Personality Trait Estimation from Text Poses Risks for Psychological Applications"

17:10 – 17:25 Kexin Zheng and T. K. Satish Kumar, "A FastMap-Based Framework for Black-Box Optimization"

17:25 – 17:40 Vernon Kok, Absalom Ezugwu and Michael Olusanya, "Adaptive Object-Aware Mapless Mobile Robot Navigation Using YOLOE and Soft Actor-Critic"

17:40 – 17:55 Jorge Jabón Domínguez, "MGGAN: Learning a Multiscale Graph-based GAN for non-parametric 3D single-shape deformation and variation"

Education Track

17:55 – 18:05 Rommel Regis, "Prompt Engineering for Learning Mathematical Optimization with ChatGPT"

Oral Presentation

18:05 – 18:15 Vera Kalinichenko, "Vibe Coded Applications in the Industry Settings"

Short Paper

18:15 – 18:25 Mireia Louzan Calo, Sergio Consoli, Alessandro Zani, Mario Ceresa and Vicenç Gómez, "Federated Learning for Multimodal Wearable Human Activity Recognition: Preliminary Experiments"

Late Breaking Papers

18:25 – 18:40 Hasibul Karim Shanto and R. M. Shahbab, "Towards Robust and Stable Single Decision Trees under Distributional Shift"

18:40 – 18:55 Alejandro Aybar Flores, Rocío Maehara and Luis Benites, "Explaining and Segmenting Inequality Perceptions in Peru: Regression Evidence and Clustering from ENADES 2024"

Short Paper

18:55 – 19:05 Mohamed El Moutaouadia El Idrissi, "Mohamed Sraitih, Parisa Ghodous, Guillaume Besch, Romain Bronnert and Amir Hajjam El Hassani, "Safe Active Learning for Patient-Specific Pharmacodynamic Identification in Propofol Sedation"

Oral Presentation

19:05 – 19:15 Pankaj Kumar Chauhan, Sachchida Nand Chaurasia and Shyam Sundar, "A Reinforcement Learning based Evolutionary Hyperheuristic for the Coloured Balanced TSP"