FAAI25 : Foundations and applications of artificial intelligence

FAAI25 program

Monday, October 20, 2025

Time	Event
13:30 - 13:40	Registration
13:40 - 14:00	Welcome address
14:00 - 15:30	Oral presentations
14:00 - 14:30	› From Machine Learning to Large Language Models for Spatial Human Mobility - Ana-Maria Olteanu-Raimond, Laboratoire en Sciences et Technologies de l'Information Géographique, France
14:30 - 15:00	› Preventive Maintenance Solution Based on AI Acoustic & Vibration Fingerprinting - Adrian Brezulianu, "Gheorghe Asachi" Technical University of Iasi, Romania
15:00 - 15:30	› DP-MicroAdam: Private and Frugal Algorithm for Training and Fine-tuning - Mihaela Hudisteanu, Alexandru Ioan Cuza University of Iași, Romania
15:30 - 15:50	Coffee break
15:50 - 17:20	Oral presentations
15:50 - 16:20	› Exploring the Facets of Cross-Disciplinarity in Data Creation and Completion - Corina Dimitriu, "Alexandru Ioan Cuza" University of Iași, Faculty of Computer Science, Romania
16:20 - 16:50	› Inferring Missing Trajectory Data with Temporal Convolutional Networks - Ilinca Tiriblecea, Oxford Technical Solutions, UK

Time	Event
16:50 - 17:20	› Convergence of a stochastic gradient algorithm and an application for a modified MAB - Stefana-Lucia Anita, Octav Mayer Institute of Mathematics of the Romanian Academy, Iasi, Romania
17:20 - 17:40	Coffee break
17:40 - 19:10	Oral presentations
17:40 - 18:10	› The Generative AI Economy: A Multi-Framework Analysis of Business Models, Market Dynamics, and Ethical Governance - Sabina Necula, Alexandru Ioan Cuza University of Iași, Romania
18:10 - 18:40	› Exploratory Analysis of Dependencies and Formant Correlations in Continuous Romanian Speech - Maria Nestor, Faculty of Computer Science, Iaşi, Romania
18:40 - 19:10	› Deep hedging with low data usage and transaction costs - Gabriel Turinici, CEREMADE, Université Paris Dauphine - PSL, Paris, France

Venue

The "Aula Academiei" and “Main Hall” rooms are located in the Romanian Academy Iasi branch,Bulevardul Carol I, nr. 8 România, Iași, 700505.

From Machine Learning to Large Language Models for Spatial Human Mobility

Ana-Maria Olteanu-Raimond*†[1]

[1]Laboratoire en Sciences et Technologies de l'Information Géographique – Ecole Nationale des Sciences Géographiques, IGN, Université Gustave Eiffel – 73 Avenue de Paris, 94165 Saint-Mandé, France

Abstract

Nowadays, mobility data are available in various forms. For instance, transport operators publish open-access datasets (e.g., Vélib, bicycle counters), while crowdsourced GNSS data generated by citizens are shared across multiple platforms (e.g., Strava, IGNRando, OpenStreetMap, city walking tour videos).
In this presentation, after a brief introduction of spatial mobility nowadays challenges, two key research goals are addressed: **(i)** spatial data quality assessment, and **(ii)** transportation mode detection.

To achieve the first goal, a method to automatically detect outliers and assess the geometric accuracy of GNSS trajectories using a machine learning based classification approach is proposed. The method was validated on real crowdsourced GNSS trajectories collected in mountainous areas during outdoor activities such as running, walking, biking (Ivanovic et al. 2019).

To address the second goal, a new pipeline based on Large Language Models (LLMs) to detect transportation modes, such as car, bike, or walking, from urban GNSS trajectories was defined. The pipeline was first tested on the benchmark GeoLife dataset in Beijing, China, and then evaluated for transferability on a GNSS trajectory dataset collected in Paris, France (Badawi et al., 2025).

References:

Stefan Ivanovic, Ana-Maria Olteanu-Raimond, Sébastien Mustière, Thomas Devogele, A Filtering-Based Approach for Improving Crowdsourced GNSS Traces in a Data Update Context, ISPRS International Journal of Geo-Information, 2019, 8 (9), pp.380.
Amir Badawi, Ana-Maria Olteanu-Raimond, Arnaud Le Guilcher, Karine Zeitouni, Harnessing Large Language Models for Predicting Mobility Modes2025 26th IEEE International Conference on Mobile Data Management (MDM), June 2025, Irvine, France. pp.222-227

Keywords: spatial mobility, GNSS trajectory, machine learning, Large Language Model, outliers, update, transportation mode detection

*Speaker

†Corresponding author: ana-maria.raimond@ign.fr

Preventive Maintenance Solution Based on AI Acoustic & Vibration Fingerprinting

Adrian Brezulianu*[1]

[1]"Gheorghe Asachi" Technical University of Iasi – Bulevardul Profesor Dimitrie Mangeron 67, Iași 700050, România, Romania

Abstract

Modern industrial equipment generates unique acoustic and vibration patterns during normal operation. By capturing these "fingerprints" with advanced sensors, it is possible to create a baseline of healthy machine behavior. Artificial Intelligence (AI) techniques, particularly in signal processing and machine learning, can continuously compare real-time data against this baseline to detect subtle deviations that may indicate early signs of wear, misalignment, or impending failures.
This preventive maintenance solution provides a non-invasive, real-time method for monitoring assets such as motors, pumps, compressors, turbines, or rotating machinery. Unlike traditional maintenance schedules that rely on fixed intervals or visual inspections, AI-based fingerprinting enables predictive insights, alerting operators to anomalies before they escalate into costly breakdowns.

Key benefits include:

- *Increased equipment uptime* through early detection and timely interventions.

- *Reduced maintenance costs* by shifting from reactive repairs to condition-based servicing.

- *Enhanced safety and reliability* by preventing catastrophic failures.

- *Scalability across assets* since models can adapt to various machines and environments.
By integrating acoustic and vibration fingerprinting into existing maintenance strategies, organizations gain a powerful tool for asset longevity and operational efficiency. Over time, the AI system becomes more accurate as it learns from historical patterns, making it a continuously improving safeguard for critical infrastructure

Keywords: Preventive Maintenance, AI, Acoustic and Vibration Fingerprinting

*Speaker

DP-MicroAdam: Private and Frugal Algorithm for Training and Fine-tuning

Mihaela Hudisteanu*[1], Edwige Cyffers[2], Nikita Kalinin[2], and Ionut-Vlad Modoranu[2]

[1]Alexandru Ioan Cuza University of Iași = Universitatea Alexandru Ioan Cuza din Iași – Bulevardul Carol I 11, Iași 700506, Romania

[2]Institute of Science and Technology [Klosterneuburg, Austria] – IST Austria Am Campus 1 3400 Klosterneuburg Austria, Austria

Abstract

Machine learning models are often trained on sensitive data. Differential Privacy (DP) has become the standard framework to address these risks, providing mathematical guarantees that the output of an algorithm does not depend significantly on any single data point. The most widely used method for private training is DP-SGD, which enforces privacy through per-sample gradient clipping, addition of calibrated Gaussian noise, and early stopping. While effective, DP-SGD is sensitive to hyperparameter choices and often converges slowly. Adaptive optimizers, such as Adam, are attractive alternatives because of their strong performance in non-private training. However, in the differentially private regime, the combination of clipping and noise introduces bias into the moment estimates, leading to degraded performance.

We propose DP-MicroAdam, a differentially private optimizer that integrates standard DP mechanisms with the sparse, memory-efficient update strategy of MicroAdam. By combining top-k gradient selection, error feedback, and stable moment estimation, DP-MicroAdam reduces bias amplification and enhances robustness to noise. We prove that DP-MicroAdam converges in stochastic non-convex optimization at the optimal O(1/sqrt(T)) rate, up to privacy-dependent constants.
Empirically, it consistently outperforms existing adaptive baselines and remains competitive with, or surpasses, DP-SGD across several tasks, including CIFAR-10 and ImageNet training, as well as private fine-tuning of pretrained transformers. These results show that sparsity-aware adaptive methods can effectively recover much of the performance typically lost in differentially private optimization.

Keywords: Differential Privacy, Private Fine, Tuning, Adam

*Speaker

Exploring the Facets of Cross-Disciplinarity in Data Creation and Completion

Corina Dimitriu*[1]

[1]"Alexandru Ioan Cuza" University of Iași, Faculty of Computer Science – Strada General Henri Mathias Berthelot 16, Iași 700483, Romania

Abstract

This analysis centers on cross-domain representation learning for data augmentation and prediction enhancement. To be more specific, I am investigating the use of modern generative techniques for producing new features, while also prioritizing the integration of intrinsic interpretability into these methods, in the sense of inferring the features' structure and significance. I drew inspiration from complex architectures designed for semantic inpainting or missing-word prediction, as these are particular cases of feature synthesis: for semantic inpainting, the missing pixels can be viewed as missing features, and in my research their generation is carried out by a combination of GANs and GNNs, whereas for word completion, the missing tokens can be seen as absent features, and masked language modeling is employed to fill them in. My target to analyze is the reliable generalization of these methods, originally designed for one specific, widely approached and resourceful modality – video or text – to produce features for out-of-domain, perhaps underrepresented tasks (e.g., tasks requiring statistical features, electromyographic features or audio features). Overall, my work aims at solving a gap-filling representation learning objective, while implicitly studying the latent connections between different modalities and domains – the widely approached ones and the underrepresented ones – via knowledge transfer.

Keywords: feature generation, data augmentation, underrepresented modalities, generative modeling, text, video, representation learning, transfer learning.

*Speaker

Inferring Missing Trajectory Data with Temporal Convolutional Networks

Ilinca Tiriblecea*[1]

[1]Oxford Technical Solutions – Park Farm Business Centre Middleton Stoney OX25 4AL, United Kingdom

Abstract

This work tackles the issue of trajectory inpainting, which involves reconstructing missing subsequences in a partially observed trajectory. It uses a Temporal Convolutional Network (TCN) architecture. We train the model on fully observed land-vehicle trajectories. To create missing intervals, we apply binary masks to continuous subsequences. A masking-aware loss function limits error calculation to the masked areas. This ensures the network concentrates on rebuilding missing intervals while using contextual information from the observed parts of the trajectory. We evaluate this method on trajectory datasets within the same domain. The results show that TCNs offer a strong framework for accurate and consistent trajectory completion. This highlights the potential of TCN-based models for reliable inference in real-world situations where sequential data is missing.

Keywords: trajectory inpainting, trajectory inference, temporal convolutional network, TCN

*Speaker

Convergence of a stochastic gradient algorithm and an application for a modified MAB

Stefana-Lucia Anita*[1] and Gabriel Turinici[2]

[1]Octav Mayer Institute of Mathematics of the Romanian Academy – bd Carol I, nr 8, Romania
[2]Universite Paris Dauphine – université Paris Dauphine, PSL Resarch University – Paris, France

Abstract

We discuss the convergence of the Stochastic Gradient Descent (SGD) when the learning rate follows an inverse time decay schedule. Then we apply the ideas to the convergence of a modified form of Multi-Armed Bandit (MAB).

Keywords: stochastic gradient algorithm, reinforcement learning, multi armed bandit

*Speaker

The Generative AI Economy: A Multi-Framework Analysis of Business Models, Market Dynamics, and Ethical Governance

Sabina Necula*[1]

[1]Alexandru Ioan Cuza University of Iași = Universitatea Alexandru Ioan Cuza din Iași – Bulevardul Carol I 11, Iași 700506, Romania

Abstract

The commercial deployment of large language models and image generation systems since 2022 has catalyzed unprecedented business model innovation across industries, raising fundamental questions about value creation, competitive dynamics, and societal implications. This research systematically examines generative AI's transformative impact through an integrated multi-framework analysis combining Business Model Canvas, Porter's Five Forces, Stakeholder Theory, and Responsible AI frameworks.

Our investigation analyzes 156 peer-reviewed sources, 45 organizational case studies, and regulatory documents spanning 2020-2025, identifying six distinct business model archetypes: AI-as-a-Service, AI-Enhanced Products, Content Marketplaces, Data Licensing, Vertical AI Solutions, and Co-Creation Platforms. Porter's Five Forces analysis reveals extreme market concentration, with 5-10 foundation model providers facing capital requirements of $100 million to $1+ billion, creating structural barriers that favor incumbents despite democratization forces from open-source models and efficiency innovations.

Economic impact projections indicate substantial value creation potential-Goldman Sachs estimates $7 trillion GDP contribution over 10 years with 1.5 percentage point annual productivity growth-alongside critical uncertainties, as 70% of AI projects fail to scale beyond pilot stage. Stakeholder analysis demonstrates asymmetric impacts across workers (80% workforce exposure), content creators (licensing opportunities versus displacement), consumers (democratized access versus privacy risks), and society (digital divide amplification).

Ethical governance analysis through Responsible AI frameworks identifies persistent challenges in bias mitigation (34.7% error rate disparities documented), transparency (technical opacity of billion-parameter models), privacy (training data and model memorization risks), and safety (3-10% hallucination rates). The EU AI Act represents emerging regulatory responses, imposing requirements for high-risk systems with penalties up to €35 million or 7% global revenue.

Our findings reveal a bifurcated future: concentration forces driven by extreme capital requirements and data advantages coexist with democratization pressures from open-source models and architectural innovations. This paradox necessitates proactive governance frameworks balancing innovation incentives with fairness, transparency, and accountability principles. We provide strategic recommendations for business leaders navigating deployment decisions and policymakers designing regulatory interventions that promote responsible AI development while capturing economic value.

*Speaker

Keywords: Generative AI, Business Models, Porter's Five Forces, Market Concentration, AI Ethics, Responsible AI, AI Governance

Exploratory Analysis of Dependencies and Formant Correlations in Continuous Romanian Speech

Maria Nestor*[1]

[1]Faculty of Computer Science [Iaşi, Roumania] – Facultatea de Informatică - Universitatea "Al. I. Cuza" - General Berthelot, 16 - IAŞI 700483, Romania

Abstract

*

Keywords: formants, Pearson correlation, histograms, z, test, Fisher z, transformation

*Speaker

Deep hedging with low data usage and transaction costs

Pierre Brugière[1] and Gabriel Turinici*[2]

[1]CEntre de REcherches en MAthématiques de la DEcision – Université Paris Dauphine-PSL, Centre National de la Recherche Scientifique – Place du Maréchal de Lattre de Tassigny 75775 - Paris Cedex 16, France

[2]CEntre de REcherches en MAthématiques de la DEcision – Université Paris Dauphine-PSL, Centre National de la Recherche Scientifique – Place du Maréchal de Lattre de Tassigny 75775 - Paris Cedex 16, France

Abstract

Deep hedging in financial applications is a domain where advances in neural networks can play an important role; but this is hindered by the non-stationarity of the data which limits the training efficiency and out-of-sample accuracy. To mitigate such circumstance we present a situation where scarce data allow to obtain results competitive with standard approaches that assume known model parameters.

Keywords: deep hedging, low data requirements, discrete time hedging, distributional reinforcement learning

*Speaker