Archimedes 3min PhD Talks + Posters 2025
Abstracts

Alexandros Lolos
Title: *Online Fair Division for Personalized 2-Value Instances*
Abstract: We study online fair division with additive agents where goods arrive sequentially and must be allocated immediately and irrevocably. Restricting our focus to two-value instances, we design an algorithm that guarantees approximate fairness at every step and improves over time. With the use of limited lookahead, we achieve stronger fairness guarantees through matching-based methods.

Minas-Marios Sotiriou
Title: *EFX in multigraphs*
Abstract: We show that EFX allocations always exist multigraphs and general monotone valuations if any of the following hold: either each agent has at most $\lceil n/4 \rceil - 1$ neighbors, where n is the total number of agents, or the shortest cycle with non-parallel edges has length at least 6.

Sotiris Kanellopoulos
Title: *Approximation Schemes for Subset Sum and Partitioning problems*
Abstract: We develop fully polynomial-time approximation schemes (FPTAS) for k-Subset Sum Ratio, i.e., the problem of finding k disjoint subsets of n positive intergers, such that the sums of the subsets are as close as possible to each other. We also extend an FPTAS to k-way Number Partitioning, a practical problem related to fair division.

Christos Pergaminelis
Title: *Finite Pinwheel Scheduling: the k-Visits problem*
Abstract: Pinwheel Scheduling is a fundamental scheduling problem, receiving as input n task deadlines and asking whether there exists an infinite sequence of task executions with no deadline ever being violated. We study a finite version of this problem, proving that it is strongly NP-complete even for two turns and identifying tractable special cases.

Thanos Tolias
Title: *Online posted-price mechanisms for budget feasible procurement auctions*
Abstract: We study online procurement auctions, where agents sequentially arrive in random order, each with privately held service costs. We focus on the posted-price setting, in which the buyer proposes a take-it-or-leave-it price to offer each agent upon arrival. We present the first constant-competitive mechanism for this setting.

Giannis Vlahos
Title: *The communication complexity of combinatorial auctions in graphs*
Abstract: We study truthful and non-truthful protocols for combinatorial auctions in which every item can be allocated to one of two agents, or more generally to a fixed number of agents. We show both positive and impossibility results for the communication complexity of approximating the optimal social welfare for general, subadditive, or XOS valuations.

Marina Kontalexi
Title: *Learning in Stackelberg Games*
Abstract: We study a bi-level toll pricing game in which operators (leaders) set tolls on network links, while non-atomic players (followers) minimize their transportation costs, given by the sum of link delays and tolls along their chosen path. Our focus is on the game's dynamics when operators employ learning algorithms.

Vasilis Christoforidis
Title: *Learning-Augmented Coordination Mechanisms*
Abstract: We study the fundamental problem of reducing the inefficiency of equilibria in congestion games via coordination mechanisms. We adopt the learning-augmented framework and design a coordination mechanism that modifies the latency functions based on a predicted demand vector. The mechanism offers an equilibrium strategy to the players that achieves strong performance guarantees.

Andreas Kontogiannis
Title: *Advances in Multi-Agent Learning, Decision Making and Computation*
Abstract: We aim to explore advances in the foundations of computation, online learning and multi-agent learning. The main pillar of research aims at showing positive results in the context of (a) learning in games (e.g., polyhedral games and Markov games) with structure, and (b) studying the computational complexity of non-convex optimization problems.

Vasilis Pollatos
Title: *Learning in Non-Convex Games: Algorithms, Dynamics, and Convergence Guarantees*
Abstract: The main goal of this PhD will be to design implementable algorithms capable of negotiating the fundamental obstructions faced by gradient methods in min-max problems (or, more generally, equilibrium problems) without a convex structure. This will involve combining techniques from zeroth-order optimization and multi-armed bandits (stochastic or non-stochastic) with the analysis of standard gradient methods for regularized learning in games.

Vasileios Moustakas
Title: *Design of analog and mixed-signal circuits for Machine Learning and Artificial Intelligence*
Abstract: Design of analog and mixed-signal circuits for Machine Learning focuses on creating energy-efficient, high-speed hardware that integrates the analog domain. Our research develops circuits enabling low-power, high-performance computation, supporting real-time learning and inference in AI applications, particularly for edge devices requiring compact, scalable, and intelligent hardware solutions.

Dimitrios Dimitropoulos
Title: *Efficient Approximate Nearest Neighbour Search in Vector Databases*
Abstract: We have developed a two-phase algorithm that first generates candidate sets from B+ trees like structures on each dimension, aggregating them (the candidates) via a scoring function. The second phase performs a precise, full-dimensional Euclidean distance evaluation on this pruned candidate set (only the top candidates) to determine the final approximate nearest neighbors.

Panagiotis Rigas
Title: *GARLIC: GAussian Representation LearnIng for spaCe partitioning*
Abstract: We introduce a space partitioning method that aims to efficiently query approximate nearest neighbors. We introduce information theoretic objectives and adaptive refinement to guide our solutions and balance between structural consistency and performance.

Nikos Theologis
Title: *How Unfair is Sticking to your own Opinion? Fairness in Opinion Formation Dynamics*
Abstract: We introduce fairness in opinion formation dynamics under the Friedkin–Johnsen model. Defining influence as a node's impact on final opinions, we propose minimal interventions in order to balance influence across competing groups. We derive closed-form equations, design algorithms, and validate them on synthetic and real networks.

Christos Tsapelas
Title: *Towards a Neural Database Execution Engine*
Abstract: We propose a new analytical query engine that extends relational algebra with AI-powered operators, enabling seamless querying across structured and unstructured data such as text, images, audio, and tables. We introduce a formal algebraic framework and a modality-agnostic execution model, allowing declarative, cross-modal analytics while integrating both foundation models (e.g., LLMs) and specialized ML/DL models for efficiency and flexibility, paving the way to advanced semantic analytics

Konstantinos Plas
Title: *LLM spatial reasoning and optimizations in geospatial RDF stores*
Abstract: We study how optimizations in triple stores and geospatial knowledge graphs impact large language model performance in tasks such as augmented retrieval and question answering. Our approach leverages techniques including LLM-driven knowledge graph summarization, view materialization, and the construction of "cleaner", more efficient geospatial knowledge graphs.

George Papadoulis
Title: *AQQUA: a private and auditable digital payment system*
Abstract: Privacy and auditability are often viewed as conflicting properties in payment systems. We design a decentralized, censorship-resistant payment system that achieves both simultaneously, enabling private transactions with verifiable audits.

Christos Fragkathoulas
Title: *Counterfactual explanations*
Abstract: AI makes decisions in loans and jobs. Models can inherit or introduce bias, so we audit group fairness using group counterfactual explanations. Our approach FACEGroup builds a feasibility graph, partitions groups into subgroups, and optimizes cost/coverage trade-offs (MIP/Greedy). New metrics reveal subgroup disparities, while FACEGroup yields feasible, compact counterfactual sets.

Nikos Spyrou
Title: *Counterfactual video generation and evaluation*
Abstract: We propose CSVC, a framework for causally faithful counterfactual video generation using latent diffusion models. By embedding causal graphs into prompts and guiding generation with a vision-language model loss, CSVC can produce realistic, causally consistent counterfactuals. It improves causal effectiveness without fine-tuning, enabling broad applications in media and healthcare.

Thomas Melistas
Title: *Benchmarking Counterfactual Image Generation*
Abstract: Generative AI enables powerful visual editing, but realistic counterfactual image generation demands causal consistency and lacks clear ground truth, complicating evaluation. This work introduces a unified benchmarking framework, integrating existing methods, datasets, and causal graphs. Results highlight Hierarchical VAEs' superiority, with a user-friendly, extensible Python package for future community research.

Efstathios Karypidis
Title: *Semantic Future Prediction*
Abstract: Semantic future prediction refers to forecasting how scene semantics like objects, depth, and spatial structure will evolve over time. As a fundamental capability of world-models, it enables accurate, robust, and efficient foresight, allowing autonomous systems and intelligent agents to operate safely, reliably, and effectively in complex, dynamic environments.

Giannis Kalogeropoulos
Title: *Learning in weight spaces*
Abstract: Weight space learning explores NN parameters as a new data modality. It enables model analysis, synthesis, and parameter generation, offering pathways to automate ML pipelines and unify representations across applications, including neural fields and implicit representations. In our work, we propose ScaleGMN that adapts the MPNN paradigm by incorporating scaling symmetries, achieving state-of-the-art results.

Andreas Zamanos
Title: *Self-supervised learning for generalizable particle picking in cryo-EM micrographs*
Abstract: Cryo-EMMAE is a self-supervised method designed to overcome the need for manually annotated cryo-EM data. cryo-EMMAE leverages the representation space of a masked autoencoder to pick particle pixels through clustering of the MAE latent representation. Evaluation across different datasets demonstrates that cryo-EMMAE outperforms state-of-the-art supervised methods in terms of generalization.

Daphne Tsolissou
Title: *VLMs for medical imaging*
Abstract: The application and adaptation of recent Large Vision-Language Models (LVLMs) for stroke risk stratification is examined, using carotid ultrasound data from patients with atherosclerosis and specialized prompts. During zero-shot evaluation, 3 out of 6 models recognized the imaging modality and anatomy but provided incorrect risk assessment. Adaptation with additional clinical data significantly improved performance, suggesting potential future clinical use.

Triantafillos Papadopoulos
Title: *Finance Market Prediction Using LLMs*
Abstract: We introduce Plutus-ben, the first Greek financial NLP benchmark, and Plutus-8B, a Greek financial LLM fine-tuned on domain-specific data. Covering five key tasks with expert-annotated datasets, our evaluation of 22 LLMs highlights linguistic and domain challenges, underscoring the need for Greek-specific financial models and multilingual inclusivity.

Dimitris Tsirmpas
Title: *Synthetic moderation / facilitation of online spaces using LLMs*
Abstract: Social media have become hostile spaces, where abuse, polarization and misinformation harm individuals and endanger democracy. Hostile LLM actors have also introduced online spaces to hybrid warfare. LLMs are the only scalable tool capable of helping online users. We use them to simulate users, in order to train LLM facilitators.

Dimitris Koutsianos
Title: *Synthetic Speech Source Detection using Metric Learning*
Abstract: This study investigates synthetic speech source tracing by adapting speaker recognition methodologies. Experimental results demonstrate that a ResNet-based, metric-learning pipeline achieves performance competitive with, and at times superior to, self-supervised learning systems. These findings underscore the efficacy of ResNet architectures for audio forensic analysis of manipulated media.

Vivian Platanou
Title: *From Stone to Paper and from Paper to Pixel: Deep learning-based Inpainting for Fragmentary Epigraphic Texts*
Abstract: This study addresses the problem of restoring fragmentary ancient inscriptions, where material deterioration often renders the text incomplete. The aim is to investigate the application of models, and in particular inpainting methods, for the digital reconstruction of lost parts. The proposed methodology seeks to enhance the existing tools of philologists by offering a computational approach to the study and interpretation of texts.

Angelos Semoglou
Title: *Confidence in Assignments via K-Partition Ensembles*
Abstract: Clustering often lacks measures of assignment reliability, with algorithms like k-means sensitive to noise and initialization. We introduce CAKE, a framework combining stability and geometric consistency across clustering ensembles into an interpretable [0,1] pointwise confidence score.

Theodoros Tsironis
Title: *Landscape Complexity for the Empirical Risk of Generalized Linear Models*
Abstract: We calculate, to first exponential order, the number of critical points of certain families of random loss functions for gaussian data with a covariance structure. We focus in the high dimensional asymptotic regime, in which the number of data is proportional to their dimension. This is the only asymptotic regime in which the problem is non-trivial and consequently it is the relevant regime for modern data science.

Symeon Mastrakoulis (no poster)
Title: *MMS for Subadditive Valuations with Few Agents*
Abstract: We study the problem of fairly allocating a set of indivisible items among a set of agents. We consider the notion of (approximate) maximin share (MMS) and we provide an improved lower bound of 1/2 for the case of subadditive valuations when the number of agents is at most four.

Thodoris Kouzelis (no poster)
Title: *Boosting generative image modeling via joint image-feature synthesis*
Abstract: Latent diffusion models (LDMs) dominate high-quality image generation, yet integrating representation learning with generative modeling remains a challenge. We introduce a novel generative image modeling framework that seamlessly bridges this gap by leveraging a diffusion model to jointly model low-level image latents (from a variational autoencoder) and high-level semantic features (from a pretrained self-supervised encoder like DINO).

Anastasios Gerontopoulos (no poster)
Title: *Multi-Token Prediction Needs Registers*
Abstract: We propose MuToR, a simple and effective approach that leverages register tokens to predict future targets and enrich supervision for autoregressive transformers. Our method introduces only a negligible number of additional parameters and requires no architectural changes—ensuring compatibility with off-the-shelf pretrained language models.

Antonis Korinthios (no poster)
Title: *Gene Clustering from Single-Cell Data using Community Detection*
Abstract: Data generated using single-cell RNA sequencing is characterized by high dimensionality, extreme sparsity, and technical noise, making its analysis both challenging and computationally demanding. To address these challenges, we propose a method for constructing cell and gene graphs and analyze the application of our message-passing community detection algorithm to single-cell data. In addition, we propose new methods for selecting Highly Variable Genes (HVGs).