Publication List: Emmanouil V. Vlatakis-Gkaragkounis

Current Projects

Smoothed Analysis of Hidden Monotone Games and Neural Kelly auctions
(joint work with M. Jordan)

We delve into a scenario where N-players participate in a monotone game, each steered by their respective neural network advisors. Through the lens of smoothed analysis, we aim to discern the feasibility of reaching a Nash Equilibrium in such environments. As an illustrative case, we highlight Kelly auctions to demonstrate the intricate dynamics that emerge when players utilize neural networks as their bidding strategists.

TLDR: Using smoothed analysis, we explore the likelihood of reaching Nash Equilibrium in monotone games where players are guided by neural network advisors, spotlighting the real-world complexities in Neural Kelly auctions.
Toolbox: Smoothed Analysis, Game theory, Convex Optimization, Auction Theory

The Influencer's Problem: Optimal Revenue Dynamic Linear Contracts
(joint work with G. Guruganesh, I. Cohen, J. Schneider, J. Wang, M. Weinberg, Y. Kolumbus)

In a strategic game, we study the decisions made between an agent (the learner) and a principal (the optimizer). The agent has n actions, each with costs known only to the principal. Over T rounds, the principal offers contracts based on m possible outcomes. These offers can be simple linear contracts or more complex ones. Meanwhile, the agent uses various learning algorithms, ranging from no-regret types to the more intricate no-swap-regret. This setup blends strategy, adaptive learning, and economic tactics, creating a dynamic decision-making environment.

TLDR: Optimal-Revenue Dynamic Linear Contracts for adaptive learning Agents.
Toolbox: Contract Theory, No-Regret Learning Dynamics, High-Dimensional Geometry

Statistical inference for model parameters in overparametrized networks
(joint work with R. Chohan, M. Jordan)

Within deep learning, blending traditional convex functions such as least squares with the intricacies of non-convex neural networks presents a significant challenge. Our research primarily focuses on inferring model parameters within these combined landscapes. We've developed a statistical method to create trustworthy confidence intervals for each neuron's weight in overparametrized networks.

TLDR: Statistical parameter inference in large highly-expressive neural networks. .
Toolbox: Probability, Ergodic Theory, Markov Chains, Optimization, Deep Learning Theory

Bounded-$\exists\mathbb{Q}$: The Computational Complexity of Tenth Hilbert's Problem and Rational NE
(joint work with M. Christ, M. Yannakakis)

The Tenth Hilbert Problem posed the intriguing question of determining the solvability of Diophantine equations. Our research ventures into this enigma, presenting insights into its computational complexity. We introduce and delve into the Bounded-$\exists\mathbb{Q}$ class, a computational domain echoing the characteristics of the ETR, yet uniquely tailored for rationals confined within specified bounds. Er establishes a newfound connection between the intricacies of tenth Hilbert's problem and Last Fermat's theorem from Number Theory and the existence of Rational Nash Equilibria.

TLDR: We dissect the Tenth Hilbert Problem's complexity, linking it to the Bounded-$\exists\mathbb{Q}$ class — the counterpart of the ETR class for bounded rationals.
Toolbox: Computational Complexity, Number Theory, Numerical Real Analysis

Breaking the Efficiency Frontiers for Online Portofolio and Quantum State Learning
(joint work with F. Vasconcelos, K. Tarun, U. Vazirani, M. Jordan)

This research showcases that using barycentric regularizer, it's possible to achieve polylogarithmic regret with feasible computational requirements. Importantly, we further adapt our findings to a broader issue: online learning of quantum states using log loss. This extended algorithm is a pioneering solution, offering polylogarithmic regret in the realm of quantum state learning.

TLDR: Leveraging barycentric regularizer, we develop the first time-efficient no-regret quantum state online algorithm under log-loss objective.
Toolbox: Quantum Computing, Online Learning, Convex Analysis

Langrangian in High Resolution ODE for Games
(joint work with A. Wibisono, N. Wadia, M. Jordan)

In our study, we delve into game dynamics using a continuous-time lens. Through the Lagrangian framework, we enhance our understanding of these dynamics by harnessing the power of high-resolution ordinary differential equations (ODEs). This approach serves as a bridge connecting theoretical continuous dynamics with tangible real-world applications.

TLDR: We extract Langragian functionals for modern game dynamics using HRDEs.
Toolbox: Differential Equations, Dynamical Systems, Game Theory.

TangoTek: AI-supported Tango Music Composing
(joint work with M. Strofalis, P. Liang, M. Jordan)

TangoTek presents a groundbreaking approach to music composition, particularly in the realm of Tango music. Utilizing AI algorithms and deep learning models, we've developed a system that can craft authentic Tango tunes. By blending traditional music theory with advanced machine learning techniques, TangoTek offers a harmonious fusion of human creativity and artificial intelligence.

TLDR: TangoTek uses AI machinery to compose authentic Tango music.
Toolbox: Non-convex Optimization, Transformers, Signal Processing.

The Competitive 2-Player Pandora's Box Problem
(joint work with E. Gergatsouli, Z. Bruno, D. Pizzaro)

The Competitive Two Player Pandora's Box Problem explores the dynamics of two players who take turns opening unknown reward boxes. These players must decide whether to claim a revealed prize or move forward, despite costs associated with opening each box. The challenge lies in optimizing their strategy considering both the unknown rewards and the opponent's actions. We detail modifications to Weitzman's strategy to pinpoint the dominant strategy for this two-player variant.

TLDR: Two players strategically open reward boxes, balancing between known costs and potential gains while anticipating each other's moves. We adapt Weitzman's strategy to this competitive setting.
Toolbox: Game Theory, Strategic Decision Making, Weitzman's Optimization.

PUBLICATIONS

Chaos persists in large-scale multi-agent learning despite adaptive learning rates
(joint work with L. Flokas, G. Piliouras)
Submitted to Mathematics of Operations Research Journal (**MOR**)

In the realm of multi-agent learning, achieving equilibrium in self-play is a complex endeavor. While adaptive learning rates have improved convergence in small games, their effectiveness in large-population settings remains elusive. Our study uncovers that, even when implementing these dynamic rates, chaos still dominates large population congestion games.

TLDR: Despite the promise of adaptive learning rates, chaos remains prevalent in large multi-agent learning scenarios, even with standard algorithms and minimal strategies.
Toolbox: Chaos Theory, Discrete Dynamical Systems, Asymptotic Calculus

Curvature-Independent Last-Iterate Convergence for Games on Riemannian Manifolds
(joint work with Y. Cai, A. Oikonomou T. Lin, M. Jordan)
Submitted to Journal of Machine Learning Research (**JMLR**)

In many machine learning applications, we aim to find a point of balance or equilibrium on curved spaces known as Riemannian manifolds. In this work, we make a significant advancement for the seminal Riemannian gradient descent (RGD) method. Our breakthrough, a first-of-its-kind, reveals that RGD converges to NE at a linear rate in strongly monotone settings, unaffected by distance distortions or the curvature of the space.

TLDR: RGD computes NE with a step size that is agnostic to the underlying curvature.
Toolbox: Differential Geometry, Game Theory, Convex Optimization, Potential Analysis

Smoothed Complexity of SWAP in Local Graph Partitioning
(joint work with X. Chen, C. Guo, M. Yannakakis,)
35th ACM-SIAM Symposium on Discrete Algorithms (**SODA**) *2024*

Local search is a widely used method for tackling optimization problems. However, even though it often works well in real-world scenarios, theoretical evaluations sometimes suggest it could be very slow. This contradiction arises because these theoretical tests often involve unlikely, extreme cases. In this study, we discovered that when minor unpredictable changes (or "noise") are introduced to the problem input, algorithms like SWAP, 1-FLIP, and 2-FLIP efficiently handle a range of binary optimization challenges.

TLDR: Despite theoretical concerns, local search algorithms like SWAP and 2-FLIP handle provably real-world optimization tasks efficiently.
Toolbox: Smoothed Analysis, Applied Probability, Combinatorics, Graph Theory

Exploiting Hidden Structures in Non-convex Games for Convergence to Nash Equilibrium
(joint work with L. Flokas, I. Sakos, P. Mertikopoulos, G. Piliouras)
37th Conference on Neural Information Processing Systems (**NeurIPS**) *2023*

From adversarial attacks to multi-agent reinforcement learning, in vast majority of ML tasks, we're essentially playing multi-agent games striving for a Nash equilibrium. Beneath the surface of these complex non-convex games' landscapes often lies a simpler convex substructure. This research unveils that core using a pioneering method, the 'hidden gradient descent,' illuminating the hidden framework and nudging the game towards the Nash equilibrium.

TLDR: 'Hidden gradient descent': A novel method for computing Nash equilibrium in ML complex non-convex landscapes with hidden convex cores.
Toolbox: Game Theory, Convex Optimization, Dynamical Systems

Stochastic Methods in Variational Inequalities: Ergodicity, Bias and Refinements
(joint work with A. Giannou, Y. Chen, Q. Xie)
37th Conference on Neural Information Processing Systems (**NeurIPS**) *2023*

In machine learning tasks, the Stochastic Extragradient (SEG) and Stochastic Gradient Descent Ascent (SGDA) algorithms have stood out in min-max optimization and variational inequalities (VIP). While constant step-size versions offer benefits like ease of tuning, their convergence patterns remain intricate. Through Markov Chains, our study provides pivotal insights establishing first-of-its-kind Law of Large Numbers and the Central Limit Theorem, demonstrating their consistent performance. Finally, we show how to improve the solution accuracy in VIPs using Richardson-Romberg extrapolation.

TLDR: Stochastic EG and GDA algorithms are governed by the Law of Large Numbers and the Central Limit Theorem and their accuracy can be boosted using Richardson extrapolation.
Toolbox: Probability, Ergodic Theory, Markov Chains, Optimization, Dynamical Systems

Quadratic Speedup in Finding Nash Equilibria of Quantum Zero-Sum Games
(joint work with F. Vasconcelos, P. Mertikopoulos, G. Piliouras, U. Vazirani, M. Jordan)
7th Conference on Quantum Techniques in Machine Learning (QTML) **Oral** *2023*

Quantum game theory, which has ties to non-local games and quantum networks, has recently regained attention due to advancements in related fields and machine learning. One key goal is to find the best strategies for players, known as Nash equilibria in fully competitive settings. A 2008 algorithm by Jain and Watrous found these equilibria for quantum games at $O(1/\sqrt{T})$. Our new method, the Single-Call Matrix Mirror Prox (1-MMP), offers quadratic speedup, setting a new standard in the field.

TLDR: State-of-art method for Computing Nash Equilibria in Quantum Zero-Sum Games.
Toolbox: Quantum Computing, Game Theory, Convex Optimization

The Computational Complexity of Multi-player Concave Games and Kakutani Fixed Points
(joint work with C. Papadimitriou, M. Zampetakis)
24th ACM Conference on Economics and Computation (EC) *2023*

Kakutani's Fixed Point theorem is pivotal in game theory and economics, but its computational versions have been limited. This study offers a broad computational view of the theorem, showing it's PPAD-complete. It delves into multi-player concave games, highlighting the challenges of finding equilibriums, and sheds light on the Walrasian Equilibrium's complexity by proving an advanced version of Berge's maximum theorem.

TLDR: Computing Kakutani's Fixpoints is PPAD-complete. Both Rosen-Nash Eq. in concave games and Walrasian Eq. in exchange markets share similar complexity for general utilities.
Toolbox: Computational Complexity, Game Theory, Convex Analysis

Algorithms and Complexity for Computing Nash Equilibria in Adversarial Team Games
(joint work with P. Kalogiannes, I. Anagnostides, Y. Panageas)
24th ACM Conference on Economics and Computation (EC) *2023*

Adversarial team games involve a team competing against a single opponent in a win-lose situation. In this research, we show that finding an approximate Nash equilibrium in these games is as complex as a category called continuous local search (CLS) by showing that the Moreau envelop of a suitable best response function acts as a potential under certain natural gradient-based dynamics.

TLDR: Computing Nash Equilibria in adversarial team games is CLS-complete.
Toolbox: Computational Complexity, Game Theory, Convex Analysis, Linear Programming

Efficiently Computing Nash Equilibria in Adversarial Team Markov Games
(joint work with P. Kalogiannes, I. Anagnostides, S. Stavroulakis, V. Chatziafratis, Y. Panageas)
11th International Conference on Learning Representations (ICLR) **Oral** *2023*

In multi-agent reinforcement learning, predicting the optimal strategy or Nash equilibrium is crucial but challenging. In this research, we focus on games where a team, without coordinating among themselves, plays against a single opponent. Such a setup captures both competitive and cooperative elements. We present a new algorithm that can find near-optimal strategies for these games efficiently.

TLDR: In multi-agent learning, we introduce an efficient algorithm to find near-optimal strategies for teams playing against a single opponent, blending both competition and cooperation settings.
Toolbox: Reinforcement Learning, Game Theory, Non-Convex Programming

Teamwork makes von Neumann work: Min-Max Optimization in 2-Team Zero-Sum Games
(joint work with P. Kalogiannes, Y. Panageas)
11th International Conference on Learning Representations (ICLR) *2023*

In ML, two-team games are popular, especially in e-sports and multi-GANs. In these games, teams compete without coordinating internally. This study shows that common methods like GDA, OGDA, and EG can't find the best strategies. Finally, we propose novel game dynamics, inspired by control-theory, that can pinpoint an optimal strategy in specific scenario.

TLDR: In e-sports and multi-GANs ML team games, traditional methods often fail to find the best strategies. New control-theory-inspired approach that works better in certain situations.
Toolbox: Game Theory, Control Theory, Dynamical Systems

First-Order Algorithms for Min-Max Optimization in Geodesic Metric Spaces
(joint work with T. Lin, M. Jordan)
36th Conference on Neural Information Processing Systems (**NeurIPS**) **Oral** *2022*

In machine learning, we often tackle problems on curved spaces, known as Riemannian manifolds. While we have solutions for similar challenges in flat spaces, curved ones remained a puzzle. In this work, we design techniques like the Riemannian corrected extragradient (RCEG) method, whose performance guarantees match the corresponding Euclidean ones.

TLDR: Closing a 100-years open problem in this work, Von Neumann's minimax points can be computed at a linear rate in the geodesically strongly-convex-concave case, matching the Euclidean result via the Riemannian corrected extragradient (RCEG) method.
Toolbox: Differential Geometry, Game Theory, Convex Optimization, Duality Theory

On the convergence of Policy Gradient methods to Nash equilibria in General Stochastic Games
(joint work with A. Giannou, K. Lotidis, P. Mertikopoulos)
36th Conference on Neural Information Processing Systems (**NeurIPS**) *2022*

In this work, we developed a versatile framework for analyzing policy gradient methods, accommodating everything from perfect policy gradients to estimates like the REINFORCE algorithm. Within this framework, we discovered that strategically stable Nash policies are likely to emerge locally. Moreover, Nash policies with strong concave curvature can converge at a fast squared distance rate of $O(1/\sqrt{n})$. Intriguingly, with just a few modifications, we can achieve local convergence to deterministic Markov Nash equilibria in constant time, despite the game's unpredictability.

TLDR: Local Convergence Guarantees of Policy Gradient Methods in Multi-Agent RL setting.
Toolbox: Multi-Agent RL, Convex Optimization, Applied Probability, Dynamical Systems

Near-Optimal Statistical Query Lower Bounds for Agnostically Learning Intersections of Halfspaces with Gaussian Marginals
(joint work with D. Hsu, R. Servedio, C. Sanford)
35th Annual Conference on Learning Theory (**COLT**) *2022*

We investigate the challenge of learning intersections of halfspaces under Gaussian distribution in the agnostic learning model. Our contribution tightens the Statistical Query (SQ) lower bounds to $n^{-\Omega(\log k)}$. These results approach the best-known upper bounds. Additionally, our methodology applies to SQ learning for classes such as "convex subspace juntas" and sets characterized by a bounded Gaussian surface area.

TLDR: Tight lower bounds for learning intersections of halfspaces under Gaussian distribution under SQ learning.
Toolbox: Boolean Learning Theory, Fourier Analysis, High-Dimensional Probability

On the Rate of Convergence of Regularized Learning in Games:
From Bandits to Optimism and Beyond (joint work with A. Giannou, P. Mertikopoulos)
35th Conference on Neural Information Processing Systems (**NeurIPS**) *2021*

In this work, we investigate the convergence properties of "Follow the Generalized Leader" (FTGL) methods – a fusion of both classic and optimistic no-regret dynamics – towards Nash Equilibria in N-player generic games. Notably, the convergence rate of FTGL algorithms to strict Nash equilibria isn't influenced by player uncertainties but only on the geometry of the regularizer near the equilibrium.

TLDR: Rates of local convergence for FTGL methods to Nash Equilibria in generic games.
Toolbox: Game Theory, Convex Optimization, Dynamical Systems, Applied Probability

Solving Min-Max Optimization with Hidden Structure via Gradient Descent Ascent
(joint work with L. Flokas, G. Piliouras)
35th Conference on Neural Information Processing Systems (**NeurIPS**) *2021*

We investigate gradient descent ascent (GDA) dynamics in specific non-convex non-concave zero-sum games, which we term "hidden zero-sum games." In these games, players control inputs of non-linear functions feeding into a convex-concave game. When the internal game is strictly convex-concave, GDA tends to converge to its von-Neumann equilibrium. If not, GDA may not converge, but with certain techniques, we can ensure convergence. Notably, our convergence results are non-local, a unique finding for these types of games. Finally we connect our continuous dynamics findings to generative adversarial networks training.

TLDR: Identifying convergence conditions for GDA in hidden zero-sum games, highlighting links to generative adversarial networks learning.
Toolbox: Game Theory, Dynamical Systems, Lyapunov Potentials

Reconstructing of weighted voting schemes from power indices
(joint work with H. Benett, A. De, R. Servedio)
34th Annual Conference on Learning Theory (**COLT**) *2021*

Computational Social Choice focuses on analyzing various voting systems. While earlier research sought to deduce voting systems from the influence of each voter, this study dives deeper, tackling situations where only partial voter influence is known. To address this, the paper introduces two innovative methods centered on Chow parameters and Shapley indices, offering solutions to the missing data challenge in voting systems.

TLDR: Algorithms deduce voting systems from limited voter influence data.
Toolbox: Boolean Learning Theory, Fourier Analysis, Computational Social Choice

On the Approximation Power of Two-Layer Networks of Random ReLUs
(joint work with D. Hsu, R. Servedio, C. Sanford)
34th Annual Conference on Learning Theory (**COLT**) *2021*

We explore the capability of depth-two ReLU networks, initialized with random bottom-level weights, to represent smooth functions. Through meticulous analysis, we present near-matching upper and lower bounds for L_2 and Sobolev norms-approximation considering factors like the Lipschitz constant, accuracy goal, and problem's dimension.

TLDR: Tight Upper-Lower bounds for Depth-two ReLU networks, with random initial weights for Representability of smooth functions
Toolbox: Deep Learning Theory, Ridgelet Fourier Analysis, Applied Probability

Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information (joint work with A. Giannou, P. Mertikopoulos)
34th Annual Conference on Learning Theory (**COLT**) *2021*

In this research, we delve into how different learning algorithms find a Nash equilibrium, in multi-player games. Central to our exploration is the "follow the regularized leader" (FTRL) algorithm, which faces varying levels of uncertainty based on the feedback it receives. We uncover a direct relationship between the stability of a Nash equilibrium and its uniqueness: Nash equilibrium is stable and attracting under FTRL with arbitrarily high probability if and only if it is strict.

TLDR: A Nash equilibrium is stable and attracting under Follow the Regularized Leader (FTRL) algorithms. with arbitrarily high probability if and only if it is strict.
Toolbox: Game Dynamics, Stochastic Stability Analysis, Dynamical Systems

Optimal Private Median Estimation under Minimal Distributional Assumptions
(joint work with C. Tzamos, I. Zadik)
34th Conference on Neural Information Processing Systems (**NeurIPS**) **Spotlight** *2020*

Estimating the median of a distribution from limited samples, while preserving data privacy, is a foundational problem in computer science. We address this problem even for distributions with minimal assumptions and potential outliers. We devise a time-efficient, privacy-preserving algorithm that achieves this optimal rate for median's estimation, leveraging a versatile Lipschitz Extension Lemma.

TLDR: Optimal Sample and Efficient algorithm for private median estimation.
Toolbox: Differential Privacy, Algorithmic Statistics, Dynamic Programming

No-regret learning and mixed Nash equilibria: They do not mix
(joint work with L. Flokas, T. Lianeas, P. Mertikopoulos, G. Piliouras)
34th Conference on Neural Information Processing Systems (**NeurIPS**) **Spotlight** *2020*

In our study of N-player games, we explored the convergence patterns of no-regret learning dynamics towards Nash equilibria. While it's recognized that play frequency gravitates towards coarse correlated equilibria, the relationship with Nash equilibria remains ambiguous. We proved that using "follow-the-regularized-leader" dynamics, mixed Nash equilibria and no-regret learning are fundamentally misaligned –only strict Nash equilibria reliably surface as stable endpoints in such scenarios.

TLDR: Mixed NE are unstable under no-regret dynamics in any N-player games.
Toolbox: Game Theory, Dynamical Systems, Invariant Measures Analysis, Topology

Smoothed Complexity of Local Max-Cut and Binary Max-CSP
(joint work with X. Chen, C. Guo, M. Yannakakis, X. Zhang)
51st ACM Symposium on Theory of Computing (**STOC**) *2019*

Local search is a widely used method for tackling optimization problems. However, even though it often works well in real-world scenarios, theoretical evaluations sometimes suggest it could be very slow. This contradiction arises because these theoretical tests often involve unlikely, extreme cases. In this study, we discovered that when minor unpredictable changes (or "noise") are introduced to the problem input, algorithms like FLIP methods efficiently handle a range of binary optimization challenges.

TLDR: Despite theoretical concerns, local search algorithms based on single flips handle provably real-world optimization tasks efficiently.
Toolbox: Smoothed Analysis, Applied Probability, Combinatorics, Graph Theory

Poincaré Recurrence, Cycles, and Spurious Equilibria in Gradient-Descent-Ascent for Non-Convex Non-Concave Zero-Sum Games
(joint work with L. Flokas, G. Piliouras)
33th Conference on Neural Information Processing Systems (**NeurIPS**) **Spotlight** *2019*

We analyze the behavior of standard Gradient Descent-Ascent dynamics in the context of hidden bilinear games. Players in these games control the inputs of a function tied to a bilinear zero-sum game, similar to the competition in Generative Adversarial Networks. Our findings reveal that GDA dynamics in these games might not always converge to the game-theoretically meaningful solution. Instead, they can display recurrent behaviors or even converge to spurious non-min-max equilibria.

TLDR: In several ML non-convex scenarios, standard GDA dynamics show behaviors that don't converge to expected game solutions, including periodic behavior and convergence to non-min-max equilibria.
Toolbox: GANs, Game Dynamics, Invariant Measures, Lyapunov Potentials, Topology

Efficiently avoiding saddle points with zero order methods: No gradients required
(joint work with L. Flokas, G. Piliouras)
33th Conference on Neural Information Processing Systems (**NeurIPS**) *2019*

Pushing the boundaries of gradient-free optimization in machine learning, we studied methods that don't need gradient information, only function values. Designing a novel "zero-order method," we efficiently reach good solutions and avoid problematic saddle points. Our method matches the speed of traditional approaches without introducing unnecessary complexities seen in earlier approaches.

TLDR: A simple derivative-free method efficiently avoids saddle points.
Toolbox: Non-convex Optimization, Applied Probability, Dynamical Systems

Pattern Search Multidimensional Scaling
(joint work with G. Paraskevopoulos, E. Tzinis, A. Potamianos)
Submitted to Transactions on Machine Learning Research (**TMLR**)

Nonlinear manifold learning is crucial in ML for unveiling hidden data structures. We offer a new twist on the classic multi-dimensional scaling (MDS) method by utilizing derivative-free optimization techniques. Specifically, instead of the usual gradient descent, we sample and evaluate potential "moves" within a set sphere for every point in the embedded space.

TLDR: Derivative-free methods for multi-dimensional scaling via of complex data structures.
Toolbox:Manifold Learning, Non-convex & Black-box Optimization, Local Search Heuristics