



Panayotis Mertikopoulos

CNRS Researcher (*chargé de recherche*)
Laboratoire d'Informatique de Grenoble (LIG)
700 av. Centrale, Domaine Universitaire
38401 Saint-Martin d'Hères
Grenoble, France

e-mail: panayotis.mertikopoulos@imag.fr
www: <http://polaris.imag.fr/panayotis.mertikopoulos>

CURRICULUM VITÆ

PANAYOTIS MERTIKOPOULOS

May 15, 2021

CONTENTS

1	PERSONAL INFORMATION	2
2	EDUCATION	2
3	PROFESSIONAL EXPERIENCE	2
4	DISTINCTIONS AND AWARDS	2
5	RESEARCH INTERESTS	3
5.1	Game theory, online optimization, and multi-agent learning	3
5.2	Data science, machine learning, and artificial intelligence	3
5.3	Computer networks and distributed systems	4
6	ADMINISTRATIVE DUTIES AND SERVICE TO THE COMMUNITY	4
6.1	Coordination, management, and administrative responsibilities	4
6.2	Editorial activities and program chairing	5
6.3	Conference organization	6
6.4	Recruitment committees	6
6.5	Thesis committees	6
7	RESEARCH SUPERVISION AND TEACHING	7
7.1	Post-doctoral fellows	7
7.2	PhD students	7
7.3	Teaching	8
8	GRANTS AND COLLABORATIONS	8
8.1	Awarded grants (above 80 K€)	8
8.2	Awarded grants (below 80 K€)	9
8.3	Participation in research projects and networks	9
8.4	Mobility and scientific stays abroad	10
9	INVITED TALKS AND TUTORIALS	10
10	REFERENCES	13
11	PUBLICATIONS AND SCIENTIFIC OUTPUT	13

1. PERSONAL INFORMATION

Name / Surname Panayotis Mertikopoulos
Birthplace / date Athens, Greece, 10/10/1980
Nationality Greek
E-mail panayotis.mertikopoulos@imag.fr
Home page <http://polaris.imag.fr/panayotis.mertikopoulos>



2. EDUCATION

2019 **Université Grenoble Alpes** Grenoble, FR
Habilitation à Diriger des Recherches (HDR)
 Topic: *Online Optimization and Learning in Games: Theory and Applications*
 Referees (hyperlinks lead to reports): Jérôme Bolte, Nicolò Cesa-Bianchi, Sylvain Sorin

2007–2010 **University of Athens, Department of Physics** Athens, GR
Doctorate of Philosophy (PhD)
 Topic: *Stochastic Perturbations in Game Theory and Applications to Networks*
 Supervised by Aris L. Moustakas

2003–2006 **Brown University, Department of Mathematics** Providence, RI, USA
Master of Science in Mathematics with a GPA of 4.0/4.0 (May 2005)
Adm. Cand. for the PhD degree in Mathematics (M.Phil. equivalent; Sept. 2005)

1998–2003 **University of Athens, Department of Physics** Athens, Greece
Ptychion degree in Physics (July 2003); graduated valedictorian with a grade of 9.1/10
 Major topic: *Gauss's law and Residue Calculus in the Framework of de Rham Cohomology*

3. PROFESSIONAL EXPERIENCE

2011–present **CNRS – French National Center for Scientific Research** Grenoble, FR
 Tenured researcher (*chargé de recherche*) in the Laboratoire d'Informatique de Grenoble

2019, fall **École Polytechnique Fédérale de Lausanne (EPFL)** Lausanne, CH
 Visiting scholar at the Laboratory for Information and Inference Systems (LIONS)

2018, spring **UC Berkeley – Simons Institute for the Theory of Computing** Berkeley, CA, USA
 Visiting scientist for the spring semester of 2017-2018

2016, fall **LUISS Guido Carli University** Rome, Italy
 Visiting professor for the fall semester of 2016-2017

2010–2011 **École Polytechnique, Department of Economics** Paris, FR
 Post-doctoral researcher in game theory

4. DISTINCTIONS AND AWARDS

2021 Long talk at ICML 2021 for [52]
 2020 Nominated for the bronze medal of the CNRS
 2020 Two spotlight awards at NeurIPS 2020 for [54, 56]
 2020 Spotlight award at ICLR 2020 for [62]
 2018 Outstanding reviewer award at NeurIPS 2018

2017	Finalist for the INFORMS George Nicholson award for [7]
2012	Best paper award at NetGCoop 2012 for [105]
2003	Valedictorian at the Physics Department of the University of Athens

5. RESEARCH INTERESTS

The central theme of my research is the **design and analysis of adaptive algorithms for data-driven systems** – ranging from large-scale computer networks to adversarial neural nets in artificial intelligence. In this regard, my research work is aligned along the following broad axes:

1. Optimization, online learning, and game theory
2. Data science and machine learning
3. Networks and distributed systems

Much of my work is theoretical but otherwise strongly motivated and nurtured by real-world problems. In particular, I make a strong point of collaborating closely with practitioners to provide a sanity check for my modeling assumptions and to ensure the applicability of the theoretical tools developed.

5.1. GAME THEORY, ONLINE OPTIMIZATION, AND MULTI-AGENT LEARNING

My research in game-theoretic learning revolves around the following fundamental question:

Does multi-agent online learning lead to a stable, equilibrium state?

My work on this question has been aligned along two basic axes: **continuous-** vs. **discrete-time** considerations. From a practitioner's viewpoint, the latter is often more interesting than the former: a differential equation can hardly be considered an implementable control algorithm, and one could argue that modeling computer-aided decision processes as a continuous-time system is folly (at least, on the surface). However, from an analytical standpoint, the continuous- and discrete-time approaches comprise two synergistic research thrusts that dovetail in a unique and singular manner: viewing an iterative learning algorithm as a discrete version of a continuous dynamical system sheds new light on the properties of the algorithm, offers tools for its asymptotic analysis, and suggests new classes of algorithms altogether.

Thanks to the theory of stochastic approximation, insights gained in continuous time can be used to prove discrete-time results that would otherwise be inaccessible. This has allowed my work on **continuous-time learning models** [21–24, 26, 30, 36, 39, 41, 45, 79, 86, 87] to serve as the linchpin for developing the intuition and analytical tools required to attack bona fide, **discrete-time learning models**. Using these tools and techniques, I have worked extensively on **algorithms for multi-agent learning and control**, with a special focus on identifying broad classes of games where **learning converges to equilibrium** [20, 25, 29, 39, 62, 64, 65, 67, 83, 84, 87], and **overcoming learning impediments** like feedback loss, delays, asynchronicities, and the like [39, 65, 66, 73, 74, 81, 83, 89].

Much of this work has been carried out as part of the **ANR ORACLESS** project (which I am coordinating), the **EU COST action GAMENET** (where I act as working group coordinator), and the **ANR GAGA** project (where I participated as a co-PI). This was also the theoretical backbone of my **Habilitation à Diriger des Recherches (HDR)**, which I successfully defended in 2019 [2].¹

5.2. DATA SCIENCE, MACHINE LEARNING, AND ARTIFICIAL INTELLIGENCE

A recent and very active thread of my research focuses on the applications of optimization theory and game theory to deep learning and artificial intelligence systems. The adversarial match-up of deep learning mechanisms – typically in the form of *generative adversarial networks* (GANs) – has led to extraordinary advances in the field of artificial intelligence, not the least of which is the ability to pass a specific version

¹The referee reports can be found at <http://polaris.imag.fr/panayotis.mertikopoulos/publications/d03-hdr>.

of Turing’s test (the automatic generation of images that can fool a human observer). However, despite the highly promising results they provide, our theoretical understanding of GANs is still at an embryonic stage: the community has a partial idea of “what” works in practice, but not the “why” or the “how”. My contributions in the field have provided novel insights on:

1. Improving the performance of GANs and deep learning architectures through the use of asynchronous, distributed computing models [12, 16, 66, 74, 77, 82]
2. Cycling and non-convergent behaviors that arise in GAN training, and how to overcome such obstacles through the use of “optimistic”, extra-gradient techniques [52, 54, 62, 64, 65, 65, 67, 75, 79].

Much of this work has been carried out in the context of (i) my participation in the **Grenoble Alpes Multidisciplinary Institute in Artificial Intelligence** (where I am part of the “Optimization & Learning” endowed chair coordinated by J. Malick); (ii) my collaboration with the **Computer Network Architecture and Performance Engineering Lab of Stanford University** (where I mentored Z. Zhou); and (iii) my ongoing collaboration with the **Laboratory for Information and Inference Systems at EPFL**.

5.3. COMPUTER NETWORKS AND DISTRIBUTED SYSTEMS

I have been working on the applications of learning and game theory to network science and distributed systems since my PhD, and I have continued to be actively involved in the community throughout my tenure in the CNRS. In more detail, much of my work on network system design has focused on the following essential components of next-generation networks:

1. **The Internet of Things (IoT)**: here, the use of game-theoretic and learning methodologies provide scalable algorithmic tools with dramatic runtime reductions and adaptability characteristics relative to conventional optimization methods [6, 18, 28, 70–72, 78, 85, 90–92].
2. **Massive MIMO** (multiple-input and multiple-output) arrays in wireless communications: here, advanced optimization and learning methodologies can lead to significant improvements in spectrum utilization and energy efficiency [10, 15, 25, 32, 34, 35, 40, 44, 63, 89, 95, 99, 102, 103, 106].
3. **Dynamic spectrum access**: here, game-theoretic techniques can yield considerable performance gains in terms of overall throughput and power utilization [27, 37, 38, 42, 46, 76].

Much of this work has been carried out as part of the **ANR ORACLESS** project (which I am coordinating), the **Huawei FLAGSHIP project ULTRON** (which I coordinated until its end in 2018), and the **ANR NETLEARN** project (where I participated as a partner).

6. ADMINISTRATIVE DUTIES AND SERVICE TO THE COMMUNITY

6.1. COORDINATION, MANAGEMENT, AND ADMINISTRATIVE RESPONSIBILITIES

I make a point to be actively involved in the collective needs of the community at all possible levels. At the EU/international level, I am endeavoring to promote game-theoretic research and the role of the French game theory school through the European Game Theory Network; at the national level, I am trying to facilitate exchanges between researchers in decision sciences and optimization through my mandate in the SMAI-MODE steering committee; and, at the local level, by providing assistance and guidance to my lab’s PhD students as a graduate students liaison. I am detailing these activities below.

2017–present **GAMENET** **Working group coordinator and core group member for GAMENET**, the European Network for Game Theory (EU COST Action CA 16228; <https://gametheorynetwork.com>). GAMENET’s aim is to facilitate collaborations between different groups of game theorists in Europe (computer scientists, applied mathematicians, economists, etc.), and to provide game-theoretic expertise to industrial stakeholders. To that end, it brings together almost 200 researchers from 35 countries and fosters interdisciplinary cross-pollination

by providing funding for meetings, thematic schools, and short-term scientific missions (STSMs).

My responsibilities as WG coordinator and core group member is to balance the financial co-sponsorship of WG-related scientific events, and to ensure a healthy equilibrium between meetings and STSMs. This requires close coordination with GAMENET's core group and other WGs, in the form of constant communication with the action's chair, regular core group meetings, and biannual "townhall" meetings with a quorum of GAMENET's management committee (where I represent France).

2014–2020
SMAI-MODE **Member of the steering committee (*comité de liaison*) of SMAI-MODE**, the optimization and decision theory group ("*groupe thématique*") of the French Society for Industrial and Applied Mathematics (SMAI). SMAI is the foremost French scientific society in the area of applied mathematics, and its goal is to contribute to the development of the discipline and its applications. SMAI-MODE in particular is dedicated to the fields of operations research, optimization, game theory, discrete mathematics, control theory, decision sciences, and, especially, **the interface between these fields and computer science**.

The *comité de liaison* of SMAI-MODE consists of 12 elected members with a limit of up to two successive mandates (I was elected in 2014 and subsequently re-elected in 2017). Through regular meetings, it coordinates the activities of the SMAI in the above areas, and it is responsible for balancing the budgetary needs of sponsored or co-sponsored events in France. During my first mandate, I proposed a funding instrument for strengthening the representation of the group's activities in thematic mini-symposia (e.g., at the biannual SMAI congress). During my second mandate, I co-organized (jointly with J. Malick) the SMAI-MODE conference in 2018 (<http://smai-mode2018.imag.fr>), a biannual event that brought together approx. 100 researchers from different disciplines to Grenoble.

2011–2019
CHMI-DOC **Graduate students liaison (*chargé de mission doctorants*)** at the Laboratoire d'Informatique de Grenoble. The role of the "mission doctorants" is to provide assistance and support to the LIG's PhD students throughout their tenure at the lab, from the moment they begin their doctoral studies to the moment they graduate (and beyond). I sought to play an active part in this effort from the moment that I joined the LIG in 2011.

During my mandate, the number of PhD students at the LIG fluctuated around 90–100, and getting to know and interacting with them was an immensely rewarding experience. Part of my duties as a graduate students liaison was to provide guidance to new students in navigating bureaucratic or other difficulties, to animate the PhD community with annual events (such as the "welcome day" for incoming PhD students and the "LIG PhD day" which provided a forum for 2nd year PhD students to present their work), the organization of mini-workshops (e.g., dedicated to scientific writing), and, events "just for fun" such as the "café doctorants" and outdoor activities coordinated with LIG's student associations.

6.2. EDITORIAL ACTIVITIES AND PROGRAM CHAIRING

Assoc. editor Operations Research Letters (ORL; 2021–present)
RAIRO Operations Research (2021–present)
EURO Journal on Computational Optimization (EJCO; 2020–present)
Journal on Dynamics and Games (JDG; 2018–present)
Methodology and Computing in Applied Probability (MCAP; 2019–present)

PC chair	Area chair for <i>ICLR 2021</i> Area chair for <i>NeurIPS 2019–2021</i> TPC co-chair for <i>NetGCoop 2020</i> (postponed) TPC co-chair for <i>WiOpt 2014</i>
Reviewer	<i>Advances in Applied Probability, Annals of Operations Research, Dynamic Games and Applications, Games and Economic Behavior, IEEE Access, IEEE Journal on Selected Areas in Communications, IEEE Transactions on Information Theory / Signal Processing / Communications / Wireless Communications, IEEE/ACM Transactions on Networking, Journal of Economic Theory, Journal of Optimization Theory and Applications, Mathematics of Operations Research, Mathematical Programming, Operations Research, SIAM Journal on Control and Optimization, SIAM Journal on Optimization, Theoretical Economics, ...</i>
PC member	NeurIPS, ICML, ICLR, COLT, AAAI

6.3. CONFERENCE ORGANIZATION

2020	Co-organizer of the 7th workshop on “Stochastic Methods in Game Theory” (postponed)
2019	Co-organizer of the workshop “Twenty years of the Price of Anarchy”, Chania, GR
2018	Co-organizer of the 2018 Paris Symposium on Game Theory, Paris, FR
2018	General co-chair of the <i>Journées SMAI–MODE 2018</i> , Autrans, FR
2018	Co-organizer of the 2018 Workshop on Games, Dynamics and Optimization, Vienna, AT
2016	General co-chair of the <i>2016 Workshop on Games, Evolution and Learning</i> , Luchon, FR
2015	Organizer of the mini-symposium “Games, Learning and Applications” at SMAI 2015
2014	Co-organizer of the track “Dynamics and Learning in Games” in IFORS 2014
2013	General co-chair of the 2013 Workshop on Algorithmic Game Theory, Grenoble, FR
2013	Publications chair of <i>WiOpt 2013</i> , Tsukuba, JP
2012	Publications chair of <i>Valuetools 2012</i> , Cargèse, FR

6.4. RECRUITMENT COMMITTEES

2021	Université Gustave Eiffel Member of selection committee (CoS) for an MCF position on “Sciences des données, Apprentissage, Méthode mathématiques”
2020	Université Grenoble-Alpes / Grenoble IAE Member of selection committee (CoS) for an MCF position on “Intelligence des données: de l'extraction d'information à l'aide à la décision”

6.5. THESIS COMMITTEES

2020	Ya-Ping Hsieh [recipient of the EPFL EDEE thesis award] EPFL (<i>rapporteur</i>) Topic: “Convergence without convexity: Sampling, optimization, and games”
2020	Xavier Fontaine U. Paris-Saclay (<i>rapporteur</i>) Topic: “Sequential learning and stochastic optimization of convex functions”
2020	Rafael Pinot U. Paris-Dauphine (<i>rapporteur</i>) Topic: “On the impact of randomization on robustness in machine learning”

- 2018 **Adil Salim** TELECOM ParisTech (*examineur*)
Topic: “*Random monotone operators and application to stochastic optimization*”
- 2014 **Tatiana Seregina** U. Toulouse III (*examineur*)
Topic: “*Applications of game theory to distributed routing and delay-tolerant networking*”

7. RESEARCH SUPERVISION AND TEACHING

7.1. POST-DOCTORAL FELLOWS

- 2020–present **Dong Quan Vu** [funding: ANR ALIAS]
Topic: “*Optimization methods for online path planning problems*”
Joint output: 1 submitted paper
- 2018–2020 **Olivier Bilenne** [funding: ANR ORACLESS]
Topic: “*Gradient-free optimization with applications to MIMO systems*”
Joint output: 2 papers [15, 63]
Position after post-doc: post-doctoral researcher at Maastricht University
- 2017–2018 **Amélie Héliou** [funding: Huawei ULTRON grant]
Topic: “*Multi-agent online learning with bandit feedback*”
Joint output: 4 papers [55, 59, 81, 84], 1 submitted
Position after post-doc: senior researcher at Criteo AI Lab
- 2017–2018 **Luigi Vigneri** [funding: Huawei ULTRON grant]
Topic: “*Scalable latency minimization algorithms for next-generation networks*”
Joint output: 1 paper [70]
Position after post-doc: senior research scientist at the IOTA Foundation
- 2014–2015 **Ioannis Stiakogiannakis** [funding: ANR NETLEARN]
Topic: “*Online optimization algorithms for dynamic MIMO systems*”
Joint output: 2 papers [95, 97], 1 preprint [10]
Position after post-doc: senior machine learning scientist at Huawei, then Cisco Systems
- 2014–2015 **Nof Abuzainab** [funding: Inria fellowship]
Topic: “*Game-theoretic analysis of cognitive radio systems*”
Position after post-doc: Virginia Tech

7.2. PHD STUDENTS

- 2019–present **Yu-Guan Hsieh** [funding: Grenoble MIAI]
Co-supervised with J. Malick (DR CNRS / LJK) and F. Iutzeler (MCF UGA)
Topic: “*Extra-gradient methods for variational inequalities and machine learning*”
Joint output: 3 papers [5, 54, 65], 3 submitted
- 2018–present **Benjamin Roussillon** [funding: MESR fellowship]
Co-supervised with P. Loiseau (CR Inria / LIG)
Topic: “*Classification en présence de données adverses : modèles et solutions*”
Joint output: 1 paper submitted
- 2017–present **Kimón Antonakopoulos** [funding: ANR ORACLESS]
Co-supervised with E. V. Belmega (MCF ENSEA / U. Cergy-Pontoise)
Topic: “*Online learning for variational inequality problems*”
Joint output: 3 papers [53, 62, 64], 3 submitted

- 2017–2020 **Bruno Donassolo** [funding: CIFRE fellowship]
 Co-supervised with A. Legrand (DR CNRS / LIG) and I. Fajjari (Orange)
 Topic: “Decentralized management of applications in Fog computing environments”
 Joint output: 3 papers [6, 71, 72]
 Position after PhD: Inria research engineer
- 2015–2019 **Alexandre Marcastel** [funding: MESR fellowship]
 Co-supervised with E. V. Belmega (MCF ENSEA / U. Cergy-Pontoise)
 Topic: “Allocation de puissance en ligne dans un réseau IoT dynamique et non-prédictible”
 Joint output: 5 papers [18, 68, 90–92]
 Position after PhD: Professeur agrégé

7.3. TEACHING

- Graduate **EPFL**: optimization for data science (*cours magistral*, 2019)
ENS Lyon: game theory, learning, optimization (*cours magistral*, 2017–2018)
- Undergrad **U. Athens**: stochastic processes, advanced algorithms (2014–present)
- Schools **French Days on Optimization and Decision Science (2020)**: *Algorithmic game theory: From multi-agent optimization to online learning* (fall 2020; [available online](#))
Trinity College Dublin (2019): *Machine learning for communications* (fall 2019)
UC Berkeley (2018): *Real-time decision-making* (spring 2018)
RESCOM (2012): Applications of Game Theory to Data Networks (spring 2012)

8. GRANTS AND COLLABORATIONS

8.1. AWARDED GRANTS (ABOVE 80 K€)

- 2020–2024 **ANR ALIAS – Adaptive learning for interacting agents and systems**
 French National Research Agency intl’ grant (ANR PRCI); **co-PI** **Budget: 284 k €**
 This project is a bilateral PRCI grant between France and Singapore, jointly funded by the ANR and Singapore’s National Research Foundation (NRF) with matching funds on either side. The project was initiated by B. Pradelski and myself on the French side (with a co-PI structure), and G. Piliouras and I. Panageas from the Singapore University of Technology and Design (SUTD). The project was built around an existing collaboration between G. Piliouras and myself and its aim is to determine whether adaptive learning leads to stable outcomes in multi-agent systems.
- 2016–2020 **ANR ORACLESS – Online resource allocation for unpredictable large-scale wireless systems**
 French National Research Agency starting grant (ANR JCJC); **PI** **Budget: 207 k €**
 This project is an ANR starting grant aiming to develop highly adaptive resource allocation schemes that are provably capable of tracking unpredictable changes in communication networks. As the project’s PI, I am coordinating its research activities and scientific output.
- 2017–2018 **ULTRON – Ultra-low latency scheduling via online learning**
 Huawei FLAGSHIP grant; **PI** **Budget: 195 k €**
 ULTRON was a Huawei Flagship project aiming to develop highly adaptive learning policies for achieving ultra-low latencies in 5G mobile systems. Owing to the applications of my work on multi-agent learning to computer networks and communications, I was invited to submit this proposal in 2016. The project ran successfully between 2017 and 2018 and led to novel resource allocation schemes for software-defined networks. L. Vigneri and A. Héliou were recruited as post-doctoral fellows to work on different work-packages of the ULTRON project.

- 2017–2021 **GAMENET** – *European Network for Game Theory*
 COST action; working group coordinator & core group member **Budget: 625 k €**
 The European Network for Game Theory is an EU COST action initiated by M. Staudigl (Maastricht University, the Netherlands, action chair), M. Scarsini (LUISS, Rome, action vice-chair), and myself. It is bringing together almost 200 researchers from 35 countries and it seeks to foster interdisciplinary cross-pollination in game theory.
- 2014–2017 **ANR GAGA** – *Geometric aspects of games*
 ANR grant; **co-PI** **Budget: 85 k €**
 This project was an ANR starting grant aiming to explore the role of geometric structures in game theory and learning. It was coordinated by V. Perchet with a “primus inter pares”, co-PI structure (typical of ANR starting grants in mathematics). I was responsible for coordinating the activities of Work-Package 2: “Geometric algorithms and dynamics for learning in games”.

8.2. AWARDED GRANTS (BELOW 80 K€)

- 2020–2021 **DISCMAN** – *Distributed control for multi-agent systems and networks*
 “Investissements d’avenir” project (ANR-IDEX); **PI** **Budget: 10 k €**
- 2017–2018 **HEAVY.NET** – *Optimization and analysis of heavily congested networks*
 PGM0/PRMO grant; **PI** **Budget: 12 k €**
- 2018 **MixedGAN** – *Mixed-strategy generative adversarial networks*
 CNRS exploratory grant (PEPS I3A); participant **Budget: 10 k €**
- 2016 **REAL.net** – *Resource allocation in dynamic network environments*
 CNRS exploratory grant (PEPS JCJC); **PI** **Budget: 10 k €**
- 2014–2015 **GATHERING** – *Game theory, evolution and randomness in networks and graphs*
 CNRS exploratory grant (PEPS HuMaIn); **PI** **Budget: 10 k €**
- 2012–2013 **LACODS** – *Learning algorithms for control and optimization in distributed systems*
 MSTIC (French competitiveness pole) career development grant; **PI** **Budget: 10 k €**

8.3. PARTICIPATION IN RESEARCH PROJECTS AND NETWORKS

- 2016–2018 **LEARN** – *Learning algorithms for games and applications*
 Franco-Chilean Network of Excellence, co-financed by ECOS-Sud and CONICYT
- 2013–2017 **NETLEARN** – *Learning algorithms orchestration for mobile networks resource management*
 Research project financed by the French National Research Agency (ANR)
- 2012–2015 **NEWCOM#** – *Network of excellence in wireless communications*
 Network of Excellence formed under FP7
- 2012–2016 **ADGO** – *Algorithms and dynamics in games and optimization*
 Franco-Chilean network funded by the Chilean National Research Agency (CONICYT)
- 2012–2016 **CROWN** – *Optimal control of self-organized wireless networks*
 Research project co-funded by Greece and the EU under the THALES initiative
- 2006–2009 **NET-REFOUND** – *Network research foundations and trends*
 Specific Targeted Research Project funded by the EU under FP6

8.4. MOBILITY AND SCIENTIFIC STAYS ABROAD

- 2019 **École Polytechnique Fédérale de Lausanne (EPFL)** Lausanne, CH
I spent the fall semester of 2019-2020 at the Laboratory for Information and Inference Systems (LIONS lab) of EPFL, hosted by V. Cevher (the lab's director). This stay kicked off a very fruitful collaboration, with several joint research projects and exchanges planned in the near future. While at EPFL, I also taught a PhD-level seminar course on mathematical optimization for data science and machine learning.
- 2018 **UC Berkeley** Berkeley, CA, USA
I spent most of the 2017-2018 spring semester at the Simons Center for the Theory of Computing at UC Berkeley. I was invited to participate in the “Real-Time Decision-Making” thematic program as an expert on multi-agent and game-theoretic learning.
- 2017 **Stanford Unviversity** Palo Alto, CA, USA
I spent a month in the 2016-2017 spring semester at the Computer Network Architecture and Performance Engineering Lab at Stanford University. This was part of an ongoing collaboration with N. Bambos and the beginning of my mentoring (unofficial supervision) of Z. Zhou.
- 2016 **LUISS Guido Carli University** Rome, Italy
I spent a month in the 2016-2017 fall semester at the Operations Research department of the LUISS Guido Carli University. This was part of an ongoing collaboration project with M. Scarsini and the starting point of the GAMENET proposal.

9. INVITED TALKS AND TUTORIALS

I am regularly invited to give talks in conferences, workshops and top-tier research institutions / universities. I am providing below an itemized list of invited talks and tutorials.

- 2021 **RWTH Aachen – Mathematics and Information Processing Seminar** Aachen, DE
“Generalized Robbins-Monro algorithms for min-min and min-max optimization”
- 2021 **Télécom ParisTech – Signal, Statistics & Learning Seminar** Paris, FR
“Online optimization: A unified view through the lens of stochastic approximation”
- 2021 **TSE – MAD-Stat Seminar** Toulouse, FR
“Dynamics, (min-max) optimization, and games”
- 2021 **Montréal Machine Learning and Optimization Seminar** Montréal, CA
“Spurious attractors in min-max optimization”
- 2020 **National Technical University of Athens (CoreLab seminar)** Athens, GR
“Games, dynamics, and spurious attractors”
- 2020 **French Days on Optimization and Decision Science (invited tutorial)** Paris, FR
“Algorithmic game theory: from multi-agent optimization to online learning”
- 2020 **One Wolrd Optimization Seminar / One World Game Theory Seminar** Virtual
“Games, Dynamics, and Optimization”
- 2020 **GDO 2020 – Games, Dynamics and Optimization** Rome, IT
“Learning in time-varying games”
- 2020 **Paris Game Theory Seminar (Institut Henri Poincaré)** Paris, FR
“Learning with scarce feedback”
- 2019 **LUISS Guido Carli University** Rome, IT
“From Hannan to Nash: cycles, learning, and equilibrium”

2019	CONNECT Summer School on Machine Learning for Communications "Online learning and optimization for wireless systems"	Dublin, IE
2019	NPCG 2019 – Network, Population and Congestion Games "No-regret learning in games"	Paris, FR
2019	GDO 2019 – Games, Dynamics and Optimization "Hessian barrier algorithms for linearly constrained optimization problems"	Cluj-Napoca, RO
2019	OSL 2019 – Optimization and Statistical Learning "Extra-gradient methods for variational inequalities"	Les Houches, FR
2019	École polytechnique fédérale de Lausanne (EPFL) "Going the extra (gradient) mile in GAN training"	Lausanne, Switzerland
2019	Criteo AI Lab "Applications of multi-agent learning to computational advertising"	Paris, FR
2018	PGMO Days 2018 "Learning dynamics for routing problems"	Paris, FR
2018	Trinity College "Efficient network utility maximization algorithms"	Dublin, IE
2018	National Technical University of Athens (Athens Polytechnic) "Traffic in congested networks: Equilibrium, efficiency, and dynamics"	Athens, GR
2018	GDO 2018 – Games, Dynamics and Optimization "Bandit learning in concave N-person games"	Vienna, AT
2018	Google Inc. "Accelerated and optimistic methods for learning"	Mountain View, CA, USA
2018	UC Berkeley – Simons Institute for the Theory of Computing "Online learning in games"	Berkeley, CA, USA
2017	University of Aix-Marseille "Convergence and non-convergence in game-theoretic learning"	Marseille, FR
2017	PGMO Days 2017 "The price of anarchy in high and low traffic"	Paris, FR
2017	GDR ISIS workshop on Game Theory, Optimization and Learning "Game theory meets signal processing (and feels no regret)"	Paris, FR
2017	Emergent and Self-Adaptive Systems Workshop – panelist "Design and validation of future computer systems: Theory and practice"	Lancaster, UK
2017	Lancaster University "Multi-agent online learning: Game theory meets machine learning"	Lancaster, UK
2017	Paris Game Theory Seminar (Institut Henri Poincaré) "No-regret learning in games"	Paris, FR
2017	Erice 2017 – Stochastic Methods in Game Theory "How bad is selfish routing in highly congested networks?"	Erice, IT
2017	Stanford University "Learning in games via reinforcement and regularization"	Stanford, CA, USA
2016	University of Vienna "On the convergence of gradient flows with noisy gradient input"	Vienna, AT
2016	Saclay Algorithmics Seminar (Université Paris-Sud) "Learning in games with continuous action spaces"	Paris, FR

2016	Sapienza University of Rome “Game-theoretic learning with noisy first-order input”	Rome, IT
2016	LUISS Guido Carli University “Learning in games with imperfect information”	Rome, IT
2016	Paris Optimization Seminar (Institut Henri Poincaré) “Learning in concave games”	Paris, FR
2016	CROWNCOM 2016 (two-part tutorial) “Game theory, learning, and cognitive radio”	Grenoble, FR
2016	ADGO 2016 – Algorithms and Dynamics for Games and Optimization “Robust optimization and online learning in games”	Santiago, CL
2015	University of Maastricht “Learning in concave N -person games”	Maastricht, The Netherlands
2015	ORANGE workshop on Learning and Networks (invited tutorial) “Online optimization for wireless communication systems”	Paris, FR
2015	2015 Paris Symposium on Game Theory “Learning in games with unknown payoff functions”	Paris, FR
2015	Huawei Algorithmic Sciences Laboratory “Online optimization for wireless networks”	Paris, FR
2015	PGMO Conference on Optimization & Practices in Industry “Geometric game dynamics and regularization”	Paris, FR
2015	Toulouse School of Economics “Regularization methods for learning in games”	Toulouse, FR
2014	MAGTA 2014 – Mathematical Aspects of Game Theory “Geometric aspects of game dynamics”	Roscoff, FR
2014	University of Wisconsin–Madison “Regularized best responses and reinforcement learning in games”	Madison, WI, USA
2014	University of Athens, Department of Mathematics “Robust learning in games”	Athens, GR
2013	University of Avignon “Transmit without regrets: online optimization for MIMO cognitive radio”	Avignon, FR
2013	ADGO 2013 – Algorithms and Dynamics for Games and Optimization “Inertial game dynamics and applications to constrained optimization”	Playa Blanca, CL
2013	Erice 2013 – Stochastic Methods in Game Theory “Robust learning in games”	Erice, IT
2013	Hausdorff Research Institute for Mathematics “Entropy-driven game dynamics”	Bonn, DE
2013	SMAI 30 years conference – “Games” Symposium “Hessian Riemannian game dynamics”	Seignosse, FR
2013	University of Athens “Strange bedfellows: Riemann, Gibbs, and vector multiple access channels”	Athens, GR
2013	Universidad de Chile “Robust entropy-driven learning procedures in games”	Santiago, CL

2012	Paris Game Theory Seminar (Institut Henri Poincaré) “Higher order game dynamics”	Paris, FR
2012	École Polytechnique “The role of memory in game-theoretic learning”	Paris, FR
2012	ENSEA “Matrix exponential learning: distributed optimization in MIMO systems”	Paris, FR
2011	MAGTA 2011 – Mathematical Aspects of Game Theory “Higher order game dynamics”	Toulouse, FR
2010	University of Vienna “Diffusion processes over polytopes and the stochastic replicator dynamics”	Vienna, AT
2010	Supélec (École Supérieure d’Électricité) “How robust is selfish routing?”	Paris, FR
2010	Erice 2010 – Stochastic Methods in Game Theory “The stochastic replicator dynamics: an assorted zoology”	Erice, IT
2010	BQGT 2010 – Behavioral & Quantitative Game Theory “Balancing traffic in stochastically fluctuating networks”	Newport, CA, USA
2010	Paris Game Theory Seminar (Institut Henri Poincaré) “Learning in the presence of noise”	Paris, FR
2009	Athens University of Economics and Business “The emergence of rationality in the presence of stochastic perturbations”	Athens, GR

10. REFERENCES

The following persons can be contacted to provide feedback on my research work and contributions:

1. **Nicholas Bambos:** R. Weiland Professor in the School of Engineering at Stanford University, head of the Computer Network Architecture and Performance Engineering Lab, and Chairman of the Department of Management Science & Engineering at Stanford University.
2. **Nicolò Cesa-Bianchi:** Professor of Computer Science at Università degli Studi di Milano. He is the author of the seminal book “Prediction, Learning, and Games”, and one of the world’s leading computer scientists on online machine learning and multi-armed bandits.
3. **Christos H. Papadimitriou:** Donovan Professor of Computer Science at Columbia University. Among several other distinctions and awards, he is the recipient of the IEEE von Neumann Medal in 2016, the Gödel prize in 2012 for the notion of the price of anarchy (PoA), and the Knuth Prize in 2002 .
4. **Sylvain Sorin:** Professor Emeritus at Sorbonne Université. He is the author of the highly influential book “Repeated Games”, his many contributions in the field span over four decades, and he is widely regarded as the *pater familias* of the French game theory school.

11. PUBLICATIONS AND SCIENTIFIC OUTPUT

I have co-authored 35 journal papers, 67 conference papers, and 7 more papers that are currently under review. Owing to the overall higher quality of reviews, I typically target the flagship machine learning and theoretical computer science conferences (NeurIPS, ICML, ICLR, COLT, SODA,...), the leading journals in optimization and game theory (Mathematical Programming, SIAM SIOPT / SICON, MOR, GEB / JET,...), and the relevant IEEE transactions (IEEE TSP, IEEE TAC, IEEE JSAC, and the like).

Author ordering depends on the conventions of each field and with whom the article was written. In papers where all authors are at a comparable seniority level, I prefer an alphabetical ordering unless there is a significant variance in each author's contribution. Otherwise, when the paper involves a PhD student, post-doc or junior researcher needing to build a reputation, I generally insist on their name appearing first (again, unless there is a significant variance in contributions).

SOFTWARE (1)

- [1] P. Mertikopoulos, "GameSeer: visualization software for game dynamics." Available under the GNU public license at: <http://polaris.imag.fr/panayotis.mertikopoulos/files/GameSeer.zip>, 2014.

DISSERTATIONS (3)

- [2] P. Mertikopoulos, *Online Optimization and Learning in Games: Theory and Applications*. Habilitation à Diriger des Recherches (HDR), Université Grenoble-Alpes, December 2019.
- [3] P. Mertikopoulos, *Stochastic Perturbations in Game Theory and Applications to Networks*. PhD thesis, National and Kapodistrian University of Athens, November 2010.
- [4] P. Mertikopoulos, *Gauss's Law and Residue Calculus in the Framework of de Rham Cohomology*. Major thesis, National and Kapodistrian University of Athens, May 2003.

WORKING/SUBMITTED PAPERS (7)

- [5] Y.-G. Hsieh, F. Iutzeler, J. Malick, and P. Mertikopoulos, "Multi-agent online optimization with delays: Asynchronicity, adaptivity, and optimism." <https://arxiv.org/abs/2012.11579>, 2020.
- [6] B. Donassolo, A. Legrand, P. Mertikopoulos, and I. Fajjari, "Online reconfiguration of IoT applications in the fog: The information-coordination trade-off." <https://hal.inria.fr/hal-02636987/document>, May 2020.
- [7] Z. Zhou, P. Mertikopoulos, N. Bambos, P. W. Glynn, and C. Tomlin, "Multi-agent online learning with imperfect information." Under review, 2018.
- [8] B. Duvocelle, P. Mertikopoulos, M. Staudigl, and D. Vermeulen, "Learning in time-varying games." <https://arxiv.org/abs/1809.03066>, 2018.
- [9] E. V. Belmega, P. Mertikopoulos, R. Negrel, and L. Sanguinetti, "Online convex optimization and no-regret learning: Algorithms, guarantees and applications." <https://arxiv.org/abs/1804.04529>, 2018.
- [10] I. Stiakogiannakis, P. Mertikopoulos, and C. Touati, "Power control via online learning in non-stationary MIMO networks." <http://arxiv.org/abs/1503.02155>, 2018.
- [11] P. Mertikopoulos, A. L. Moustakas, and A. Tzanakaki, "Boltzmann meets Nash: Energy-efficient routing in optical networks under uncertainty." <https://arxiv.org/abs/1605.01451>, 2016.

JOURNAL PAPERS (35)

- [12] Z. Zhou, P. Mertikopoulos, N. Bambos, P. W. Glynn, and Y. Ye, "Distributed stochastic optimization with large delays," *Mathematics of Operations Research*, to appear.
- [13] R. I. Boç, P. Mertikopoulos, M. Staudigl, and P. T. Vuong, "Minibatch stochastic forward-backward-forward methods for solving stochastic variational inequalities," *Stochastic Systems*, to appear.
- [14] Z. Zhou, P. Mertikopoulos, A. L. Moustakas, N. Bambos, and P. W. Glynn, "Robust power management via learning and game design," *Operations Research*, vol. 69, pp. 331–345, January 2021.
- [15] O. Bilenne, P. Mertikopoulos, and E. V. Belmega, "Fast optimization with zeroth-order feedback in distributed multi-user MIMO systems," *IEEE Trans. Signal Process.*, vol. 68, pp. 6085–6100, October 2020.
- [16] Z. Zhou, P. Mertikopoulos, N. Bambos, S. P. Boyd, and P. W. Glynn, "On the convergence of mirror descent beyond stochastic convex programming," *SIAM Journal on Optimization*, vol. 30, no. 1, pp. 687–716, 2020.
- [17] R. Colini-Baldeschi, R. Cominetti, P. Mertikopoulos, and M. Scarsini, "When is selfish routing bad? The price of anarchy in light and heavy traffic," *Operations Research*, vol. 68, pp. 411–434, March 2020.

- [18] A. Marcastel, E. V. Belmega, P. Mertikopoulos, and I. Fijalkow, "Online power optimization in feedback-limited, dynamic and unpredictable IoT networks," *IEEE Trans. Signal Process.*, vol. 67, pp. 2987–3000, June 2019.
- [19] I. M. Bomze, P. Mertikopoulos, W. Schachinger, and M. Staudigl, "Hessian barrier algorithms for linearly constrained optimization problems," *SIAM Journal on Optimization*, vol. 29, no. 3, pp. 2100–2127, 2019.
- [20] P. Mertikopoulos and Z. Zhou, "Learning in games with continuous action sets and unknown payoff functions," *Mathematical Programming*, vol. 173, pp. 465–507, January 2019.
- [21] P. Mertikopoulos and M. Staudigl, "Stochastic mirror descent dynamics and their convergence in monotone variational inequalities," *Journal of Optimization Theory and Applications*, vol. 179, pp. 838–867, December 2018.
- [22] P. Mertikopoulos and W. H. Sandholm, "Riemannian game dynamics," *Journal of Economic Theory*, vol. 177, pp. 315–364, September 2018.
- [23] P. Mertikopoulos and M. Staudigl, "On the convergence of gradient-like flows with noisy gradient input," *SIAM Journal on Optimization*, vol. 28, pp. 163–197, January 2018.
- [24] M. Bravo and P. Mertikopoulos, "On the robustness of learning in games with stochastically perturbed payoff observations," *Games and Economic Behavior*, vol. 103, John Nash Memorial issue, pp. 41–66, May 2017.
- [25] P. Mertikopoulos, E. V. Belmega, R. Negrel, and L. Sanguinetti, "Distributed stochastic optimization via matrix exponential learning," *IEEE Trans. Signal Process.*, vol. 65, pp. 2277–2290, May 2017.
- [26] J. Kwon and P. Mertikopoulos, "A continuous-time approach to online optimization," *Journal of Dynamics and Games*, vol. 4, pp. 125–148, April 2017.
- [27] A. S. Shafiq, P. Mertikopoulos, S. Glisic, and Y. M. Fang, "Semi-cognitive radio networks: A novel dynamic spectrum sharing mechanism," *IEEE Trans. on Cogn. Commun. Netw.*, vol. 3, pp. 97–111, March 2017.
- [28] S. D'Oro, L. Galluccio, P. Mertikopoulos, G. Morabito, and S. Palazzo, "Auction-based resource allocation in OpenFlow multi-tenant networks," *Computer Networks*, vol. 115, pp. 29–41, March 2017.
- [29] S. Perkins, P. Mertikopoulos, and D. S. Leslie, "Mixed-strategy learning with continuous action sets," *IEEE Trans. Autom. Control*, vol. 62, pp. 379–384, January 2017.
- [30] P. Mertikopoulos and W. H. Sandholm, "Learning in games via reinforcement and regularization," *Mathematics of Operations Research*, vol. 41, pp. 1297–1324, November 2016.
- [31] A. L. Moustakas, P. Mertikopoulos, and N. Bambos, "Power optimization in random wireless networks," *IEEE Trans. Inf. Theory*, vol. 62, pp. 5030–5058, September 2016.
- [32] B. Gaujal and P. Mertikopoulos, "A stochastic approximation algorithm for stochastic semidefinite programming," *Probability in the Engineering and Informational Sciences*, vol. 30, pp. 431–454, July 2016.
- [33] P. Mertikopoulos and Y. Viossat, "Imitation dynamics with payoff shocks," *International Journal of Game Theory*, vol. 45, pp. 291–320, March 2016.
- [34] P. Mertikopoulos and E. V. Belmega, "Learning to be green: Robust energy efficiency maximization in dynamic MIMO-OFDM systems," *IEEE J. Sel. Areas Commun.*, vol. 34, pp. 743 – 757, April 2016.
- [35] P. Mertikopoulos and A. L. Moustakas, "Learning in an uncertain world: MIMO covariance matrix optimization with imperfect feedback," *IEEE Trans. Signal Process.*, vol. 64, pp. 5–18, January 2016.
- [36] R. Laraki and P. Mertikopoulos, "Inertial game dynamics and applications to constrained optimization," *SIAM Journal on Control and Optimization*, vol. 53, pp. 3141–3170, October 2015.
- [37] S. D'Oro, P. Mertikopoulos, A. L. Moustakas, and S. Palazzo, "Interference-based pricing for opportunistic multi-carrier cognitive radio systems," *IEEE Trans. Wireless Commun.*, vol. 14, pp. 6536–6549, December 2015.
- [38] G. Bacci, E. V. Belmega, P. Mertikopoulos, and L. Sanguinetti, "Energy-aware competitive power allocation for heterogeneous networks under QoS constraints," *IEEE Trans. Wireless Commun.*, vol. 14, pp. 4728–4742, September 2015.
- [39] P. Coucheney, B. Gaujal, and P. Mertikopoulos, "Penalty-regulated dynamics and robust learning procedures in games," *Mathematics of Operations Research*, vol. 40, pp. 611–633, August 2015.
- [40] P. Mertikopoulos and E. V. Belmega, "Transmit without regrets: online optimization in MIMO-OFDM cog-

- native radio systems,” *IEEE J. Sel. Areas Commun.*, vol. 32, pp. 1987–1999, November 2014.
- [41] R. Laraki and P. Mertikopoulos, “Higher order game dynamics,” *Journal of Economic Theory*, vol. 148, pp. 2666–2695, November 2013.
- [42] P. Mertikopoulos, E. V. Belmega, A. L. Moustakas, and S. Lasaulce, “Distributed learning policies for power allocation in multiple access channels,” *IEEE J. Sel. Areas Commun.*, vol. 30, pp. 96–106, January 2012.
- [43] C. Pawlowitsch, P. Mertikopoulos, and N. Ritt, “Neutral stability, drift, and the diversification of languages,” *Journal of Theoretical Biology*, vol. 287, pp. 1–12, July 2011.
- [44] P. Kazakopoulos, P. Mertikopoulos, A. L. Moustakas, and G. Caire, “Living at the edge: a large deviations approach to the outage MIMO capacity,” *IEEE Trans. Inf. Theory*, vol. 57, pp. 1984–2007, April 2011.
- [45] P. Mertikopoulos and A. L. Moustakas, “The emergence of rational behavior in the presence of stochastic perturbations,” *The Annals of Applied Probability*, vol. 20, pp. 1359–1388, July 2010.
- [46] P. Mertikopoulos and A. L. Moustakas, “Correlated anarchy in overlapping wireless networks,” *IEEE J. Sel. Areas Commun.*, vol. 26, pp. 1160–1169, September 2008.

CONFERENCE PAPERS (67)

- [47] W. Azizian, F. Iutzeler, J. Malick, and P. Mertikopoulos, “The last-iterate convergence rate of optimistic mirror descent in stochastic variational inequalities,” in *COLT ’21: Proceedings of the 34th Annual Conference on Learning Theory*, 2021.
- [48] A. Giannou, E. V. Vlatakis-Gkaragkounis, and P. Mertikopoulos, “Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information,” in *COLT ’21: Proceedings of the 34th Annual Conference on Learning Theory*, 2021.
- [49] Y.-G. Hsieh, K. Antonakopoulos, and P. Mertikopoulos, “Adaptive learning in continuous games: Optimal regret bounds and convergence to Nash equilibrium,” in *COLT ’21: Proceedings of the 34th Annual Conference on Learning Theory*, 2021.
- [50] N. Hallak, P. Mertikopoulos, and V. Cevher, “Regret minimization in stochastic non-convex learning via a proximal-gradient approach,” in *ICML ’21: Proceedings of the 38th International Conference on Machine Learning*, 2021.
- [51] A. Héliou, M. Martin, P. Mertikopoulos, and T. Rahier, “Zeroth-order non-convex learning via hierarchical dual averaging,” in *ICML ’21: Proceedings of the 38th International Conference on Machine Learning*, 2021.
- [52] Y.-P. Hsieh, P. Mertikopoulos, and V. Cevher, “The limits of min-max optimization algorithms: Convergence to spurious non-critical sets,” in *ICML ’21: Proceedings of the 38th International Conference on Machine Learning*, 2021.
- [53] K. Antonakopoulos, E. V. Belmega, and P. Mertikopoulos, “Adaptive extra-gradient methods for min-max optimization and games,” in *ICLR ’21: Proceedings of the 2021 International Conference on Learning Representations*, 2021.
- [54] Y.-G. Hsieh, F. Iutzeler, J. Malick, and P. Mertikopoulos, “Explore aggressively, update conservatively: Stochastic extragradient methods with variable stepsize scaling,” in *NeurIPS ’20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [55] A. Héliou, M. Martin, P. Mertikopoulos, and T. Rahier, “Online non-convex optimization with imperfect feedback,” in *NeurIPS ’20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [56] L. Flokas, E. V. Vlatakis-Gkaragkounis, T. Lianas, P. Mertikopoulos, and G. Piliouras, “No-regret learning and mixed Nash equilibria: They do not mix,” in *NeurIPS ’20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [57] P. Mertikopoulos, N. Hallak, A. Kavis, and V. Cevher, “On the almost sure convergence of stochastic gradient descent in non-convex problems,” in *NeurIPS ’20: Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- [58] A. Alacaoglu, Y. Malitsky, P. Mertikopoulos, and V. Cevher, “A new regret analysis for Adam-type algorithms,” in *ICML ’20: Proceedings of the 37th International Conference on Machine Learning*, 2020.

- [59] A. Héliou, P. Mertikopoulos, and Z. Zhou, “Gradient-free online learning in continuous games with delayed rewards,” in *ICML ’20: Proceedings of the 37th International Conference on Machine Learning*, 2020.
- [60] T. Lin, Z. Zhou, P. Mertikopoulos, and M. I. Jordan, “Finite-time last-iterate convergence for multi-agent learning in games,” in *ICML ’20: Proceedings of the 37th International Conference on Machine Learning*, 2020.
- [61] P. Mertikopoulos, H. H. Nax, and B. S. R. Pradelski, “Quick or cheap? Breaking points in dynamic markets,” in *EC ’20: Proceedings of the 21st ACM Conference on Economics and Computation*, 2020.
- [62] K. Antonakopoulos, E. V. Belmega, and P. Mertikopoulos, “Online and stochastic optimization beyond Lipschitz continuity: A Riemannian approach,” in *ICLR ’20: Proceedings of the 2020 International Conference on Learning Representations*, 2020.
- [63] O. Bilenne, P. Mertikopoulos, and E. V. Belmega, “Derivative-free optimization over multi-user MIMO networks,” in *NetGCoop ’20: Proceedings of the 2020 International Conference on Network Games, Control and Optimization*, 2020.
- [64] K. Antonakopoulos, E. V. Belmega, and P. Mertikopoulos, “An adaptive mirror-prox algorithm for variational inequalities with singular operators,” in *NeurIPS ’19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019.
- [65] Y.-G. Hsieh, F. Iutzeler, J. Malick, and P. Mertikopoulos, “On the convergence of single-call stochastic extragradient methods,” in *NeurIPS ’19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 6936–6946, 2019.
- [66] N. Liakopoulos, A. Destounis, G. Paschos, T. Spyropoulos, and P. Mertikopoulos, “Cautious regret minimization: Online optimization with long-term budget constraints,” in *ICML ’19: Proceedings of the 36th International Conference on Machine Learning*, 2019.
- [67] P. Mertikopoulos, B. Lecouat, H. Zenati, C.-S. Foo, V. Chandrasekhar, and G. Piliouras, “Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile,” in *ICLR ’19: Proceedings of the 2019 International Conference on Learning Representations*, 2019.
- [68] A. Marcastel, E. V. Belmega, P. Mertikopoulos, and I. Fijalkow, “Gradient-free online resource allocation algorithms for dynamic wireless networks,” in *SPAWC ’19: Proceedings of the 2019 IEEE International Workshop on Signal Processing Advances in Wireless Communications*, 2019.
- [69] M. Staudigl and P. Mertikopoulos, “Convergent noisy forward-backward-forward algorithms in non-monotone variational inequalities,” in *LSS ’19: Proceedings of the 15th IFAC Symposium on Large Scale Complex Systems*, 2019.
- [70] L. Vigneri, G. Paschos, and P. Mertikopoulos, “Large-scale network utility maximization: Countering exponential growth with exponentiated gradients,” in *INFOCOM ’19: Proceedings of the 38th IEEE International Conference on Computer Communications*, 2019.
- [71] B. Donassolo, I. Fajjari, A. Legrand, and P. Mertikopoulos, “A fog-based framework for IoT service provisioning,” in *CCNC ’19: Proceedings of the 16th IEEE International Conference on Consumer Communications & Networking*, 2019.
- [72] B. Donassolo, I. Fajjari, A. Legrand, and P. Mertikopoulos, “Load-aware provisioning of IoT services on fog computing platform,” in *ICC ’19: Proceedings of the 2019 IEEE International Conference on Communications*, 2019.
- [73] M. Bravo, D. S. Leslie, and P. Mertikopoulos, “Bandit learning in concave N -person games,” in *NeurIPS ’18: Proceedings of the 32nd International Conference of Neural Information Processing Systems*, 2018.
- [74] Z. Zhou, P. Mertikopoulos, S. Athey, N. Bambos, P. W. Glynn, and Y. Ye, “Learning in games with lossy feedback,” in *NIPS ’18: Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018.
- [75] R. I. Boş, P. Mertikopoulos, M. Staudigl, and P. T. Vuong, “On the convergence of stochastic forward-backward-forward algorithms with variance reduction in pseudo-monotone variational inequalities,” in *NIPS ’18: Workshop on Smooth Games, Optimization and Machine Learning (SGO&ML)*, 2018.
- [76] A. Ward, Z. Zhou, P. Mertikopoulos, and N. Bambos, “Power control with random delays: Robust feedback averaging,” in *CDC ’18: Proceedings of the 57th IEEE Annual Conference on Decision and Control*, 2018.
- [77] Z. Zhou, P. Mertikopoulos, N. Bambos, P. W. Glynn, Y. Ye, J. Li, and F.-F. Li, “Distributed asynchronous opti-

- mization with unbounded delays: How slow can you go?,” in *ICML '18: Proceedings of the 35th International Conference on Machine Learning*, 2018.
- [78] M. Leconte, G. Paschos, P. Mertikopoulos, and U. Kozat, “A resource allocation framework for network slicing,” in *INFOCOM '18: Proceedings of the 37th IEEE International Conference on Computer Communications*, 2018.
- [79] P. Mertikopoulos, C. H. Papadimitriou, and G. Piliouras, “Cycles in adversarial regularized learning,” in *SODA '18: Proceedings of the 29th annual ACM-SIAM Symposium on Discrete Algorithms*, 2018.
- [80] R. Colini-Baldeschi, R. Cominetti, P. Mertikopoulos, and M. Scarsini, “The asymptotic behavior of the price of anarchy,” in *WINE 2017: Proceedings of the 13th Conference on Web and Internet Economics*, 2017.
- [81] J. Cohen, A. Héliou, and P. Mertikopoulos, “Learning with bandit feedback in potential games,” in *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- [82] Z. Zhou, P. Mertikopoulos, N. Bambos, S. P. Boyd, and P. W. Glynn, “Stochastic mirror descent for variationally coherent optimization problems,” in *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- [83] Z. Zhou, P. Mertikopoulos, N. Bambos, P. W. Glynn, and C. Tomlin, “Countering feedback delays in multi-agent learning,” in *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.
- [84] J. Cohen, A. Héliou, and P. Mertikopoulos, “Hedging under uncertainty: Regret minimization meets exponentially fast convergence,” in *SAGT '17: Proceedings of the 10th International Symposium on Algorithmic Game Theory*, 2017.
- [85] A. L. Moustakas, P. Mertikopoulos, Z. Zhou, and N. Bambos, “Least action routing: Identifying the optimal path in a wireless relay network,” in *PIMRC'17: 28th annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, 2017.
- [86] P. Mertikopoulos and M. Staudigl, “Convergence to Nash equilibrium in continuous games with noisy first-order feedback,” in *CDC '17: Proceedings of the 56th IEEE Annual Conference on Decision and Control*, 2017.
- [87] Z. Zhou, P. Mertikopoulos, A. L. Moustakas, N. Bambos, and P. W. Glynn, “Mirror descent learning in continuous games,” in *CDC '17: Proceedings of the 56th IEEE Annual Conference on Decision and Control*, 2017.
- [88] Z. Zhou, P. Mertikopoulos, A. L. Moustakas, S. Mehdian, N. Bambos, and P. W. Glynn, “Power control in wireless networks via dual averaging,” in *GLOBECOM '17: Proceedings of the 2017 IEEE Global Telecommunications Conference*, 2017.
- [89] P. Mertikopoulos, E. V. Belmega, and L. Sanguinetti, “Distributed learning for resource allocation under uncertainty,” in *GlobalSIP '16: Proceedings of the 2016 IEEE Global Conference on Signal and Information Processing*, 2016.
- [90] A. Marcstel, E. V. Belmega, P. Mertikopoulos, and I. Fijalkow, “Online power allocation for opportunistic radio access in dynamic OFDM networks,” in *VTC '16-Fall: Proceedings of the 2016 IEEE Vehicular Technology Conference*, 2016.
- [91] A. Marcstel, E. V. Belmega, P. Mertikopoulos, and I. Fijalkow, “Online interference mitigation via learning in dynamic IoT environments,” in *GLOBECOM '16: Proceedings of the 2016 IEEE Global Telecommunications Conference*, 2016.
- [92] A. Marcstel, E. V. Belmega, P. Mertikopoulos, and I. Fijalkow, “Interference mitigation via pricing in time-varying cognitive radio systems,” in *NetGCoop '16: Proceedings of the 2016 International Conference on Network Games, Control and Optimization*, 2016.
- [93] A. S. Shafiq, P. Mertikopoulos, and S. Glisic, “A novel dynamic network architecture model based on stochastic geometry and game theory,” in *ICC '16: Proceedings of the 2016 IEEE International Conference on Communications*, 2016.
- [94] S. D'Oro, P. Mertikopoulos, A. L. Moustakas, and S. Palazzo, “Cost-efficient power allocation in OFDMA cognitive radio networks,” in *EUCNC'15: Proceedings of the 2015 European Conference on Networks and Communications*, 2015.
- [95] I. Stiakogiannakis, P. Mertikopoulos, and C. Touati, “No more tears: A no-regret approach to power control in dynamically varying MIMO networks,” in *WiOpt '15: Proceedings of the 13th International Symposium and*

- Workshops on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, 2015.
- [96] E. V. Belmega and P. Mertikopoulos, "Energy-efficient power allocation in dynamic multi-carrier systems," in *VTC '15-Spring: Proceedings of the 2015 IEEE Vehicular Technology Conference*, (Glasgow, Scotland), May 2015.
 - [97] I. Stiakogiannakis, P. Mertikopoulos, and C. Touati, "No regrets: Distributed power control under time-varying channels and QoS requirements," in *Allerton '14: Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing*, 2014.
 - [98] S. D'Oro, P. Mertikopoulos, A. L. Moustakas, and S. Palazzo, "Adaptive transmit policies for cost-efficient power allocation in multi-carrier systems," in *WiOpt '14: Proceedings of the 12th International Symposium and Workshops on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, 2014.
 - [99] P. Coucheney, B. Gaujal, and P. Mertikopoulos, "Distributed optimization in multi-user MIMO systems with imperfect and delayed information," in *ISIT '14: Proceedings of the 2014 IEEE International Symposium on Information Theory*, 2014.
 - [100] G. Bacci, E. V. Belmega, P. Mertikopoulos, and L. Sanguinetti, "Energy-aware competitive link adaptation in small-cell networks," in *WiOpt '14: Proceedings of the 12th International Symposium and Workshops on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, 2014.
 - [101] P. Mertikopoulos and A. L. Moustakas, "Entropy-driven optimization dynamics for Gaussian vector multiple access channels," in *ICC '13: Proceedings of the 2013 IEEE International Conference on Communications*, 2013.
 - [102] P. Mertikopoulos and A. L. Moustakas, "Riemannian-geometric optimization methods for MIMO multiple access channels," in *ISIT '13: Proceedings of the 2013 IEEE International Symposium on Information Theory*, 2013.
 - [103] P. Mertikopoulos and E. V. Belmega, "Adaptive spectrum management in MIMO-OFDM cognitive radio: An exponential learning approach," in *ValueTools '13: Proceedings of the 7th International Conference on Performance Evaluation Methodologies and Tools*, 2013.
 - [104] J. Lepping, P. Mertikopoulos, and D. Trystram, "Accelerating population-based search heuristics by adaptive resource allocation," in *GECCO '13: Proceedings of the 15th ACM Annual Conference on Genetic and Evolutionary Computation*, pp. 1165–1172, 2013.
 - [105] P. Mertikopoulos, "Strange bedfellows: Riemann, Gibbs and vector Gaussian multiple access channels," in *NetGCoop '12: Proceedings of the 2012 International Conference on Network Games, Control and Optimization*, 2012.
 - [106] P. Mertikopoulos, E. V. Belmega, and A. L. Moustakas, "Matrix exponential learning: Distributed optimization in MIMO systems," in *ISIT '12: Proceedings of the 2012 IEEE International Symposium on Information Theory*, pp. 3028–3032, 2012.
 - [107] P. Mertikopoulos and A. L. Moustakas, "Selfish routing revisited: Degeneracy, evolution and stochastic fluctuations," in *ValueTools '11: Proceedings of the 5th International Conference on Performance Evaluation Methodologies and Tools*, 2011.
 - [108] P. Mertikopoulos, E. V. Belmega, A. L. Moustakas, and S. Lasaulce, "Dynamic power allocation games in parallel multiple access channels," in *ValueTools '11: Proceedings of the 5th International Conference on Performance Evaluation Methodologies and Tools*, 2011.
 - [109] P. Mertikopoulos and A. L. Moustakas, "Learning in the presence of noise," in *GameNets '09: Proceedings of the 1st International Conference on Game Theory for Networks*, 2009.
 - [110] P. Kazakopoulos, P. Mertikopoulos, A. L. Moustakas, and G. Caire, "Distribution of MIMO mutual information: a large deviations approach," in *ITW '09: Proceedings of the 2009 IEEE Information Theory Workshop*, 2009.
 - [111] P. Mertikopoulos, A. L. Moustakas, and N. Dimitriou, "Vertical handover between wireless service providers," in *WiOpt '08: Proceedings of the 6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, 2008.
 - [112] N. Dimitriou, P. Mertikopoulos, and A. L. Moustakas, "Vertical handover between wireless standards," in *ICC '08: Proceedings of the 2008 IEEE International Conference on Communications*, 2008.
 - [113] P. Mertikopoulos and A. L. Moustakas, "The simplex game: Can selfish users learn to operate efficiently

in wireless networks?” in *ValueTools '07: Proceedings of the 2nd International Conference on Performance Evaluation Methodologies and Tools*, 2007.