Ioannis Mitliagkas
Curriculum Vitae

Contact Information

Department of Computer Science and Operations Research
University of Montréal
Pav. André-Aisenstadt
CP6128, Succ. Centre-Ville
Montreal (QC) H3C 3J7

E-mail: ioannis@iro.umontreal.ca
Web: mitliagkas.github.io

Academic appointments

University of Montréal — Started September 2017

Assistant Professor, Department of Computer Science and Operations Research

Stanford University — 2015-2017

Postdoctoral Research Fellow, Departments of Statistics and Computer Science

Supervised by: Assistant Prof. Christopher Ré, Adjunct Prof. Lester Mackey

Research Interests

Statistical machine learning, optimization, high-dimensional statistics, MCMC methods, large-scale and distributed learning systems.

Education

The University of Texas at Austin

PhD, ECE department. — Awarded in August 2015

Advised by: Prof. Constantine Caramanis and Prof. Sriram Vishwanath

Thesis topic: Resource-Constrained, Scalable Learning

Technical University of Crete, Chania, Greece

MSc., ECE department. — 2008 - 2010

Successfully defended thesis in the summer of 2010.

Advisor: Professor Nikos D. Sidiropoulos

Area of Study: Optimization Problems in Wireless Telecommunications

Technical University of Crete, Chania, Greece

Diploma, Electronic and Computer Engineering (5 year degree), — 2002 - 2008

Advisor: Professor Nikos D. Sidiropoulos

Thesis topic: Convex Approximation-based Joint Power and Admission Control for Cognitive Underlay Networks

GPA: 9.01/10, second in class.

Scholarships, Awards

Gerondelis Foundation Inc.: Graduate Scholarship, 2014

The University of Texas at Austin: Microelectronics and Computer Development (MCD) Fellowship, 2009-2011

Technical University of Crete: Undergraduate excellence award, 2008

State Scholarships Foundation (Greece): Undergraduate excellence award, 2005

Technical Chamber of Greece: Undergraduate excellence award, 2005

Research and Teaching

The University of Texas at Austin **Spring 2012**
Teaching Assistant—Information Theory

The University of Texas at Austin **2009-2015**
Research Assistant

Technical University of Crete **Fall 2008**
Teaching Assistant—Telecommunication Networks

Technical University of Crete **May 2007 to August 2008**
Undergraduate Research Assistant

Publications

T. Kurth, J. Zhang, N. Satish, I. Mitliagkas, E. Racah, M.A. Patwary, T. Malas, N. Sundaram, W. Bhimji, M. Smorkalov, J. Deslippe, M. Shiryaev, S. Sridharan, P. Dubey Deep Learning at 15PF: Supervised and Semi-Supervised Classification for Scientific Data. Accepted, *Supercomputing (SC), 2017*.

I. Mitliagkas, L. Mackey Improving Gibbs Sampler Scan Quality with DoGS. International Conference on Machine Learning (ICML), 2017.

I. Mitliagkas, C. Zhang, S. Hadjis, C. Ré Asynchrony begets Momentum, with an Application to Deep Learning. *Allerton Conference on Communication, Control, and Computing, 2016*.

B. He, C. De Sa, I. Mitliagkas, C. Ré Scan Order in Gibbs Sampling: Models in Which it Matters and Bounds on How Much. *Neural Information Processing Systems (NIPS) 2016*.

J. Zhang, C. De Sa, I. Mitliagkas, C. Ré Parallel SGD: When does averaging help? *Optimization Methods for the Next Generation of Machine Learning Workshop, ICML 2016, New York City*.

I. Mitliagkas, M. Borokhovich, A. Dimakis, C. Caramanis FrogWild! – Fast PageRank Approximations on Graph Engines. *VLDB*, 2015 – Preliminary version appeared at *NIPS 2014* Workshop.

D. Papailiopoulos, I. Mitliagkas, A. Dimakis, C. Caramanis. Finding dense subgraphs through low-rank approximations. *International Conference on Machine Learning, 2014 (Vol. 14, pp. 1890-1898)*.

I. Mitliagkas, C. Caramanis, P. Jain. Memory-limited Streaming PCA. Appeared in *Neural Information Processing Systems (NIPS), 2013*.

I. Mitliagkas, A. Gopalan, C. Caramanis, S. Vishwanath. User Rankings from Comparisons: Learning Permutations in High Dimensions. *Allerton Conference on Communication, Control, and Computing, 2011*.

I. Mitliagkas, N. D. Sidiropoulos, and A. Swami. Joint Power and Admission Control for Ad-hoc and Cognitive Underlay Networks: Convex Approximation and Distributed Implementation. *IEEE Transactions on Wireless Communications, 2011*.

I. Mitliagkas, S. Vishwanath. Strong Information-Theoretic Limits for Source/Model Recovery. Appeared in *Allerton Conference on Communication, Control, and Computing, 2010*.

I. Mitliagkas, N. D. Sidiropoulos, and A. Swami. Distributed Joint Power and Admission Control for Ad-hoc and Cognitive Underlay Networks. *ICASSP 2010*.

I. Mitliagkas, N. D. Sidiropoulos, and A. Swami. Convex Approximation-based Joint Power and Admission Control for Cognitive Underlay Networks. *International Wireless Communications and Mobile Computing Conference, 2008. IWCMC'08. IEEE*.

Technical reports and manuscripts in preparation

J. Zhang, I. Mitliagkas, C. Ré. YellowFin and the Art of Momentum Tuning. arXiv:1706.03471.

P. Achlioptas, O. Diamanti, I. Mitliagkas, L. Guibas. Representation Learning and Adversarial Generation of 3D Point Clouds. arXiv:1707.02392.

C. De Sa, B. He, I. Mitliagkas, C. R, P. Xu. Accelerated stochastic power iteration. arXiv:1707.02670.

S. Hadjis, C. Zhang, I. Mitliagkas, C. Ré. Omnivore: An Optimizer for Multi-device Deep Learning on CPUs and GPUs. arXiv:1606.04487.

Recent Invited Talks

Colloquium, University of Montréal	September 2017
Colloquium, The University of Texas, Austin	September 2017
AutoML workshop, ICML, Sydney	August 2017
Workshop on Advances in Computing Architectures, Stanford SystemX	April 2016
ITA workshop, UC San Diego	February 2017
AAAI 2017 Workshop on Distributed Machine Learning	February 2017
Microsoft Research, Cambridge, UK	December 2016
SystemX Stanford Alliance Fall Conference	November 2016
Microsoft Research, New England	October 2016
Allerton Conference	September 2016
Google	August 2016
MIT Lincoln Labs	August 2016
NVIDIA	July 2016

Professional Service

Reviewer for a number of journals and conferences including NIPS, ICML, COLT, AISTATS, AAAI, ICLR, Transactions on Information Theory, ISIT, ICASSP, Transactions on Wireless Communications.

Technical Skills

Languages: C, C++, Java, Python, Julia, Matlab, Scala.

Distributed programming: Worked on Caffe, TensorFlow, MapReduce, Spark, GraphLab, Amazon EC2 infrastructure. Hacked the engine of GraphLab to improve its random algorithms support. Implemented asynchronous training capability on IntelCaffe.

Parallel programming: Lock-free multi-threaded programming in C, multi-process programming in Python.

Other: Some experience in reverse software engineering and network vulnerability detection tools. Hardware design and programming: VHDL, assembly language programming (x86, MIPS, AVR).

GRADUATE COURSE HIGHLIGHTS

Algorithms: Techniques and Theory (CS department), Convex Analysis, Information Theory, Randomized Algorithms (CS department), Systems Theory, Topics in Network Sciences, Analysis and Design of Communication Networks, Theory of Probability (Math Department).

REFERENCES

Christopher Ré, Stanford University chrismre@cs.stanford.edu

Lester Mackey, MSR New England/Stanford University lmackey@stanford.edu

Constantine Caramanis, UT Austin constantine@utexas.edu

Alex Dimakis, UT Austin dimakis@austin.utexas.edu

Sriram Vishwanath, UT Austin sriram@austin.utexas.edu

Nikos D. Sidiropoulos, University of Minnesota nikos@ece.umn.edu