

Pablo Samuel Castro

41 rue Violet
Paris, France 75015
+33.6.79.44.66.06

pablosamuelcastro@gmail.com
<http://www.cs.mcgill.ca/~pcastr>

Research Interests Artificial Intelligence; Reinforcement Learning; Markov Decision Processes; Planning; Machine Learning; Data mining; Stochastic Decision Problems; Formal Verification; Model Checking; Ubiquitous computing.

Current Position Postdoctoral researcher at Institut Telecom SudParis, under the supervision of Prof. Daqing Zhang.

Education **McGill University**
Ph.D. in Computer Science
Dissertation: *On planning, prediction and knowledge transfer in Fully and Partially Observable Markov Decision Processes*
Advisors: Dr. Doina Precup and Dr. Prakash Panangaden October 2011

McGill University
M.Sc. in Computer Science
Thesis: *Bayesian Exploration in Markov Decision Processes*
Advisors: Dr. Doina Precup and Dr. Prakash Panangaden Fall 2007

McGill University
B.Sc. with Honors in Computer Science, minor in Mathematics
Honors project title: Computer music improvisation using Reinforcement Learning
Advisor: Dr. Doina Precup Spring 2001

Honors **NSERC D3 Postgraduate Scholarship** Fall 2007-Fall 2010
Tenured at McGill University.

FQRNT B2 3 years doctoral research scholarship 2007
Declined.

School of Computer Science TopUp Scholarship Fall 2007–Winter 2010

McGill Graduate Studies Fellowship Fall 2007

Recruitment Excellence Fellowship Winter 2008

Research experience **Post-doctoral Research:** I am in charge of a project that aimed to extract the underlying behaviours and dynamics of a city's road network, by using the GPS logs of around 5000 taxis over a year in a large city in China. I developed a novel algorithm based on inverse Reinforcement Learning for learning behaviours from multiple *ranked* experts; this algorithm was used for extracting good passenger-finding strategies. I developed a mechanism for predicting future traffic conditions, as well as automatically detecting the capacity of the different roads in a city. Finally, I co-supervised two Ph.D. students working on characterizing and automatically detecting anomalous taxi routes.

Doctoral Research: I performed an analysis of various equivalence notions for Partially Observable Markov Decision Processes with large observation spaces. By considering a subset of these observations for updating the agent's uncertainty over states

as well as another for measuring the agent’s performance, I proved some surprising results that have direct consequences for representations of states based on future predictions. I developed algorithms for performing policy transfer in Markov Decision Processes using bisimulation metrics, with theoretical guarantees as well as good empirical performance. Additionally, I developed a method for automatically constructing temporally extended actions in an MDP by means of bisimulation metrics. I developed methods for performing smarter model sampling in model-based Bayesian Reinforcement Learning, with empirical performance superior to other state-of-the-art methods.

Masters Research: I proposed various algorithms for performing Bayesian exploration in Markov Decision Processes. These algorithms approximated the countably-infinite set of all possible futures to construct their exploratory policies, with guaranteed asymptotic convergence to the optimal exploratory policy.

Undergraduate RA: I developed a two-layer recurrent Neural Network that I was training to improvise in jazz. By providing it common chord progressions and grading its performance, the system was able to improve its musicality and produce reasonable improvisations over full songs.

**Supervisory
experience**

I have been co-supervising two Ph.D. students, Lin Sun and Chao Chen, since April 2011. I have participated actively in the technical aspects of their research and in the writing of the research papers.

**Refereed
Publications**

P.S. Castro, D. Zhang, S. Li (2011). “Urban traffic modelling and prediction using large scale taxi GPS traces”. To appear in *Proceedings of the 10th International Conference on Pervasive Computing*.

C. Chen, D. Zhang, **P.S. Castro**, N. Li, L. Sun, S. Li (2011). “Real-time Detection of Anomalous Taxi Trajectories from GPS Traces”. In *Proceedings of the 8th International ICST Conference on Mobile and Ubiquitous Systems*. 31.7% acceptance rate. **Runner up for best paper award.**

P.S. Castro, D. Precup (2011). “Automatic construction of temporally extended actions for MDPs using bisimulation metrics”. In *Proceedings of the European Workshop on Reinforcement Learning (EWRL-2011)*.

P.S. Castro, D. Precup (2010). “Smarter Sampling for Bayesian Reinforcement Learning”. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD 2010)*. Pages 200–214. 18.23% acceptance rate.

P.S. Castro, D. Precup (2010). “Using bisimulation for policy transfer in MDPs”. In *Proceedings of the 24th AAAI Conference (AAAI-10)*. Pages 1065–1070. 26.9% acceptance rate.

P.S. Castro, D. Precup (2010). “Using bisimulation for policy transfer in MDPs”. Extended abstract in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*. Pages 1399–1400. 43% acceptance rate.

P.S. Castro, D. Precup (2010). “Using bisimulation for policy transfer in MDPs”. In *10th Adaptive and Learning Agents Workshop (ALA-10)*.

P.S. Castro, P. Panangaden, D. Precup (2009). “Equivalence relations in Fully and Partially Observable Markov Decision Processes”. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-09)*. Pages 1653 – 1658. 25.7% acceptance rate.

P.S. Castro, D. Precup (2007). “Using linear programming for Bayesian exploration in Markov Decision Processes”. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*. Pages 2437–2442. 34.7% acceptance rate.

N. Ferns, **P.S. Castro**, D. Precup, P. Panangaden (2006). “Methods for Computing State Similarity in MDPs”. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence (UAI 2006)*. Pages 174–181. 32% acceptance rate.

**Publications
under review**

P.S. Castro, D. Zhang, C. Chen, S. Li (2012). “From Taxi GPS Traces to Social and Community Dynamics: A Survey”. Submitted to *ACM Computing Survey*.

Lin Sun, **P.S. Castro**, D. Zhang, C. Chen, S. Li (2011). “Detecting and Analyzing Anomalous Trips in a GPS-Based Intelligent Taxi System”. Submitted to *the 10th International Conference on Mobile Systems, Applications, and Services*.

Talks

Real-time Detection of Anomalous Taxi Trajectories from GPS Traces
8th International ICST Conference on Mobile and Ubiquitous Systems (Runner up for best paper award)
December, 2011

Inverse Reinforcement Learning with Multiple Ranked Experts
McGill University (Host: Doina Precup)
September, 2011

Automatic construction of temporally extended actions for MDPs using bisimulation metrics
European Workshop on Reinforcement Learning
September, 2011

On Prediction and Planning in Partially Observable Markov Decision Processes with Large Observation Sets
Robotics Institute, Carnegie Mellon University (Host: Stéphane Ross)
February, 2011

Smarter Sampling in Model-Based Bayesian Reinforcement Learning
ECML-PKDD 2010, Barcelona, Spain
September 2010

On Prediction and Planning in Partially Observable Markov Decision Processes with Large Observation Sets
AI group, Universitat Pompeu Fabra, Barcelona, Spain (Host: Hector Geffner)
September 2010

Using bisimulation for policy transfer in MDPs
Twenty-Fourth AAAI Conference, Atlanta, GA
July 2010

Using bisimulation for policy transfer in MDPs
Workshop on Adaptive and Learning Agents (ALA), in AAMAS-2010
May 2010

Using bisimulation for knowledge transfer in MDPs
Oxford University, UK (Host: Michael Benedikt)
February 2010

Equivalence relations in Fully and Partially Observable Markov Decision Processes
Oxford University, UK (Host: Marta Kwiatkowska)
February 2010

Using bisimulation for knowledge transfer in MDPs
AI group, UC Berkeley, CA (Host: Alexandre Bouchard-Côté)
July 2009

Equivalence relations in Fully and Partially Observable Markov Decision Processes
Twenty-First International Joint Conference on Artificial Intelligence, Pasadena, CA
July 2009

Using bisimulation for knowledge transfer in MDPs
AI group, Universitat Pompeu Fabra, Barcelona, Spain (Host: Anders Jonsson)
June 2009

Using bisimulation for knowledge transfer in MDPs
IDSIA Research Institute, Lugano, Switzerland (Host: Juergen Schmidhuber)
June 2009

Teaching experience

Institut Telecom SudParis

Guest lecturer

Fall 2011

Delivered a lecture overviewing existing research on mining large scale taxi GPS traces to a Masters level course on Ambient Intelligence.

McGill University

Teaching assistant

Fall 1999-Spring 2001, Fall 2005–Present

I have been a teaching assistant for several courses. Responsibilities included grading assignments, holding office hours to answer students' questions, as well as some extra tasks specified below. The numbers in parentheses after each course indicate the number of times I have been a teacher's assistant for that course.

- COMP-202: Introduction to Computing (1)
- COMP-250: Introduction to Computer Science (3). Gave a tutorial on trees as data structures and relevant algorithms.
- COMP-302: Programming Languages & Paradigms (4). Assisted in developing some of the assignments.
- COMP-330: Theoretical Aspects of Computer Science (3). Delivered three class lectures on bisimulation and model checking.
- COMP-431: Data Structures and Algorithms (1)
- COMP-524: Theory of Programming Languages (1)
- COMP-525: Formal Verification (3). Delivered three class lectures on Information Systems.
- COMP-526: Probabilistic Reasoning and AI (1)

Service

Program Committee member for AAAI 2011. Reviewer for Annals of Operations Research (AnOR), Computational Intelligence and the 7th International Colloquium on Theoretical Aspects of Computing (ICTAC 2010).

Volunteered for the 26th International Conference on Machine Learning (ICML-09) verifying the formatting of final accepted papers.

Other Work Experience

Software engineer, CAE, Montreal, QC Summer 2006

- Traveled to Memphis, TN to perform a full system update for a flight simulator from a GOULD system to a Linux box. Tasks involved porting software, interfacing with old hardware and ensuring simulator passed FAA certification.
- Acted as group leader for the entire update

Software Engineer, CAE, Montreal, QC July 2001-August 2005

- In charge of performing full and partial system updates for flight simulators as well as avionic software updates.
- Driver development to interface with old hardware.
- In most projects acted as group leader for the entire update.
- Traveled to Trenton, ON; Charlotte, NC; Dallas, TX; Maastricht, The Netherlands; Hyderabad, India; Hong Kong; Shanghai, China; Toowoomba, Australia.

Web designer and OOP instructor, Quito, Ecuador Summers of 1998,1999,2000

- Designed the corporate webpage for GMS.
- Gave an intensive course on Object Oriented Programming to the technical staff of GMS.
- Developed user interactive forms for CYBERWEB's corporate website.
- Designed the user interface of an internet banking software product for MACOSA.

Other activities

I have always been actively involved in music. To date I have recorded three demo CDs and two full length albums. I have played many festivals and theatres around Quebec, including the Montreal International Jazz Festival. One of my compositions, *Tierra Nuestra* won the third prize in the Latin category of the 17th Annual Billboard World Song Contest 2010.

I am the musical director and co-leader of my current band, Ecos de Portoalegre. I am in charge of writing most of the music, lyrics and arrangements, obtaining shows, and maintaining the finances of the group.

In high school one of my compositions came second place in a city-wide competition. Because of this, my high school awarded me a half-scholarship for my senior year.

Personal

Born in Quito, Ecuador in 1979. Aside from Quito, lived in Miami, FL for 4.5 years, in Puerto Rico for 1.5 years, in Montreal, QC for 11 years, in Gatineau, QC for 3 years, and in Paris, France for 6 months.

Citizenship: Canadian and Ecuadorian.

Languages: Spanish (maternal), English (fluent), French (mid-level), Portuguese (beginner).