

EMILIO PARISOTTO

Website: <http://www.cs.toronto.edu/~eparisotto/>

Nationality: Canadian

Languages: English, French

EDUCATIONAL QUALIFICATIONS

PhD in Machine Learning, University of Toronto

- Supervisor: Ruslan Salakhutdinov
- 2015-Present

Honours Bachelor of Science, Computer Science Specialist, Trinity College, University of Toronto

- 2012-2015

Quebec College Diploma, Pure and Applied Science, Marianopolis College of Montreal

- 2010-2012 (Transferred to University of Toronto)

St. George's High School of Montreal

- 2005-2010

AWARDS

Natural Sciences and Engineering Research Council (NSERC) CGS-M

- University of Toronto (2015-2016)

Drew Thompson Graduation Scholarship

- Trinity College, University of Toronto (2015)

Provost's Graduation Scholarship

- Trinity College, University of Toronto (2015)

James Scott Scholarship

- Trinity College, University of Toronto (2014)

Drew Thompson Scholarship

- Trinity College, University of Toronto (2013)

Dean's List Scholar at the University of Toronto

- University of Toronto (2013-2015)

NSERC Undergraduate Student Research Award (declined)

- University of Toronto (2014)

Rensselaer Medalist (declined)

- St. George's High School (2010)

RESEARCH PROJECTS

- Multitask and Transfer Reinforcement Learning,
E. Parisotto, J. Ba, R. Salakhutdinov. "Actor-Mimic: Deep Multitask and Transfer Reinforcement Learning". Submitted to International Conference on Learning Representations (ICLR) 2016.

<http://arxiv.org/abs/1511.06342>. Presented at Neural Information Processing Systems (NIPS) 2015 Deep Reinforcement Learning workshop (spotlight + poster).

- Used model compression techniques to train a multitask deep policy network that can play several Atari 2600 games at once, at a near-expert level.
- We show that the representations learnt by the multitask deep policy network are capable of generalizing to new tasks, thereby achieving transfer.
- Captions to Images,
E. Mansimov, **E. Parisotto**, J. Ba, R. Salakhutdinov. “Generating Images From Captions with Attention”. Submitted to International Conference on Learning Representations (ICLR) 2016.
<http://arxiv.org/abs/1511.02793>. Presented at Neural Information Processing System (NIPS) 2015 Multimodal Machine Learning workshop (poster) and Reasoning, Attention, Memory (RAM) workshop (oral).
 - Given a descriptive caption of a visual scene, we define a conditional generative model capable of generating an image corresponding to that caption.
 - Our model combines both visual and textual attention mechanisms, and is trained on the challenging Microsoft COCO dataset.
 - After training on the Microsoft COCO dataset, our proposed model can generate the visual structure of a described scene. In addition, it can in some cases even generalize to textual descriptions of scenes that are not physically plausible, such as a stop sign in the sky.

UNDERGRADUATE RESEARCH PROJECTS

- Electroencephalography (EEG) classification,
E. Parisotto, Y.A. Ghassabeh, S. Freydoonnejad, F. Rudzicz. “EEG Dimensionality Reduction in Automatic Identification of Synonymy”. In Proceedings of the 40th International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2015.
Advisor: Professor Frank Rudzicz, Director of SPOClab, Department of Computer Science, University of Toronto
 - Devised a method of automatically classifying whether a person heard two synonymous or non-synonymous words given their EEG signals, with up to 86.84% accuracy.
 - Feature selection, based on p-values, selected a subset from an extremely large pool of features which were then classified using a support vector machine, with the effects of principal components analysis examined on the classification accuracy.
- Magnetoencephalography (MEG) classification,
E. Parisotto, Y.A. Ghassabeh, M.J. MacDonald, A. Cozma, E.W.Pang, F. Rudzicz. “Automatic Identification of Received Language in MEG”. In Proceedings of Interspeech 2015.
Advisor: Professor Frank Rudzicz, Director of SPOClab, Department of Computer Science, University of Toronto
 - Devised a method of predicting whether a person heard English or Romanian on a semantic processing task given their MEG signals.

- Each subject was tested twice and their MEG recorded, once before learning Romanian and once after a two week daily Romanian language-learning program.
- Up to 90% mean classification rate achieved on MEG data before language-learning, with statistically significant decreases in performance afterwards.
- Independent components analysis was performed on the raw MEG signals, after which a p-value-based feature selection procedure reduced an extremely large pool of features into a manageable subset. This feature set was further reduced with principal components analysis and classification was done using a support vector machine.

GRADUATE COURSE PROJECTS

- CS2542 “Topics in Knowledge Representation and Reasoning: Automated Planning and Reasoning about Action”, Summer 2014
Instructor: Professor Sheila McIlraith, Department of Computer Science, University of Toronto
 - Wrote a literature review covering Inverse Reinforcement Learning and several of its main algorithms.

WORK EXPERIENCE

- CSC384 Teaching Assistant (Course Development), “Introduction to Artificial Intelligence”, 2014
Instructors: Professor Sheila McIlraith and Professor Fahiem Bacchus, Department of Computer Science, University of Toronto
 - Redeveloped two programming assignments for the course from prolog to python 3.
 - Designed and created the automated marking code.

TOOLS

- Scripting: bash, perl, python, scheme/lisp, DOS, matlab
- Programming Languages: C, C++, Java, C#, prolog, CUDA
- Query Languages: SQL
- Hardware Description Languages: Verilog
- Operating Systems: Windows, Android, Linux
- Graphics APIs: DirectX, OpenGL

OTHER ACTIVITIES AND INTERESTS

- Drawing, creating video games and designing and making custom electronics