

Shixiang (Shane) Gu

Jesus College, Jesus Lane
Cambridge, Cambridgeshire
CB5 8BL, United Kingdom
sg717@cam.ac.uk
website: <http://sg717.user.srcf.net/>

RESEARCH INTERESTS

Deep Learning, Reinforcement Learning, Robotics, Probabilistic Models and Approximate Inference, Causality, and other Machine Learning topics.

EDUCATION

PhD in Machine Learning **2014 Present**

University of Cambridge, UK.

Advised by Prof. Richard E. Turner and Prof. Zoubin Ghahramani FRS.

Max Planck Institute for Intelligent Systems, Germany.

Advised by Prof. Bernhard Schölkopf.

BASc in Engineering Science, Electrical and Computer major **2009 2013**

University of Toronto, Canada.

CGPA: 3.93. Highest Rank 1st of 264 students (2010 Winter). Recipient of Award of Excellence 2013 (awarded to top 5 graduating students from the program).

High School, Vancouver, Canada. **2006 2009**

Middle School, Shanghai, China. **2002 2006**

Primary School, Tokyo, Japan. **1997 2002**

PROFESSIONAL EXPERIENCE

Google Brain, Google Research, Mountain View, CA.

PhD Intern **June 2015 Jan. 2016, June 2016 Sept. 2016, June 2017 Sept. 2017**

Supervisors & collaborators: Ilya Sutskever, Sergey Levine, Vincent Vanhoucke, Timothy Lillicrap.

- Published 1 NIPS, 2 ICML, 3 ICLR, and 1 ICRA papers as a result of internships.
- Worked part-time during Feb. 2016—May 2016 at Google DeepMind under Timothy Lillicrap.
- Developed the state-of-the-art methods for deep RL in continuous control, and demonstrated the robots learning door opening under 3 hours.

Panasonic Silicon Valley Lab, Panasonic R&D Company of America, Cupertino, CA.

Software Engineer **June 2013 Sept. 2014**

- Pursued R&D efforts on investigating the uses of Deep Convolutional Neural Networks on a range of computer vision tasks, including: benchmarking large-scale image recognition on ImageNet; implementing real-time pedestrian detection using distributed GPU servers; and improving object detection through integration with object proposal techniques
- Initiated R&D efforts on Deep Learning at the Panasonic Silicon Valley Lab, in collaboration with Panasonic Cloud Solutions Centre (CSC) in Japan; worked on convolutional neural nets, denoising

autoencoders, recurrent neural nets, generative models, distance metric learning, margin maximization

Department of Computer Science, University of Toronto, Canada.

Thesis Student

Sept. 2012 – April 2013

Supervisor: Prof. Geoffrey E. Hinton.

- Investigated a novel model for evolving artificial neural networks to enable efficient distributed training of deep neural networks (DNN) on large data sets, against Google's DistBelief model.

EyeTap Personal Imaging Lab, Toronto, Canada.

Research Assistant (Volunteer)

Jan. 2012 – April 2012

Supervisor: Prof. Steve Mann.

- Succeeded in implementing the first quality real-time HDR (High Dynamic Range) video processing on FPGA through extensive Verilog coding and exercised multiple types of image processing techniques.
- Co-authored an IEEE CCECE paper and exhibited demo at ACM SIGGRAPH 2012 Emerging Technologies Exhibition.

The Next 36, Canada's National Program for Entrepreneurial Leadership Initiative.

Entrepreneur, Co-founder and CTO

Dec. 2010 – Aug. 2011

- Selected as one of the top 36 students among all undergraduate students across Canada
- With \$50,000 funding, launched a mobile venture in a team of four during an eight-month period in the area of mobile marketing
- Earned world-class mentorship from the faculty of professors and professionals from Canada and the US, including CEOs, successful entrepreneurs, VCs, and top professors
- Performed the role as the CTO, formulated business strategies, and performed technical execution

University of Toronto Institute for Aerospace Studies

Research Assistant

May 2010 – Aug. 2010

Supervisor: Prof. Craig Steeves

- Worked on a pilot project for a company to manufacture ultra-lightweight aircrafts, modelled and analyzed airfoil stresses based on wing design and pressure distributions using Finite Element software (ABAQUS) and simplified models on MATLAB
- Manufactured samples of carbon fibre composites and tested using a MTS load cell and strain gauges

WORKING PAPERS

1. Vitchyr Pong*, Shixiang Gu*, Murtaza Dalal, Sergey Levine. "Temporal Difference Model Learning: Model-free Deep RL for Model-based Control". Submitted. *equal contribution
2. Benjamin Eysenbach, Shixiang Gu, Julian Ibarz, Sergey Levine. "Leave no Trace: Learning to Reset for Safe and Autonomous Reinforcement Learning". Submitted.

REFEREED PUBLICATIONS

1. Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E. Turner, Bernhard Schölkopf, Sergey Levine. "Interpolated Policy Gradient: Merging On-Policy and Off-Policy Gradient Estimation for Deep Reinforcement Learning". NIPS 2017.

2. Natasha Jaques, Shixiang Gu, Dzmitry Bahdanau, Jose Miguel Hernandez Lobato, Richard E. Turner, Douglas Eck. “Sequence Tutor: Conservative fine-tuning of sequence generation models with KL-control”. ICML 2017.
3. Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E. Turner, Sergey Levine. “Q-Prop: Sample-Efficient Policy Gradient with An Off-Policy Critic”. ICLR 2017 [**Oral, ~3%**].
4. Eric Jang, Shixiang Gu, Ben Poole. “Categorical Reparametrization with Gumble-Softmax”. ICLR 2017.
5. Shixiang Gu*, Ethan Holly*, Timothy Lillicrap, Sergey Levine. “Deep Reinforcement Learning for Robotic Manipulation with Asynchronous Off-Policy Updates”. ICRA 2017. *equal contribution
6. Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, Sergey Levine. “Continuous Deep Q-Learning with Model-based Acceleration”. ICML 2016.
7. Shixiang Gu, Sergey Levine, Ilya Sutskever, Andriy Mnih. “MuProp: Unbiased Backpropagation for Stochastic Neural Networks”. ICLR 2016.
8. Shixiang Gu, Zoubin Ghahramani, Richard E. Turner. “Neural Adaptive Sequential Monte Carlo”. NIPS 2015.
9. Nilesh Tripuraneni*, Shixiang Gu*, Hong Ge, Zoubin Ghahramani. “Particle Gibbs for Infinite Hidden Markov Models”. NIPS 2015. *equal contribution
10. Steve Mann, Raymond Chun Hing Lo, Kalin Ovtcharov, Shixiang Gu, David Dai, Calvin Ngan, Tao Ai. “Realtime HDR (High Dynamic Range) Video for EyeTap Wearable Computers, FPGA-Based Seeing Aids, and GlassEyes”, IEEE CCECE 2012, Montreal, 2012 April 29 to May 2. 6 pages, to be indexed in IEEE Xplore. ACM SIGGRAPH 2012, Emerging Technologies Exhibition.

WORKSHOP PUBLICATIONS

1. Natasha Jaques, Shixiang Gu, Richard E. Turner, Douglas Eck. “Tuning Recurrent Neural Networks with Reinforcement Learning”. NIPS 2016 Deep Reinforcement Learning Workshop.
2. Shixiang Gu, Luca Rigazio. “Toward Deep Neural Network Architectures Robust to Adversarial Examples”. ICLR 2015 Workshop.

INVITED TALKS

1. Shixiang Gu. “Sample-Efficient Deep RL for Robotics”. Vector Institute, Canada, 2017.
2. Shixiang Gu. “Sample-Efficient Deep RL for Robotics”. Hosted by [Xiaoou Tang](#) and [Xiaogang Wang](#). CUHK, Hong Kong, China, 2017.
3. Shixiang Gu. “Sample-Efficient Deep RL for Robotics”. Hosted by [Masashi Sugiyama](#) and [Matsuo Yutaka](#). University of Tokyo, Japan, 2016.

4. Timothy Lillicrap, Shixiang Gu. “Deep RL methods in Robotics”. Reinforcement Learning Forum. Google, USA, 2016.
5. Shixiang Gu. “Generalized Backprop, Neural Particle Filter, and Guided Q-Learning”. Hosted by Pieter Abbeel. UC Berkeley, USA, 2015.
6. Shixiang Gu. “Algorithms for Training Deep Stochastic Neural Networks”. Hosted by Noah Goodman. Stanford University, USA, 2015.
7. Shixiang Gu, Andrey Malinin. “Long Short-Term Memory Networks”. Machine Learning RCC. Cambridge University, UK, 2015.

ACADEMIC ACTIVITIES

1. Program Committee for NIPS Deep Reinforcement Learning Symposium, 2017.
2. Reviewer for International Conference on Robotics and Automation (ICRA), 2018.
3. Reviewer for AISTATS 2018.
4. Reviewer for Neural Information Processing Systems (NIPS), 2017.
5. Reviewer for International Conference on Machine Learning (ICML), 2017.
6. Reviewer for International Conference on Learning Representations (ICLR), 2016, 2017.
7. Reviewer for IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2017.
8. Reviewer for IEEE Conference on Decision and Control (CDC), 2017.

AWARDS/HONORS

2016 — Google Focused Research Award

2014 — Cambridge-Tubingen PhD Fellowship in Machine Learning (partially funded by Facebook)

2014 — MAS Admission to MIT Media Lab (courteously declined)

2013 — Award of Excellence 1T3 (awarded to 5 graduating students for academic achievements)

2013 — Alexander Graham Bell Canada Graduate Scholarship-Master (\$17,500, courteously declined)

2009/2010/2012 — University of Toronto Scholarship (\$8,000) 2009 - Sir Isaac Newton Exam (physics competition), 2nd in Canada

2009 — BC's Brightest Minds (physics competition, \$1500), 1st in British Columbia

2009 — CAP High School Physics Prize Exam, 14th in Canada

2009 — Euclid High School Math Exam, 18th in Canada

SKILLS

- Software: Python, C++, Lua/Torch, Julia, TensorFlow, Torch, Theano
- Languages: Fluent in Japanese, Mandarin, English