

# Silviu Pitis

🇨🇦 Canadian citizen

✉ silviu.pitis@gmail.com

📺 spitis

🐦 @silviupitis

📖 Google Scholar

🏠 silviupitis.com

## Research Summary

---

My ultimate research interest lies in the normative design of general purpose artificial agency: *how should* we design AIs that solve general tasks and contribute positively to society? To date, my work has focused on two more specific questions. First, how should we think about and design goals and abstractions for decision making agents? Second, how can we design such agents to reason about and generalize across a wide range of different goals? My work focuses on large language models and reinforcement learning agents, and applies tools from decision theory, causal inference, neural networks and representation learning. Full research statement available at <https://silviupitis.com/#research>.

## Education

---

### University of Toronto

PH.D. IN COMPUTER SCIENCE, 4.0/4.0. ADVISOR: JIMMY BA.

Toronto, ON

Sep 2018 - May 2024

### Georgia Institute of Technology

M.S. IN COMPUTER SCIENCE, 4.0/4.0

Atlanta, GA

Sep 2016 - Dec 2017

### Harvard Law School

J.D., *Magna Cum Laude*, 4.2/4.0. JOHN M. OLIN FELLOW IN LAW AND ECONOMICS.

Cambridge, MA

Sep 2011 - May 2014

### Schulich School of Business, York University

B.B.A., *With Distinction*, 8.4/9.0

Toronto, ON

Sep 2006 - May 2010

## Publications

---

*All papers accessible via my Google Scholar profile*

### CONFERENCE

- **Silviu Pitis**. [Consistent Aggregation of Objectives with Diverse Time Preferences Requires Non-Markovian Rewards](#). Neural Information Processing Systems (NeurIPS). New Orleans, LA, 2023.
- Andrei Muresanu, Yongchao Zhou, Ziwen Han, **Silviu Pitis**, Harris Chan, Keiran Paster, Jimmy Ba. [Large Language Models are Human-Level Prompt Engineers](#). International Conference on Learning Representations (ICLR). Kijali, Rwanda, 2023.
- **Silviu Pitis**, Elliot Creager, Ajay Mandelkar, Animesh Garg. [MoCoDA: Model-based Counterfactual Data Augmentation](#). Neural Information Processing Systems (NeurIPS). New Orleans, LA, 2022.
- **Silviu Pitis**, Elliot Creager, Animesh Garg. [Counterfactual Data Augmentation using Locally Factored Dynamics](#). Neural Information Processing Systems (NeurIPS). Virtual, 2020. (Poster).
- **Silviu Pitis\***, Harris Chan\*, Stephen Zhao, Bradly Stadie, Jimmy Ba. [Maximum Entropy Gain Exploration for Long Horizon Multi-goal Reinforcement Learning](#). International Conference on Machine Learning (ICML). Virtual, 2020. (Pre-recorded talk).
- **Silviu Pitis\***, Harris Chan\*, Jimmy Ba. [An Inductive Bias for Distances: Neural Nets that Respect the Triangle Inequality](#). International Conference on Learning Representations (ICLR). Virtual, 2020. (Pre-recorded talk).
- **Silviu Pitis**, Michael Zhang. [Objective Social Choice: Using Auxiliary Information to Improve Voting Outcomes](#). International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). Virtual, 2020. (Pre-recorded talk).
- Kristopher De Asis, Alan Chan, **Silviu Pitis**, Daniel Graves, Richard Sutton. [Fixed-Horizon Temporal Difference Methods for Stable Reinforcement Learning](#). AAAI Conference on Artificial Intelligence (AAAI). New York, NY, 2020. (Poster).
- **Silviu Pitis**. [Rethinking the Discount Factor in Reinforcement Learning: A Decision Theoretic Approach](#). AAAI Conference on Artificial Intelligence (AAAI). Honolulu, HI, 2019. (Oral, Poster).
- **Silviu Pitis**. [Source Traces for Temporal Difference Learning](#). AAAI Conference on Artificial Intelligence (AAAI). New Orleans, LA, 2018. (Oral).
- Silviu Pitis. [Methods for Retrieving Alternative Contract Language Using a Prototype](#). International Conference on Law and Artificial Intelligence (ICAIL). London, UK, 2017. (Oral, **Best Student Paper**).

### WORKSHOP

- Mingjian Jiang, Yangjun Ruan, Sicong Huang, Saifei Liao, **Silviu Pitis**, Roger Grosse, Jimmy Ba. [Calibrating Language Models via Augmented Prompt Ensembles](#). Deployment Challenges for Generative AI, ICML 2023.
- **Silviu Pitis**. [Multi-objective Agency Requires Non-Markovian Rewards](#). The Many Facets of Preference-based Learning, ICML 2023.
- **Silviu Pitis**. [Failure Modes of Learning Reward Models for LLMs and other Sequence Models](#). The Many Facets of Preference-based Learning, ICML 2023.

## WORKSHOP, CONTINUED

- Keiran Paster\*, **Silviu Pitis\***, Sheila A. McIlraith, Jimmy Ba. [Return Augmentation Gives Supervised RL Temporal Compositionality](#). NeurIPS DeepRL Workshop 2022. NeurIPS Workshop on Foundation Models for Decision Making 2022.
- Andrei Muresanu, Yongchao Zhou, Ziwen Han, **Silviu Pitis**, Harris Chan, Keiran Paster, Jimmy Ba. [Large Language Models are Human-Level Prompt Engineers](#). NeurIPS Workshop on ML Safety 2022. NeurIPS Workshop on Foundation Models for Decision Making 2022.
- Eric Xu, Jimmy Ba, **Silviu Pitis\***, Harris Chan\*. [Temporary Goals for Exploration](#). NeurIPS DeepRL Workshop 2022.
- **Silviu Pitis**, Duncan Bailey, Jimmy Ba. [Rational Multi-Objective Agents Must Admit Non-Markov Reward Representations](#). NeurIPS Workshop on ML Safety 2022.
- Silviu Pitis, Elliot Creager, Ajay Mandlekar, Animesh Garg. [MoCoDA: Model-based Counterfactual Data Augmentation](#). Decision Awareness in Reinforcement Learning Workshop at ICML 2022. Baltimore, MD, 2022.
- **Silviu Pitis**, Elliot Creager, Animesh Garg. [Counterfactual Data Augmentation using Locally Factored Dynamics](#). Object-Oriented Learning (OOL) at ICML 2020. Virtual, 2020. (Pre-recorded contributed talk, **Outstanding Paper**).
- **Silviu Pitis\***, Harris Chan\*, Stephen Zhao, Bradly Stadie, Jimmy Ba. [Maximum Entropy Gain Exploration for Long Horizon Multi-goal Reinforcement Learning](#). Adaptive and Learning Agents Workshop at AAMAS 2020. Virtual, 2020. (Pre-recorded long talk, **Best Paper**).
- **Silviu Pitis\***, Harris Chan\*, Jimmy Ba. [ProtoGE: Prototype Goal Encodings for Multi-goal Reinforcement Learning](#). The 4th Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM 2019). Montreal, Canada, 2019. (Poster).
- **Silviu Pitis**. [Challenging the MDP Status Quo: An Axiomatic Approach to Rationality for Reinforcement Learning Agents](#). The 1st Workshop on Goal Specifications for Reinforcement Learning, held jointly at ICML/IJCAI/AAMAS 2018. Stockholm, Sweden, 2018. (Poster).
- **Silviu Pitis**. [Reasoning for Reinforcement Learning](#). Hierarchical RL Workshop, NIPS 2017. Long Beach, CA, 2017. (Poster).

## PREPRINTS / OLDER UNPUBLISHED WORK

- Yangjun Ruan, Honghua Dong, Andrew Wang, Silviu Pitis, Yongchao Zhou, Jimmy Ba, Yann Dubois, Chris J. Maddison, Tatsunori Hashimoto. [Identifying the Risks of LM Agents with an LM-Emulated Sandbox](#). Arxiv 2023.
- Silviu Pitis, Michael R. Zhang, Andrew Wang, Jimmy Ba. [Boosted Prompt Ensembles for Large Language Models](#). Arxiv 2023.
- Silviu Pitis. [An Alternative Arithmetic for Word Vector Analogies](#). 2016.
- Silviu Pitis. [Punitive Damages in International Trade](#). Advised by Mark Wu. 2014.
- Silviu Pitis. [Designing Optimal Takeover Defenses](#). Advised by Holger Spamann. 2013.
- Silviu Pitis. [Examining Expected Utility Theory from Descriptive & Prescriptive Perspectives](#). Advised by Archishman Chakraborty. 2010.

## OPEN SOURCE PROJECTS

- Silviu Pitis, Harris Chan, Stephen Zhao. [mrl: modular RL](#). [github.com/spitis/mrl](https://github.com/spitis/mrl). Github. 2020.
- Silviu Pitis. [Swoosh: A leaner, faster backend for Whoosh](#). [github.com/spitis/PyIndex](https://github.com/spitis/PyIndex). Github. 2016.

## Fellowships & Awards

---

- 2023 - 2024 **SRI Graduate Fellowship**, Schwartz Reisman Institute for Technology and Society
- 2020 - 2023 **Canada Graduate Scholarship - Doctoral (CGS D)**, Natural Sciences and Engineering Research Council of Canada
- 2019 - 2024 **Vector Institute Research Grant**, Vector Institute
- 2019 - 2020 **Ontario Graduate Scholarship (OGS)**, Government of Ontario
- 2018 - 2022 **Faculty of Arts and Science Top (FAST) Doctoral Fellowship**, University of Toronto
- 2018 **Faculty of Arts and Science Entrance Scholarship**, University of Toronto
- 2017 **Don Berman Best Student Paper**, International Conference on Law and Artificial Intelligence
- 2013 - 2014 **John M. Olin Fellowship in Law & Economics**, Olin Center for Law & Economics (1 of 3 in HLS Class of 2014)
- 2011 - 2014 **Dean's Scholar Prize (8x)**, Harvard Law School (awarded to top 1-2 students in large classes)
- 2006 - 2010 **York University Renewable Entrance Scholarship**, York University
- 2006 **Invitee**, Canadian Mathematics Olympiad (1 of 60 in Canada)

## Teaching & Supervision

---

### INSTRUCTOR

- Introduction to Machine Learning (CSC311). Fall 2020. (Joint with Roger Grosse, Chris Maddison, Juhan Bae).

### TEACHING ASSISTANT

- Neural Networks and Deep Learning (CSC413/2516). Spring 2020.
- Machine Learning and Data Mining (CSC411). Spring 2019.
- Introduction to Artificial Intelligence (CSC384). Fall 2018.

## RESEARCH SUPERVISION / MENTORING

- Andrew Wang. Variable Rate Discounting as Implicit Bayesian Reinforcement Learning. 2021-2022.
- Eric Xu. Temporary Goals for Exploration. 2021-2022. (Joint with Harris Chan).
- Ian Jiang. Architectures for Learning Local Causal Graphs. FY2021. (Joint with Elliot Creager, Animesh Garg).
- Andrew Gritsevskiy. Leveraging Heuristic Search in Goal-Conditioned RL. 2020-2021. (Joint with Harris Chan).
- Stephen Zhao. Layer-Wise Contrastive Unsupervised Representation Learning. Summer-Fall 2019. (Joint with Jimmy Ba).
- Kiarash Jamali. Metric Nearness Experiments for Deep Norms. 2019-2020. (Joint with Harris Chan).
- Bohan Zhang. Towards Unifying Deterministic and Stochastic Deep Reinforcement Learning Algorithms. Undergraduate Thesis (ESC499). Spring 2019. (Joint with Jimmy Ba).

## Service

---

### REVIEWER / PROGRAM COMMITTEE

- NeurIPS 2022
- ICML 2022
- AAAI 2022
- NeurIPS 2021
- ICML 2021
- ICLR 2021
- AAAI 2021
- NeurIPS 2020
- ICML 2020
- IJCAI 2020
- AAAI 2020
- IJCAI 2019
- NeurIPS 2019
- ICML 2019

## Industry/Professional Experience

---

### Microsoft Research

RESEARCH INTERN (LARGE LANGUAGE MODELS; WITH ALESSANDRO SORDONI, ZIANG XIAO, NICOLAS LE ROUX)

Montreal, QC  
May 2023 - May 2024

### Kirkland & Ellis LLP

ASSOCIATE (CORPORATE LAW; PRIVATE EQUITY, PUBLIC M&A AND CROSS-BORDER TRANSACTIONS)  
SUMMER ASSOCIATE

New York, NY  
Oct 2014 - Mar 2016  
May 2013 - Aug 2013

### Harvard University

RESEARCH ASSISTANT (LAW & ECONOMICS; WITH PROFESSOR LOUIS KAPLOW)

Cambridge, MA  
Feb 2012 - May 2013

### Pokerstars

PROFESSIONAL ONLINE POKER PLAYER

Toronto, ON  
Jan 2008 - Aug 2011

## Other

---

### AFFILIATIONS

- Vector Institute for Artificial Intelligence, 2018 to present.
- AAAI, 2017 to 2021.
- New York Bar, 2015 to present.
- Recurse Center, Spring 2016.
- John M. Olin Center for Law and Economics, 2013 - 2014.

### PROGRAMMING

- Poker hand evaluator and Nash equilibrium solver (Java).
- Real-time multi-user whiteboard application (Javascript).
- Contract aggregator and multi-way diff tool (C#, F#).
- Proprietary stock valuation model (Python, Tensorflow).
- Search engine (Python, 2x faster than best Python alternative).
- Deep reinforcement learning framework (Python, Pytorch).

### PRIVATE TUTORING / COACHING

- Poker. 2008 - 2011.
- Management statistics (undergraduate level). 2008.
- Math and physics (gifted high school level). 2006 - 2008.

### ENTREPRENEURIAL & RELATED

- Unchained Partners, Technology Consultant, 2022.
- Smart Contract Developer/Auditor (independent), 2021 - 2022.
- NextAI, Scientist-in-Residence, 2019 - 2020.
- DealWIP (legal tech startup), Advisor, 2019.
- iAgree LLC (alt. dispute resolution startup), Consultant, 2012.
- SCOPE. Cofounder & Director, 2006 - 2009. Not-for-profit that recruited over 60 student volunteers and ran goal-setting workshops for over 2000 Toronto high school students.