# Alberto Camacho

acamacho@cs.toronto.edu • http://www.AlbertoCamacho.com

## RESEARCH EXPERIENCE

Feb 2022 – present **X, The Moonshot Factory (formerly, Google X)**, Mountain View, California, USA.
- Senior Research Scientist
  - Technical Lead of a project that involves simulation and offline Reinforcement Learning (RL).
  - Technical Lead of a project that combines Large Language Models (LLMs) and sequential decision making.
  - Technical Project Manager of a project that combined AI Planning and LLMs for program synthesis.

Oct 2019 – Sep 2021 **Google**, New York.
- AI Resident in Google Brain
  - Web Navigation agents using Deep Reinforcement Learning.
  - Robotic Manipulation agents from pixels using Neuro-Symbolic Deep Reinforcement Learning.
  - Played an instrumental role in learning long-horizon tasks. Such project had 2+ years of prior research.

Jun – Sep 2016 **Microsoft Research**, Redmond, Seattle, USA.
- Research Intern
  - Designed a conversational agent model with more natural conversation topic flows than the baseline.
  - Mentor: Dr. Mona Habib, Principal researcher at Microsoft Research

Sep 2013 – Sep 2014 **Universitat Pompeu Fabra**, Artificial Intelligence Group.
- Research Assistant
  - Project: Development of algorithms for planning under uncertainty
  - Supervisor: Professor Hector Geffner
  - Research areas: artificial intelligence, automated planning, decision-making

## EDUCATION

2014 – 2022 **University of Toronto**, Toronto, Canada.
- Ph.D. in Computer Science, Artificial Intelligence
  - Thesis title: Automata-Theoertic Synthesis of Plans and Reactive Strategies.
  - Supervisor: Prof. Sheila A. McIlraith
  - Research areas: sequential decision-making, automated planning, reactive synthesis

2012 – 2013 **Universitat Pompeu Fabra**, Barcelona, Spain.
- M.Sc. in Intelligent Interactive Systems
  - Thesis Title: Computing Compact Policies for Fully Observable Non-Deterministic Planning Problems
  - Supervisor: Prof. Hector Geffner

**Universitat Politecnica de Catalunya**, Barcelona, Spain.

2005 – 2012 ▪ M.Sc.-level Engineering Degree in Telecommunications

2005 – 2011 ▪ M.Sc.-level Degree in Mathematics

## PROFESSIONAL EXPERIENCE

Jun 2012 – Sep 2012 **European Space Agency**, Science Operations Trainee. Madrid, Spain.
- Project topic: data visualization

Feb 2011 – Mar 2011 **NeoMetrics Analytics**. Analytics Consultant. Barcelona, Spain.
- Project topic: expansion of a supermarket brand
- Development of a tool for data visualization in Google Earth maps

Sep 2009 – Jun 2014 **ESTALMAT Professor**. Barcelona, Spain.
- Taught high school students with exceptional talent in mathematics

## PROFESSIONAL AFFILIATIONS

2018 – 2022 ▪ Vector Institute for Artificial Intelligence, Toronto, Canada.
Graduate student affiliate.

**SELECT HONORS AND AWARDS**

2023 ▪ Best PhD Dissertation Award
This award honors an outstanding PhD dissertation in any area of automated planning and scheduling.
Awarded by he International Conference on Automated Planning and Scheduling (ICAPS).

2023 ▪ X Leadership Program
12-month leadership and personal development program designed for high potential people at X.
Selected by nomination. Only 25 people are selected each year.

2018 ▪ Best System Demonstration at ICAPS-18.
Awarded by popular vote to SynKit: a webservice and API for rapid synthesis of LTL specifications.
Competed with 8 other entries, including systems by IBM and NASA.

2017 – 2018 ▪ Generation Google Scholarship
15 scholarships awarded worldwide.
Award based on leadership and academic merit to underrepresented groups in computer science.

2017 ▪ 50th Anniversary Graduate Scholarship
One of the two awards given by the Department of Computer Science of the University of Toronto.
Based on academic merit and research excellence.

2016 ▪ Microsoft //oneweek Hackathon
2nd Finalist in the Millenials category of the Microsoft //oneweek Hackathon.
International Competition open to all Microsoft employees. 16,174 hackers. 3,834 projects. 47 venues.

▪ 2015-16 DCS Best Teaching Assistant Team
Awarded by the Dept. of Computer Science at the University of Toronto. Introduction to Artificial Intelligence.

▪ Student Life Recognition
Awarded by the Office of Student Life at the University of Toronto.
Based on contributions to the University of Toronto and active participation in campus life.

2015 ▪ Banting & Best Centre for Innovation & Entrepreneurship Collaboration Fellowship
$10,000 competitive funding for a startup project. Awarded by the University of Toronto.

▪ IBM AquaHacking
One of 10 finalists of the hackathon sponsored by IBM.
Prototyped a surveillance system with drones. Press coverage:
- Environmental Monitor. Jan 19, 2018. University of Toronto Doctoral Student Sees Environmental Monitoring Future in Internet of Things.
- Environmental Monitor. Aug 13, 2015. Drone River Project Uses Artificial Intelligence To Advance Automated Monitoring.
- Phys.org. Jun 30, 2015. How emerging technologies can monitor environment, prevent disasters.
- UofT News. Jun 29, 2015. This drone just wants to help protect Canada's rivers, lakes.
- UofT News. Aug 14, 2015. Project Drone River: a crossroad between human efforts and artificial intelligence.

2014 ▪ IBM Sports Hack
Runner-up of the hackathon sponsored by IBM. Toronto, Canada.

2013 – 2014 ▪ Graduate Studies Grant, Universitat Pompeu Fabra
Awarded with one of the 30 annual grants for graduate studies

2011 ▪ Indra Future Minds Competition
Winner of the International Case Study Competition organized by consulting and technology company Indra.
Media coverage:
- Interview article in *"El Economista"*, June 15, 2011.

▪ BEST-WEC Week Engineering Competition
Winner of the Case Study competition, at the Polytechnic University of Catalonia (UPC).

2009 ▪ National Engineering Competition
Winner of the regional phase. Second finalist in the final national phase in Spain.

2007 ▪ IberoAmerican University Mathematical Olympiad
8th Finalist within Spain and Honorable Mention among all Spanish and Latin American countries.

2005 ▪ Spanish Mathematical Olympiad
8th National finalist and Silver Medal.

pre-2005 ▪ Winner of various mathematical competitions and awards
Silver Medal in the Catalan Mathematical Olympiad 2004, 2nd classified in the Kangaroo 2004, winner of the Poincare Prize to the best high-school mathematical thesis.

## PUBLICATIONS

### CONFERENCE PAPERS

[1] **Reward Machines for Vision-Based Robotic Manipulation**. A. Camacho, J. Varley, A. Zeng, D. Jain, A. Iscen, and D. Kalashnikov. In *the 2020 IEEE International Conference on Robotics and Automation* (ICRA), 2021.

[2] **LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning**. A. Camacho, R. Toro, T. Klassen, R. Valenzano, and S. McIlraith. In *28th International Joint Conference on Artificial Intelligence* (IJCAI), pp. 6065–6073, 2019.
- This paper was also presented in the Recent Published Research (RPR) Track at KR 2021.
- This paper was also presented in the KR2ML workshop at NeurIPS 2019.

[3] **Strong Fully Observable Non-Deterministic Planning with LTL and LTL-f Goals**. A. Camacho and S. McIlraith. In *28th International Joint Conference on Artificial Intelligence* (IJCAI), pp. 5523–5531, 2019.

[4] **Learning Interpretable Models Expressed in Linear Temporal Logic**. A. Camacho and S. McIlraith. In *29th Intl. Conf. Automated Planning and Scheduling* (ICAPS), pp. 621–630, 2019.
- Also presented in the KR2ML workshop at NeurIPS 2019 and RACES workshop at KR 2020.

[5] **Towards a Unified View of AI Planning and Reactive Synthesis**. A. Camacho, M. Bienvenu, S. McIlraith. In *29th Intl. Conf. on Automated Planning and Scheduling* (ICAPS), pp. 58–67, 2019.

[6] **Finite LTL Synthesis with Environment Assumptions and Quality Measures**. A. Camacho, M. Bienvenu, and S. McIlraith. In *16th International Conference on Principles of Knowledge Representation and Reasoning* (KR), pp. 454-463, November 2018.

[7] **LTL Realizability via Safety and Reachability Games**. A. Camacho, C. Muise, J. Baier, and S. McIlraith. In *27th International Joint Conference on Artificial Intelligence* (IJCAI), pp. 4683–4691. July 2018.

[8] **SynKit: LTL Synthesis as a Service**. A. Camacho, C. Muise, J. Baier, and S. McIlraith. In *System Demonstrations at 27th International Joint Conference on Artificial Intelligence* (IJCAI), pp. 5817-5819. July 2018.

[9] **Finite LTL Synthesis as Planning**. A. Camacho, J. Baier, C. Muise, and S. McIlraith. In *28th Intl. Conference on Automated Planning and Scheduling* (ICAPS), pp. 29-38. June 2018.

[10] **SynKit: Finite LTL Synthesis as a Service**. A. Camacho, C. Muise, J. Baier, and S. McIlraith. In *System Demonstrations at 28th International Conference on Automated Planning and Scheduling* (ICAPS). June 2018.

[11] **Synthesizing controllers: On the Correspondence Between LTL Synthesis and Non-Deterministic Planning**. A. Camacho, J. Baier, C. Muise, and S. McIlraith. In *Advances in Artificial Intelligence - 31st Canadian Conf. on Artificial Intelligence*, pp. 45-59, May 2018.

[12] **Non-Markovian Rewards Expressed in LTL: Guiding Search Via Reward Shaping**. A. Camacho, O. Chen, S. Sanner, and S. McIlraith. In *10th International Symposium on Combinatorial Search* (SOCS), pp. 159-160, June 2017.

[13] **Decision-Making with Non-Markovian Rewards: From LTL to automata-based reward shaping**. A. Camacho, O. Chen, S. Sanner, and S. McIlraith. In *3rd Multi-disciplinary Conference on Reinforcement Learning and Decision Making* (RLDM), pp. 279-283, June 2017.

[14] **Non-Deterministic Planning with Temporally Extended Goals: LTL over finite and infinite traces**. A. Camacho, E. Triantafillou, C. Muise, J. Baier, and S. McIlraith. In *31st AAAI Conference on Artificial Intelligence* (AAAI), pp. 3716-3724, February 2017.

[15] **From FOND to Robust Probabilistic Planning: Computing compact policies that bypass avoidable deadends**. A. Camacho, C. Muise, and S. McIlraith. In *26th International Conference on Automated Planning and Scheduling (ICAPS)*, pp. 65-69, June 2016.

**WORKSHOP PAPERS**

[16] **Temporal Logic Goal Specifications for Automated Planning**. A. Camacho, S. McIlraith. In *AAAI 2023 Spring Symposium Series*, March 2023.

[17] **SparseDice: Imitation Learning for Temporally Sparse Data via Regularization**. A. Camacho, I. Gur, M. L. Moczulski, O. Nachum, and A. Faust. In *Workshop on Unsupervised Reinforcement Learning* at ICML, July 2021.

[18] **Disentangled Planning and Control in Vision Based Robotics via Reward Machines**. A. Camacho, J. Varley, A. Zeng, D. Jain, A. Iscen, and D. Kalashnikov. In *Workshop on Deep Reinforcement Learning (Deep RL)* at NeurIPS, December 2020.
- This paper was also presented in the "Why Robots Fail to Grasp?" workshop at IROS 2020.

[19] **Towardsa Neural-Guided Program Synthesis for Linear Temporal Logic Specifications**. A. Camacho, and S. McIlraith. In *Workshop on Knowledge Representation and Reasoning Meets Machine Learning (KR2ML)* at NeurIPS, December 2019.

[20] **Non-Markovian Rewards Expressed in LTL: Guiding Search Via Reward Shaping (Extended Version)**. A. Camacho, O. Chen, S. Sanner, and S. McIlraith. In *1st Workshop on Goal Specifications for Reinforcement Learning (GoalsRL)* at IJCAI/ICML/AAMAS, July 2018.

[21] **Bridging the Gap Between LTL Synthesis and Automated Planning**. A. Camacho, J. Baier, C. Muise, and S. McIlraith. In *Workshop on Generalized Planning (GenPlan)* at ICAPS, June 2017.

[22] **Strong-Cyclic Planning when Fairness is Not a Valid Assumption**. A. Camacho and S. McIlraith. In *Workshop on Knowledge-based techniques for problem solving and reasoning (KnowProS)* at IJCAI, July 2016.

[23] **Non-Deterministic Planning with Temporally Extended Goals: Completing the story for finite and infinite LTL (Amended Version)**. A. Camacho, E. Triantafillou, C. Muise, J. Baier, and S. McIlraith. In *Knowledge-based techniques for problem solving and reasoning (KnowProS) Workshop* at IJCAI, July 2016.

[24] **From FOND to Probabilistic Planning: Guiding search for quality policies**. A. Camacho, A. Ganeshen, C. Muise, and S. McIlraith. In *Workshop on Heuristic Search and Domain Independent Planning (HSDIP)* at ICAPS, pp. 20-28, June 2015.

**SUMMER SCHOOLS**

Jul 2018 **Deep Learning and Reinforcement Learning Summer School**, Toronto

Jun 2018 **ICAPS Summer School on Automated Planning**, The Netherlands

Jul 2017 **Reinforcement Learning Summer School**, University of Montreal, Quebec

Jun 2016 **ICAPS Summer School on Automated Planning**, King's College, London

Jul 2013 **European Agent Systems Summer School (EASSS)**, King's College, London

**COMMUNITY SERVICE**

2015 – present

**Senior Program Committee**
- IJCAI 2021

**Program Committee**
- AAAI 2019, AAAI 2020, AAAI 2021, AAAI 2022, AAAI 2023
- IJCAI 2020, IJCAI 2022, IJCAI 2023
- ICAPS 2021, ICAPS 2022, ICAPS 2023

**Reviewer**
- Artificial Intelligence Journal 2023
- ICAPS 2018, IEEE RA-L and ICRA 2021

**Program Committee in Workshops**
- ICAPS 2020 Workshop on Heuristics and Search for Domain-Independent Planning (HSDIP)
- AAAI 2019 Student Abstracts
- ICAPS 2017 Workshop on Generalized Planning (GenPlan)

**Subreviewer**,
- AAAI (2015, 2016, 2017), IJCAI (2015, 2016, 2017, 2018), ICAPS (2016, 2017, 2019), KnowProS@IJCAI (2017)

2022 – present **Friquifund member**, Barcelona.
- Supporting the NGO with he mission of supporting budding geeks in developing their careers.

2015 – 2016 **UofT Ambassador**, University of Toronto
- Instrumental in the creation of a UofT-Spain elite undergraduate exchange program

## TEACHING

2014 – 2018 **University of Toronto**, Teaching Assistant.
- CSC384: Introduction to Artificial Intelligence (Fall 2014, Win 2015-16-17-18)

2013 – 2014 **Universitat Pompeu Fabra**, Teaching Assistant.
- Signals & Systems (Fall 2013, Win 2014)
- Digital Logic & Computers (Win 2014, Sum 2014)

## ADVISORY EXPERIENCE

**Oscar Chen**, Undergraduate Student. University of Toronto.

Sep 2015 – Jun 2016
- Project: Planning in Markov Decision Processes (MDPs)

Jun 2015 – Sep 2015
- Project: Design of a cognitive robot that exploits POMDPs to assist dementia patients.

**Maayan Shvo**, Undergraduate Exchange Student. University of Toronto.

Sep 2015 – Dec 2015
- Project: Improving heuristic search in the state-of-the-art algorithms
- Co-supervised with Postdoctoral Fellow Richard Valenzano

## START-UP EXPERIENCE

Experience as co-founder of technological startup (closed).

2014
- **Business of Software**, Entrepreneurship course at the University of Toronto.

2014
- **Techno Impact Centre**, Startup accelerator at the University of Toronto.

## LANGUAGES

**English:** Proficient, **French:** Basic (forgotten), **Spanish:** Native

## COMPUTER SKILLS

**Programming Languages**
- Preferred: Python.
- Other: C/C++

**Deep Learning Frameworks**
- Preferred: TF-agents, Keras / Tensorflow 2
- Other: ACME, TensorFlow, JAX