

The Fourth Annual Conference on Learning for Dynamics and Control

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

1. Introduction

The conference series on Learning for Dynamics and Control (L4DC) brings together an interdisciplinary community of researchers with a goal to develop learning-based methods that excel at large-scale dynamics modeling and control, while maintaining desirable theoretical properties and guarantees. Over the next decade, embedded devices that sense and control the physical world are likely to become one of the largest new sources of data. Tackling this surge of real-time data emerging from the physical world requires consolidation of several research areas, such as machine learning, control theory, and optimization. Over the past decade, machine learning had an unprecedented impact in several areas, such as computer vision and language modeling. However, real-time data from embedded devices requires the development of machine learning methods that go beyond pattern recognition and address problems in data driven control and decision making, as well as learning-based optimization of dynamical processes. Control theory, firmly rooted in the tradition of model-based design, could provide a strong foundation. However, the scale of data (both temporal and spatial) will require rethinking the foundations of the discipline.

While this research agenda is very important for the future of various areas (control, optimization, reinforcement learning, data-driven decision making), until recently there was no natural forum to bring together researchers from the diverse areas needed to advance the state of the art in Learning for Dynamics and Control. In recent years, automatic control conferences, such as the IEEE Conference on Decision and Control, featured dedicated sessions for researchers interested in learning-based control, safe learning, or constrained learning. These efforts, however, mainly attract researchers with a control background and are not visible to most researchers in machine learning. On the other hand, in core machine learning conferences, such as NeurIPS, ICML, and COLT, there is limited interest in learning for dynamical and decision systems. Many researchers in other related areas (signal processing, optimization, robotics) have growing efforts on the interface with machine learning, emphasizing mostly prediction and less decision making in dynamical systems. Finally, the recently established Conference on Robot Learning (CoRL) is also addressing issues at this interface, but with a focus on the application domain of robotics.

To fill this niche, our goal with the L4DC conference series is to create a forum that brings together pioneers and state of the art research in the areas of control systems, optimization, machine learning, distributed systems, physics modeling, robotics, and related disciplines. An elite conference on this topic can have tremendous impact not only scientifically by bridging distant areas, but also by creating a community that nurtures a growing number of junior researchers working on this emerging interface. By enabling interactions across control, optimization and learning, L4DC aspires to provide a natural home for professional development for students and faculty that may feel

marginalized in the current conference landscape. In the long run, the aim is to support the emergence of a new community that spans the related disciplines, asks novel questions, and develops the foundations of this new scientific area.

Following the success of the inaugural L4DC workshop held in 2019 at MIT, the 2nd Annual Conference on Learning for Dynamics and Control (L4DC2020) was hosted by U.C. Berkeley, the 3rd conference (L4DC2021) was hosted by ETH Zurich. Due to the pandemic, the conference was held in a virtual format in 2020 and 2021. This year, the 4th Annual Conference on Learning for Dynamics and Control (L4DC2022) is hosted at Stanford University.

2. Conference scope

The conference focuses on the foundations and applications of learning for dynamical and control systems. We invited submissions of papers addressing topics including, but not limited, to:

- Foundations of learning of dynamics models
- System identification
- Optimization for machine learning
- Data-driven optimization for dynamical systems
- Distributed learning over distributed systems
- Reinforcement learning for physical systems
- Safe reinforcement learning and safe adaptive control
- Statistical learning for dynamical and control systems
- Bridging model-based and learning-based dynamical and control systems
- Physics-constrained learning
- Physical learning in dynamical and control systems applications in robotics, autonomy, transportation systems, cognitive systems, neuroscience, etc.

While the conference is open to any topic on the interface between machine learning, control, optimization and related areas, its primary goal is to address scientific and application challenges in real-time physical processes modeled by dynamical or control systems.

3. Earlier L4DC events

The inaugural conference took place at MIT during May 30-31 2019, assembling invited speakers across the disciplines of machine learning, dynamics, control systems, optimization and related disciplines. Poster sessions including many other invited researchers were also featured. With a total of 400 registered attendees, the event quickly reached the maximum that the space allowed and registration had to be closed. Details of the event schedule and posters are available at <https://l4dc.mit.edu>. As part of L4DC2019, there was a discussion about the format of the conference for future years. There was great support for turning the event into a regular conference with submitted

papers for review; it was announced that the next conference will be held in Berkeley, CA, USA in 2020.

Due to the global travel restrictions, L4DC2020 had to take place as a virtual event, held during June 11–12, 2020, hosted by U.C. Berkeley. The event featured live presentations by 5 invited speakers and 14 contributed papers, as well as a virtual forum for the posters; details of the program are available at <https://sites.google.com/berkeley.edu/l4dc>. The talks were given on Zoom and streamed live on YouTube, with moderators collecting questions from the YouTube forum to convey to the speakers; attendance of the YouTube stream peaked at around 500 participants. In addition, an OpenReview forum gave the opportunity to the audience to interact with the authors of papers presented as posters.

L4DC2021 was planned as a physical event to be held in Zurich in June 2021. The pandemic once again forced a change of plans, so L4DC2021 was held as a virtual event during June 7–8, 2021, hosted by ETH Zurich. The program featured invited and contributed talks (held live on Zoom and streamed on YouTube and Gather.Town). The conference introduced live poster sessions, held virtually on Gather.Town during both days of the conference. Interest from the community has once again been strong, with at least 600 registrations for the virtual event. L4DC2021 sought submissions with an open call for papers to the relevant communities and received 140 contributions. All were reviewed by members of the program committee and debated through a rebuttal phase with the authors. At the end of the process, 14 of these contributed papers were accepted for oral presentation, and an additional 90 for poster presentation. The program also included a dedicated session at the end of the first day to discuss the format of the conference for future years, with an open invitation to the program committee and the wider audience to put forward ideas. The complete program of L4DC2021 can be found at <https://l4dc.ethz.ch/>.

4. L4DC2022

L4DC2022 sought submissions with an open call for papers to the relevant communities and received 176 contributions. All were reviewed by at least two (an in a number of cases three) members of the program committee. For contributions that received disparate feedback from reviewers, the editors initiated discussions on EasyChair urging reviewers to come to a consensus on whether to accept each contribution. At the end of the process, 16 papers were accepted as oral presentations, and an additional 90 for poster presentation. The conference event was planned for June 23-24, 2022, hosted by Stanford University, Stanford, CA, USA. The complete program of L4DC2022 can be found at <https://l4dc.stanford.edu>.

Taking into account the feedback from 2021, we increased the number of oral presentations and included them as posters as well. The goal was to ensure that there is enough time for the community to learn more about these outstanding works. We also verified that the papers selected for oral presentations span a variety of topics so that these presentations have the potential to engage each part of the community.

Our invited speakers for L4DC2022 include researchers from a broad range of backgrounds and various career stages:

- Stephen Boyd (Stanford)
- Sarah Dean (Cornell University)

- Chuchu Fan (MIT)
- Monroe Kennedy III (Stanford)
- Daniela Rus (MIT)

This year we invited an Awards Chair (Claire Tomlin) to oversee the process for selecting the outstanding papers for awards. We hope that this new aspect of the conference will motivate the researchers to submit their best work to the conference. Furthermore, highlighting these key contributions would provide the broader community with examples of outstanding work covered by the conference.

This year we also planned social and diversity events at the conference. These include campus and lab tours, focus groups for lunch and dinner to discuss research topics of interest in an informal setting, and a sponsored dinner to encourage participation from underrepresented minorities.

5. L4DC 2022 Organizing Committee

- Roya Firoozi, Postdoctoral Researcher in the Aeronautics and Astronautics department at Stanford University
- Negar Mehr, Assistant Professor in the Department of Aerospace Engineering at the University of Illinois at Urbana-Champaign
- Esen Yel, Postdoctoral Scholar in the Aeronautics and Astronautics Department at Stanford University
- Rika Antonova, Postdoctoral Scholar at the Artificial Intelligence Lab at Stanford University
- Jeannette Bohg, Assistant Professor of Computer Science at Stanford University
- Mac Schwager, Associate Professor in Aeronautics and Astronautics at Stanford University
- Mykel Kochenderfer, Associate Professor of Aeronautics and Astronautics at Stanford University
- Claire Tomlin, Charles A. Desoer Chair in Engineering at the University of California Berkeley

6. Acknowledgements

We are grateful to the program committee for all their efforts in reviewing and discussing the papers contributed to L4DC2021, in alphabetic order:

Naman Agarwal, Google AI Princeton
Mahnoosh Alizadeh, University of California Santa Barbara
Aaron Ames, California Institute of Technology
Shankar Anand, Deka University of California Berkeley
Nikolay Atanasov, University of California San Diego
Navid Azizan, MIT
Kamyar Azzizadenesheli, Purdue University
Thomas Beckers, University of Pennsylvania
Nicola Bezzo, University of Virginia

Valentina Breschi, Politecnico di Milano
Mo Chen, Simon Fraser University
Alessandro Chiuso, University of Padova
Sam Coogan, Georgia Institute of Technology
Yuchen Cui, Stanford University
Sarah Dean, Cornell University
Stefano Di Cairano, Mitsubishi Electric Research Laboratories
Florian Doerfler, ETH Zurich
Katie Driggs-Campbell, University of Illinois at Urbana-Champaign
Clemens Eppner, NVIDIA
Chuchu Fan, MIT
Aleksandra Faust, Google Brain
Mahyar Fazlyab, Johns Hopkins University
Jaime Fernández, Fisac Princeton University
Dylan Foster, MIT
David Fridovich-Keil, University of Texas at Austin
Simone Garatti, Politecnico Milano
Konstantinos Gatsis, University of Oxford
Ali Ghadirzadeh, Stanford University
Bahman Ghahsifard, Queen's University
Shromona Ghosh, Waymo
Stephanie Gil, Harvard University
Sofie Haesaert, Eindhoven University of Technology
Ankur Handa, NVIDIA
Hamed Hassani, University of Pennsylvania
Sylvia Herbert, University of California San Diego
Sandra Hirche, Technical University of Munich
Bin Hu, University of Illinois at Urbana-Champaign
Radoslav Ivanov, University of Pennsylvania
Rahul Jain, University of Southern California
Mihailo Jovanovic, University of Southern California
Yiannis Kantaros, University of Pennsylvania
Jens Kober, Delft University of Technology
Alec Koppel, Amazon
Danica Kragic, KTH
Laurent Lessard, Northeastern University
Jiachen Li, Stanford University
Na Li, Harvard University
Yingying Li, University of Illinois at Urbana-Champaign
Changliu Liu, Carnegie Mellon University
Anirudha Majumdar, Princeton University
Horia Mania, MIT
Kostas Margelos, University of Oxford
Sophie Marie, Fosso Politecnico di Torino
Georg Martius, Max Planck Institute

Nikolai Matni, University of Pennsylvania
 Anastasia Mavrommati, Mathworks
 Sayan Mitra, University of Illinois at Urbana-Champaign
 Matthias Müller, Hannover University
 Gergely Neu, Universitat Pompeu Fabra
 Takayuki Osa, Kyushu Institute of Technology
 Lionel Ott, ETH Zurich
 Necmiye Ozay, University of Michigan
 Jacopo Panerati, University of Toronto
 George Pappas, University of Pennsylvania
 Pablo Parillo, MIT
 Francesca Parise, Cornell University
 Panos Patrinos, Katholieke Universiteit Leuven
 Marco Pavone, Stanford University
 Lerrel Pinto, New York University
 Max Raginsky, University of Illinois at Urbana-Champaign
 Lillian Ratliff, University of Washington
 Anders Ratzer, Lund University
 Vicenç Rubies, Royo University of California Berkeley
 Thomas Schön, Uppsala University
 Ransalu Senanayake, Stanford University
 Shahin Shahrampour, Northeastern University
 Yuanyuan Shi, University of California San Diego
 Florian Shkurti, University of Toronto
 Milad Siami, Northeastern University
 Max Simchowitz, MIT
 Koushil Sreenath, University of California Berkeley
 Bartolomeo Stellato, Princeton University
 Zachary Sunberg, University of Colorado Boulder
 Jie Tan, Google
 Sarah Tang, Waymo
 Yuval Tassa, Google/DeepMind
 Sebastian Trimpe, RWTH Aachen University
 Kyriakos Vamvoudakis, Georgia Institute of Technology
 Melanie Zeilinger, ETH Zurich

We are also grateful to the reviewers for all their efforts in reviewing the papers contributed to L4DC2022, in alphabetic order:

Ian Abraham, Mohammad Alsalti, Liviu Aolaritei, Salar Asayesh, Dominik Baumann, Amrit Bedi, Carolin Benjamins, Petar Bevanda, Mohak Bhardwaj, Paul Bonczek, Jan Bruedigam, Mona Buisson-Fenet, Monimoy Bujarbaruah, Bertrand Charpentier, Benjamin Chasnov, Shaoru Chen, Xin Chen, Yize Chen, Ashish Cherukuri, Chih-Yuan Chiu, Jason Choi, Peter Coppens, Ryan Cosner, Noel Csomay-Shanklin, Preston Culbertson, Will Dabney, Bo Dai, Neha Das, Charles Dawson, Alexandre Didier, Victor Dorobantu, Zhe Du, Katharina Ensinger, Evan Faulkner, Malintha Fernando, Lukas Fröhlich, Andreas Rene Geist, Gautam Goel, Johan Grönqvist, Cole Hawkins,

Johannes Hendriks, Yassir Jedra, Junjie Jiao, Aris Kanellopoulos, Sydney Katz, Michael Kepler, Olle Kjellqvist, Pavel Kolev, Shreyas Kousik, Ben Kraske, Sahin Lale, Nicolas Lanzetti, Zakariya Laouar, Emanuel Laude, Simon Le Cleach, Bruce Lee, Thomas Lew, Jingqi Li, Sarah Li, Lars Lindemann, Bo Liu, Victor Lopez, Michael Lu, Eric Luxenberg, Xubo Lyu, Yecheng Ma, Udari Madhushani, Giovanni Luca Marchetti, Jincheng Mei, Kunal Menda, Yue Meng, Zak Mhammedi, Jared Miller, Youngjae Min, Ioanna Mitsioni, Ahmadreza Moradipari, Simon Muntwiler, Keiko Nagami, Siddharth Nair, Adhyyan Narang, Thanh Nguyen-Tang, Cameron Nowzari, Ibrahim Ozaslán, Vincent Pacelli, Alberto Padoan, Arun Pandey, Juan Perdomo, Philipp Pilar, Cristina Pinneri, Jian Qian, Rien Quirynen, Maxim Raginsky, Alfredo Reichlin, Allen Ren, Zhaolin Ren, Venktraman Renganathan, Antonio Horta Ribeiro, Spencer Richards, Rabih Salhab, Rajiv Sambharya, Anna Scampicchio, Mathijs Schuurmans, Ola Shorinwa, Chelsea Sidrane, Sumeet Singh, Friedrich Solowjow, Eleni Stai, Dawei Sun, Haoyuan Sun, Pouria Tajvar, Andrew Taylor, Gustaf Tegnér, Brendan Tracey, Johnathan Tucker, Berkay Turan, Joe Vincent, Abraham P. Vinod, Feicheng Wang, Yufei Wang, Patrick Washington, Michael C. Welle, Liren Yang, Hang Yin, Muhammad Aneeq Zaman, Songyuan Zhang, Yang Zheng, Ingvar Ziemann