

RESEARCH INTERESTS

I work on statistical machine learning and am most excited about solving and gaining insights about challenging reinforcement learning tasks.

EDUCATION

Doctor of Philosophy, Candidate – CARNEGIE MELLON UNIVERSITY AUGUST 2014 – SEPTEMBER 2019

Master of Science – CARNEGIE MELLON UNIVERSITY AUGUST 2014 – DECEMBER 2017

MACHINE LEARNING

GPA: 4.08 (A: 4.00, A+: 4.33)

Thesis Topic: Strategic Exploration in Reinforcement Learning - New Algorithms and Learning Guarantees

Thesis Committee: Emma Brunskill, Rémi Munos, Barnabás Póczos, Benjamin Recht, Benjamin Van Roy

Advisor: Emma Brunskill

Master of Science – TECHNICAL UNIVERSITY OF DARMSTADT NOVEMBER 2011 – APRIL 2014

Major: COMPUTER SCIENCE, Minor: OPTIMIZATION

Final grade: 1.01 (equivalent GPA: 3.99 / 4.0), Rank: 1

Master Thesis: Value-function-based Reinforcement Learning with Temporal Differences

Advisor: Jan Peters

Bachelor of Science – TECHNICAL UNIVERSITY OF DARMSTADT OCTOBER 2008 – NOVEMBER 2011

COMPUTER SCIENCE

Final grade: 1.0 (equivalent GPA: 4.0 / 4.0), Rank: 1

Bachelor Thesis: A Spatial Consistent CRF for Semantic Image Segmentation

Advisors: Stefan Roth, Peter Gehler

RESEARCH POSITIONS AND INTERNSHIPS

Research Scientist at GOOGLE RESEARCH, NEW YORK, US SINCE OCTOBER 2019

Developing and theoretically analyzing algorithms for sequential decision making that are suitable for a range of applications and can tackle real-world challenges.

Intern at GOOGLE CLOUD AI RESEARCH, SUNNYVALE, US MAY 2018 – AUGUST 2018

Developed a framework that improves accountability of current reinforcement learning methods by forcing them to provide certificates for their policies with guaranteed accuracy. Unlike in existing learning frameworks, this allows users to intervene in high-stakes applications if a certificate does not guarantee sufficient quality.

Supervisors: Wei Wei, Lihong Li

Research Intern at MICROSOFT RESEARCH, NEW YORK, US MAY 2017 – AUGUST 2017

Investigated computational feasibility of current reinforcement learning methods for sample-efficient exploration in tasks with very large observation spaces (contextual decision processes) as well as designing improved algorithms for that setting.

Supervisors: Alekh Agarwal, John Langford

Research Intern at MICROSOFT RESEARCH, CAMBRIDGE, UK JUNE 2016 – AUGUST 2016

- Developed method for estimating memory use in reinforcement learning policies
- Empirically validated the method by investigating memory requirement of DQN policies for 49 Atari games
- Derived theoretical foundation for approach by showing that a lower bound on minimum memory requirement is estimated

Supervisors: Sebastian Nowozin, Katja Hofmann

Research Intern at MICROSOFT RESEARCH, CAMBRIDGE, UK

MARCH 2014 – JUNE 2014

Computer Vision Consultant at MICROSOFT RESEARCH, CAMBRIDGE, UK

JULY 2014

- Developed and analyzed algorithms for depth inference in time-of-flight depth cameras
- Build realistic simulator for time-of-flight cameras based on Metropolis-Light-Transport 3D scene renderers
- Implemented inference and optimization algorithms efficiently in Matlab and C++

Supervisor: Sebastian Nowozin

Research Visiting Student at MASSACHUSETTS INSTITUTE OF TECHNOLOGY, US JUNE 2013 – NOVEMBER 2013

LABORATORY FOR INFORMATION AND DECISION SYSTEMS, *Aerospace Controls Laboratory*

- Identified and fixed premature convergence behavior of the iFDD⁺ (incremental feature dependency discovery) algorithm by theoretical analysis
- Developed novel feature expansion algorithm for value function representations in high-dimensional continuous state-spaces outperforming predecessor by 30%
- Became lead maintainer of the RLPy framework for reinforcement learning <http://bitbucket.org/rlpy/rlpy>, establishing rigorous testing, coding standards and implementing many components

Supervisors: Alborz Geramifard, Jonathan P. How

Research Intern at MAX-PLANCK-INSTITUTE FOR COMPUTER SCIENCE, GERMANY APRIL 2011 – OCTOBER 2011

Computer Vision and Multimodal Computing Department, Bernt Schiele

- Developed and evaluated probabilistic graphical model for semantic image segmentation
- Compared and experimentally evaluated directed and undirected graphical model formulations (trained by latent structural support vector machines and likelihood maximization)
- Implemented a distributed expectation maximization algorithm for maximum likelihood estimation for training large datasets on computing clusters.

Supervisor: Peter Gehler

PUBLICATIONS

PEER-REVIEWED CONFERENCES:

1. *Being Optimistic to Be Conservative: Quickly Learning a CVaR Policy*
Ramtin Keramati, **Christoph Dann**, Alex Tamkin, Emma Brunskill
(AAAI '20) Thirty-Fourth AAAI Conference on Artificial Intelligence [pdf]
2. *Policy Certificates: Towards Accountable Reinforcement Learning*
Christoph Dann, Lihong Li, Wei Wei, Emma Brunskill
(ICML '19) International Conference on Machine Learning [pdf]
3. *On Oracle-Efficient PAC Reinforcement Learning with Rich Observations*
Christoph Dann, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, Robert Schapire
(NeurIPS '18) Neural Information Processing Systems [with spotlight talk] [pdf]
4. *Decoupling Gradient-Like Learning Rules from Representations*
Philip S. Thomas, **Christoph Dann**, Emma Brunskill
(ICML '18) International Conference on Machine Learning [pdf]
5. *Unifying PAC and Regret: Uniform PAC Bounds for Episodic Reinforcement Learning*
Christoph Dann, Tor Lattimore, Emma Brunskill
(NeurIPS '17) Neural Information Processing Systems [with spotlight talk] [pdf]
6. *Sample Efficient Policy Search for Optimal Stopping Domains*
Karan Goel, **Christoph Dann**, Emma Brunskill
(IJCAI '17) International Joint Conference on Artificial Intelligence [pdf]
7. *Energetic Natural Gradient Descent*
Philip S. Thomas, Bruno Castro da Silva, **Christoph Dann**, Emma Brunskill
(ICML '16) International Conference on Machine Learning [pdf]

8. *Sample Complexity of Episodic Fixed-Horizon Reinforcement Learning*
Christoph Dann, Emma Brunskill
(NeurIPS '15) Neural Information Processing Systems [\[pdf\]](#)
9. *The Human Kernel*
Andrew Wilson, **Christoph Dann**, Chris Lucas, Eric Xing
(NeurIPS '15) Neural Information Processing Systems [\[pdf\]](#)
10. *Off-Policy Learning Combined with Automatic Feature Expansion for Solving Large MDPs*
Alborz Geramifard, **Christoph Dann**, Jonathan P. How
(RLDM '13) The 1st Multidisciplinary Conference on Reinforcement Learning and Decision Making [\[pdf\]](#)
11. *Pottics – The Potts Topic Model for Semantic Image Segmentation*
Christoph Dann, Peter Gehler, Stefan Roth and Sebastian Nowozin
(DAGM '12) Proceedings of the 34th DAGM Symposium, Springer, August 2012. [\[pdf\]](#)

PEER-REVIEWED JOURNALS:

1. *Automated Matching of Pipeline Corrosion Features from In-line Inspection Data*
Markus Dann, **Christoph Dann**
Reliability Engineering and System Safety, Volume 162, Pages 40-50, June 2017
2. *Bayesian Time-of-Flight for Realtime Shape, Illumination and Albedo*
Amit Adam, **Christoph Dann**, Omer Yair, Shai Mazor, Sebastian Nowozin
(TPAMI '16) IEEE Transactions on Pattern Analysis and Machine Intelligence , 2016 [\[pdf\]](#)
3. *RLPy – A Reinforcement Learning Framework for Research and Education*
Alborz Geramifard*, **Christoph Dann***, Robert H. Klein*, William Dabney, Jonathan P. How
(JMLR OSS Track '15) Journal of Machine Learning Research , 16(Aug):1573-1578, 2015 [\[pdf\]](#)
*: Equal contribution.
4. *Policy Evaluation with Temporal Differences: A Survey and Comparison*
Christoph Dann, Gerhard Neumann, Jan Peters
(JMLR '14) Journal of Machine Learning Research, 15(Mar):809–883, 2014 [\[pdf\]](#)
Also in the Journal Track of the International Conference on Automated Planning and Scheduling, ICAPS, 2015

PREPRINTS, REPORTS AND WORKSHOPS PAPERS:

- *Distributionally-Aware Exploration for CVaR Bandits*
Alex Tamkin, Ramtin Keramati, **Christoph Dann**, Emma Brunskill
Safety and Robustness in Decision Making Workshop, NeurIPS, 2019
- *Memory Lens: How Much Memory Does an Agent Use?*
Christoph Dann, Katja Hofmann, Sebastian Nowozin
European Workshop On Reinforcement Learning (EWRL), 2016 [\[pdf\]](#)
Abstract presented at Workshop on Interpretable Machine Learning, NeurIPS, 2016
- *Thoughts on Massively Scalable Gaussian Processes*
Andrew Wilson, **Christoph Dann**, Hannes Nickisch
Arxiv preprint 2015: 1511.01870 [\[pdf\]](#)

SCHOLARSHIPS AND AWARDS

- 2017 **Best Reviewer Award at NeurIPS 2017**
- 2017 **Runner Up in Citadel Datathon at CMU** with cash prize of \$2500 for team of 4 students
- 2015 **Datenlotsen Award 2015** for outstanding master's thesis in Information Technology at TU Darmstadt
- 2011–2014 **Studienstiftung des deutschen Volkes – German Academic Scholarship Foundation**
Scholarship for Master's studies
given to less than 0.5 % of all students in Germany.
- 2011 – 2012 **Deutschlandstipendium**
Scholarship for Master's studies
given to less than 1 % of all students in Germany.
- 2004 **3. Place in Bavaria, Jugend Forscht in Engineering**
most prestigious competition for research projects of high-school students in Germany
- 2003 **1. Place Landeswettbewerb Mathematik Bayern**
state-wide mathematics competition in Bavaria

CONFERENCE AND INVITED TALKS

- **Policy Certificates: Towards Accountable Reinforcement Learning,**
Spotlight Talk, ICML Conference, July 2019
- **More Efficient and Accountable Reinforcement Learning: New Guarantees and Algorithms,**
UC Berkeley, Host: Peter Bartlett, March 2019
- **On Oracle-Efficient PAC Reinforcement Learning with Rich Observations,**
Spotlight Talk, NeurIPS Conference, December 2018
- **Learning Guarantees for Sample-Efficient and More Accountable Reinforcement Learning,**
UC Berkeley, Host: Benjamin Recht, November 2018
- **Unifying PAC and Regret: Uniform PAC Bounds for Episodic RL,**
Spotlight Talk, NeurIPS Conference, December 2017
- **Unifying PAC and Regret: Uniform PAC Bounds for Episodic RL,**
Machine Learning Lunch, Stanford, October 2017
- **Towards More Practical Reinforcement Learning With Meaningful Guarantees,**
Machine Learning Lunch, CMU, April 2017
- **Value-Function Based Reinforcement Learning,**
Datenlotsen award ceremony, TU Darmstadt, November 2015
- **Policy Evaluation with Temporal Differences: A Survey and Comparison,**
International Conference on Automated Planning and Scheduling (ICAPS), June 2015
- **Policy Evaluation with Temporal Differences,**
University of Cambridge, March 2014

ACADEMIC AND COMMUNITY SERVICES

- Reviewer for Journal of Machine Learning Research (JMLR) 2015,2020, IEEE Transactions on Cybernetics 2015, Journal of Mathematics of Operations Research 2019, IEEE Transactions on Information Theory 2020
- Reviewer for ICML 2015, 2017–2020; NeurIPS 2014–2019; IROS 2014, 2015, 2017; UAI 2016–2019; COLT 2019–2020; AISTATS 2019–2020; AAAI 2019; ICLR 2019
- Student member of the *Doctoral Review Committee* 2016–2019 improving the Machine Learning PhD program at CMU
- Served in Spring 2018 on the *Admission Committee* for the Machine Learning PhD program at CMU
- Served in 2015–2016 on the *Social Committee* (1 of 3) organizing social activities in the ML Department at CMU

TEACHING

- Teaching Assistant for 10-725 Convex Optimization in Fall 2016 (CMU)
- Teaching Assistant for 15-889e Real Life Reinforcement Learning in Fall 2015 (CMU)
- Teaching Assistant for Foundation of Computing in Fall 2010 (TU Darmstadt)

COMPUTER SKILLS

Github profile <http://github.com/chrodan>
Languages Julia, Python, Matlab, \LaTeX , C++
Operating systems Linux (User and Sys-Admin), Windows, OS X

LANGUAGES

German Native speaker
English Fluent (TOEFL: 118 / 120)