



INTERESTS

I am interested in statistical machine learning, in particular in sequential decision making under uncertainty. My current focus is sample-efficient reinforcement learning.

EDUCATION

Doctor of Philosophy, Candidate – CARNEGIE MELLON UNIVERSITY SINCE AUGUST 2014

MACHINE LEARNING

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

Advisor: Emma Brunskill

Research: I work on both gaining theoretical insights into solving reinforcement learning (RL) tasks as well as designing algorithms with good empirical performance. I have characterized how difficult episodic fixed-horizon reinforcement learning is by proving almost tight upper- and lower bounds on the sample complexity.

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

PUBLICATIONS AND WORKING PAPERS

- **Christoph Dann**, Tor Lattimore, Emma Brunskill
Unifying PAC and Regret: Uniform PAC Bounds for Episodic Reinforcement Learning
to appear in Advances in Neural Information Processing Systems (NIPS) 2017
- Philip S. Thomas, **Christoph Dann**, Emma Brunskill
Decoupling Learning Rules from Representations
Arxiv preprint 2017: 1706.03100
- Karan Goel, **Christoph Dann**, Emma Brunskill
Sample Efficient Policy Search for Optimal Stopping Domains
International Joint Conference on Artificial Intelligence (IJCAI) 2017
- Markus Dann, **Christoph Dann**
Automated Matching of Pipeline Corrosion Features from In-line Inspection Data
Reliability Engineering and System Safety, Volume 162, Pages 40-50, June 2017
- **Christoph Dann**, Katja Hofmann, Sebastian Nowozin
Memory Lens: How Much Memory Does an Agent Use?
European Workshop On Reinforcement Learning (EWRL), 2016
Abstract presented at Workshop on Interpretable Machine Learning, NIPS, 2016
- Philip S. Thomas, Bruno Castro da Silva, **Christoph Dann**, Emma Brunskill
Energetic Natural Gradient Descent
International Conference on Machine Learning (ICML) 2016

- Amit Adam, **Christoph Dann**, Omer Yair, Shai Mazor, Sebastian Nowozin
Bayesian Time-of-Flight for Realtime Shape, Illumination and Albedo
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2016
- **Christoph Dann**, Emma Brunskill
Sample Complexity of Episodic Fixed-Horizon Reinforcement Learning
Advances in Neural Information Processing Systems (NIPS) 2015
- Andrew Wilson, **Christoph Dann**, Chris Lucas, Eric Xing
The Human Kernel
Advances in Neural Information Processing Systems (NIPS) 2015
- Andrew Wilson, **Christoph Dann**, Hannes Nickisch
Thoughts on Massively Scalable Gaussian Processes
Arxiv preprint 2015: 1511.01870
- Alborz Geramifard*, **Christoph Dann***, Robert H. Klein* William Dabney, Jonathan P. How
RLPy – A Reinforcement Learning Framework for Research and Education
Journal of Machine Learning Research (JMLR), 16(Aug):1573-1578, 2015
*: Equal contribution of first three authors.
Abstract presented at
 - Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2015
 - Workshop on Machine Learning Open Source Software 2013: Towards Open Workflows, NIPS, 2013
- **Christoph Dann**, Gerhard Neumann, Jan Peters
Policy Evaluation with Temporal Differences: A Survey and Comparison
Journal of Machine Learning Research (JMLR), 15(Mar):809–883, 2014
Also in the Journal Track of the International Conference on Automated Planning and Scheduling, ICAPS, 2015
- Alborz Geramifard, **Christoph Dann**, Jonathan P. How
Off-Policy Learning Combined with Automatic Feature Expansion for Solving Large MDPs
The 1st Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2013.
- **Christoph Dann**, Peter Gehler, Stefan Roth and Sebastian Nowozin
Pottics – The Potts Topic Model for Semantic Image Segmentation
Proceedings of the 34th DAGM Symposium, Springer, August 2012.

RESEARCH VISITS AND INTERNSHIPS

Research Intern at MICROSOFT RESEARCH NEW YORK 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).
Supervisors: Alekh Agarwal, John Langford

Research Intern at MICROSOFT RESEARCH CAMBRIDGE 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 MARCH 2014 – JUNE 2014
Computer Vision Consultant at MICROSOFT RESEARCH CAMBRIDGE 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 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 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

SCHOLARSHIPS AND AWARDS

- 2015 **Datenlotsen Award 2015** for outstanding master's thesis in Information Technology at TU Darmstadt
- 2015 NSF Student Travel Scholarship to attend the International Conference on Automated Planning and Scheduling (ICAPS 2015);
- 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

- **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 ACTIVITIES

- Reviewer for the *Journal of Machine Learning Research (JMLR)*; *ICML 2015, 2017*; *NIPS 2014, 2015, 2016, 2017*; *IROS 2014, 2015, 2017*; *UAI 2016, 2017*, the *Advances in Machine Learning for Sensorimotor Control Workshop 2013* and the *European Workshop on Reinforcement Learning 2012*
- Participant of the Machine Learning Summer School 2013 in Tübingen, Germany (25% acceptance rate)

COMPUTER SKILLS

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

LANGUAGES

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