CURRICULUM VITAE

Name:	**Marcus Hutter**
Occupation:	Professor, ANU
Education: (degrees)	Dr.rer.nat.habil.dipl.inf.univ. [2×BSc,MSc,PhD,Habil]
Main Interests:	Science and Philosophy
HomePage:	http://www.hutter1.net/
Email:	marcus.hutter@anu.edu.au

January 29, 2019

Mini Biography

Marcus Hutter is Professor in the RSCS at the Australian National University in Canberra, Australia. He received his PhD and BSc in physics from the LMU in Munich and a Habilitation, MSc, and BSc in informatics from the TU Munich. Since 2000, his research at IDSIA and now ANU is centered around the information-theoretic foundations of inductive reasoning and reinforcement learning, which has resulted in 150+ publications and several awards. His book "Universal Artificial Intelligence" (Springer, EATCS, 2005) develops the first sound and complete *theory* of AI. He also runs the Human Knowledge Compression Contest (50'000€ H-prize).

This document contains his detailed hyper-linked curriculum vitae.

Contents

Short Biography **2**
Experience **2**
Professional Career **3**
Academic Qualifications **4**
Grants, Prizes, Awards, Honors **5**
Community Service **6**
Outreach / Interviews / Press **9**
Attended Conferences / Workshops / Symposia / ... **10**
(Invited) Short Research Visits **11**
Lecturing **13**
Past & Current Group Members **19**
Publications **22**

Short Biography

Marcus Hutter is Professor in the Research School of Computer Science (RSCS) at the Australian National University in Canberra, Australia. He is also chair of the ongoing Human Knowledge Compression Contest and sponsor of the 50'000€ H-prize. He received a Masters degree in computer science in 1992 from the University of Technology in Munich, Germany, a PhD in theoretical particle physics in 1996, and completed his Habilitation in 2003. He worked as an active software developer for various companies in several areas for many years, before he commenced his academic career in 2000 at the Artificial Intelligence (AI) institute IDSIA in Lugano, Switzerland, where he stayed for six years. His first large project to receive national attention was a complete 3D CAD program for 8 bit computers in Assembler, which he developed during his final year at high school. In his five years in industry, he developed various algorithms for a medical software company, which are still used in equipment sold world-wide. Since 2000, he has mainly worked on fundamental questions in AI resulting in over 150 peer-reviewed research publications. His book "Universal Artificial Intelligence" (Springer, EATCS, 2005) lays down the information-theoretic foundations of inductive reasoning and reinforcement learning, and develops the first sound and complete *theory* of AI. This work has been generously supported by various research grants. He also runs and sponsors the Human Knowledge Compression Contest (50'000€ H-prize). His general interests are in universal artificial intelligence, theories of everything, statistics, philosophy of science, and mathematical and physical puzzles. His current research is centered around reinforcement learning, algorithmic information theory and statistics, universal induction schemes, adaptive control theory, and related areas. He has served (as PC member, chair, organizer) for numerous conferences, and reviews for all major conferences and journals in his research areas. He has given invited lectures at numerous universities, institutions, conferences, workshops, and companies, as well as public presentations and interviews. His work in AI and statistics was nominated for and received several awards (UAI, IJCAI-JAIR, AGI Kurzweil, Lindley).

Experience

Main Interests:	Universal Artificial Intelligence, Physical Theory of Everything, Mathematical and Physical Puzzles, Statistics, Theoretical Computer Science, Numerical Algorithms, Computer Vision&Graphics, Analytic Philosophy.
Artificial Intelligence:	reinforcement/machine learning, algorithmic complexity, optimization, game theory, genetic algorithms, neural nets,Bayesian/robust/expert/MDL/online/sequence prediction.

Engineering:	information theory, adaptive control, time-series forecasting, electronics.
Physics:	non-perturbative quantum field theory, QCD, solitons and instantons, statistical physics, path integrals, anomaly, quark and meson masses, string and brane theory.
Medical:	PencilBeam dose algorithm for radiotherapy and IMRT, Brachytherapy, CT/MT imaging.
Numerics:	Monte Carlo, simulated annealing, multidimensional optimization, finite elements, 1-3d fft & splines & advanced interpolation.
Computer Graphics:	volume & surface rendering.
Image Processing:	segmentation, smoothing, recognition, 2d-3d registration.
Statistics:	probability, Bayes, model selection, sequential decisions.
Mathematics:	discrete math, logic, algebra, analysis.
Programming Languages:	C++, Pascal, Prolog, Fortran, DBase, Forth, Lisp, Basic, Assembler, Html.

Professional Career

2014 - 2016	Associate Director (Research) in RSCS@ANU.
since 2011	Full Professor in the Research School of Computer Science (RSCS) at the Australian National University (ANU).
2011 - 2012	Sabbatical Year in the Machine Learning Laboratory at the ETHZ.
2006 - 2010	Associate Professor in the Research School of Information Sciences and Engineering (RSISE) at the Australian National University (ANU) and senior researcher in the National Information and Communication Technology of Australia (NICTA).
2003/2004	Lecturer at Munich University of Technology, Germany
10.2000 - 2006	Senior researcher and project leader at IDSIA (Research Institute for Artificial Intelligence) in Lugano, Switzerland,

05.1996 - 09.2000	Software developer and project leader at BrainLAB: (Occupation: Numerical algorithms in medical field)
	Development of a Neuro-Navigation system, a Brachytherapy planning system, a dose algorithm (PencilBeam) for radiotherapy for IMRT, a real time software volume renderer, and various image processing modules. Invention of a new image enhancement and post-antialiasing algorithm (patented). Supervision of Diploma theses in computer science, ...
08.1992 - 04.1993	Design & implementation of a protection module+organization for licensing programs in C (IABG)
06.1987 - 10.1987	Implementation of a user interface for an expert system - under GEM (IABG)
02.1986 - 01.1987	Design & implementation of a 3D-CAD-Program in Assembler (Markt & Technik)
03.1983 - 06.1983	Programming of a member organization program in DBase (for tax advisor Keller)
1988 - 1994	Private tuition of high school and university students.

Academic Qualifications

2001 - 2003	Habilitation (≈2nd PhD) in Computer Science at TU-Munich on Optimal Sequential Decisions based on Algorithmic Probability Supervisor: Prof. Wilfried Brauer.
1993 - 1996	PhD (Dr.rer.nat.) in Theoretical Particle Physics on Instantons in QCD at the University (LMU) in Munich. Supervisor: Prof. H. Fritzsch.
1989 - 1992	Masters Degree (Dipl.inform.univ) in Computer Science with Minors in Mathematics at the University of Technology in Munich.
1988 - 1991	Bachelor (Vordiplom) in General Physics at the University of Technology in Munich.
1987 - 1989	Bachelor (Vordiplom) in Computer Science with Minors in Mathematics.

Grants, Prizes, Awards, Honors

2019 - 2023	A$ 7'500'000,- *ANU Grand Challenge*. 10 CIs. Human Machine Intelligence (HMI)
2019 - 2021	US$ 276'000,- *Future of Life Project grant*. Sole CI. The Control Problem for Universal AI: A Formal Investigation (CPUAI)
2018	*First winner of the AI alignment prize round 2* for paper [P18align]: The Alignment Problem for History-Based Bayesian Reinforcement Learners.
2016	*Kurzweil prize for best AGI paper* for paper [P16selfmod]: Self-Modification of Policy and Utility Function in Rational Agents.
2016	*UAI Best student paper* for paper [P16thompgrl]: Thompson Sampling is Asymptotically Optimal in General Environments.
2015 - 2019	A$ 421'500,- *Australian Research Council DP grant*. Sole CI. Unifying Foundations for Intelligent Agents (UFIA)
2014	Honorable Mention *IJCAI-JAIR Best Paper Prize* for paper [P11aixictwx]: A Monte-Carlo AIXI Approximation.
2012 - 2015	A$ 390'000,- *Australian Research Council DP grant*. Primary CI. Feature Reinforcement Learning (FRL)
2009 - 2011	A$ 240'000,- *Australian Research Council DP grant*. Sole CI. From Universal Induction to Intelligent Agents (UAI)
2008 - 2012	A$ 270'000,- *Industrial research grant*. Sole PI. Image-based Car Damage Detection (ICAR)
2009	1st runner up of the *Kurzweil Best AGI Paper Prize* for paper [P09phimdp]: Feature Markov Decision Processes.
2007	*Lindley Prize* awarded for innovative research in Bayesian Statistics. Best paper [P07pcreg, P07pcregx] from 326 submissions to ISBA Valencia 8.
2006 - 2008	SFr 92'730,- *Swiss National Science Foundation grant*. (shared) A Bayesian approach for integrated cancer genome profiling (BIG)
2005 - 2007	SFr 248'772,- *Swiss National Science Foundation grant*. Sole PI Optimal rational AIXI agent based on algorithmic complexity (AIXI)
2003 - 2005	SFr 273'616,- *Swiss National Science Foundation grant*. Sole PI Optimal rational agents in unknown environments (ORAUE)
2001 - 2003	SFr 193'680,- *Swiss National Science Foundation grant*. Unification of universal inductive inference and sequential decision theory (UISD)

Community Service

(Co)organization

- Reinforcement Learning Seminar in Dagstuhl (EWRL 2013)
- Theory&Practice of Machine Learning Workshop in Canberra (TPML 2013)
- European Workshop on Reinforcement Learning (EWRL 2011), Chair.
- Weekly Readings Groups at ANU (AI&RL&KC 2009–2011)
- Algorithmic Learning Theory in Canberra (ALT 2010), General Chair
- Machine Learning Summer School in Canberra (MLSS 2010)
- Artificial General Intelligence in Lugano (AGI 2010), Conference Chair
- Partially Observable Reinforcement Learning Symposium in Vancouver (PORL 2009)
- Machine Learning Summer School in Kioloa (MLSS 2008)
- Kolmogorov Complexity Seminar in Dagstuhl (KC 2006)
- Weekly Theory Reading Group at IDSIA in Lugano (TRG 2004–2005)
- Universal Learning & Optimal Search Workshop at NIPS (ULAOS 2002)

Conference program committee chair

- 2nd Artificial General Intelligence in Washington (AGI 2009)
- 18th Algorithmic Learning Theory in Senday (ALT 2007)

Membership of conference program committees

- IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL 2013 & 2014 & 2016)
- ICALP 2013 Satellite Workshop on Learning Theory and Complexity (LTC@ICALP 2013), Riga, Latvia.
- Conf. on Philosophy and Theory of Artificial Intelligence (PTAI 2013 & 2017)
- The 1st&4&5&6th&7&8&9&10th Conference on Artificial General Intelligence (AGI 2008 & 2011 & 2012 & 2013 & 2014 & 2015 & 2016 & 2017 & 2018 &)
- The 25&26&28&30&32&33rd Conference on Uncertainty in Artificial Intelligence (UAI 2009 & 2010 & 2012 & 2014 & 2016 & 2017)
- 23&30th Annual Conf. on Neural Information Processing Systems (NIPS 2009 & NIPS 2016)
- The 8th Conference on Computability in Europe (CiE 2012), Turing Centenary Conference, Cambridge, UK.
- Ray Solomonoff (1926-2009) 85th Memorial Conference (SMC 2011), Melbourne, Australia.

- 22nd European Conference on Machine Learning (ECML 2011), Athens, Greece.
- The 22&24&27th Intl. Joint Conference on Artificial Intelligence (IJCAI 2011 & 2015 & 2018)
- The 28th Intl. Conference on Machine Learning (ICML 2011), Bellevue, Washington, USA.
- The 21&23&28th Australasian Joint Conference on Artificial Intelligence (AusAI 2008 & 2010 & 2015)
- The 8th biannual International Conference on Artificial Intelligence & Statistics (AISTATS 2007), San Juan, Puerto Rico.
- The 17&19&24th International Conference on Algorithmic Learning Theory (ALT 2006 & 2008 & 2013 & 2015)
- The 18&20th Annual Conference on Learning Theory (COLT 2005 & 2007)
- Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn 2002 & 2004 & 2005 & 2006)

Service to/at the ANU

Most ANU service 2014-2016 below was part of my Associate Director Research duties at RSCS@ANU:

- RSCS-TTMTR: Tenure track mid-term review committee member (3x 2018)
- RSCS-ADirR: Associate Director Research RSCS (2014–2016)
- RSCS-SM: Attend (since 2009) / present at (2015-2016) monthly RSCS staff meetings
- HDR: Chairing the Systems Theme Intl PhD Scholarship committee (2017)
- Attend RSCS retreats (2016&2017) and SML@NICTA retreats (2008&2010)
- Evaluate various RSCS Applicants (2016&2017)
- Acting Head of RSCS (2015-2016 in total for 11 weeks)
- RSCS-MEX: Biweekly meeting of (Assoc.)Dir RSCS (2014-2016)
- RSCS-IAB: Member of RSCS Industry Advisory Board (2015-2016)
- RSCS-SAC: Member of Strategic Appointment Committee (2016)
- RSCS-ER: Preparation of RSCS External Review (2016)
- RSCS-TTC: Member of the (mid-term) tenure (track) committees (2015, 2016)
- RSCS-RT: Support of development of new Research Themes/Areas and Website in RSCS (2015, 2016)
- RSCS-DTCA: Dealing with Defense Trade Control Act (2015-2016)
- NICTA-UP: Attending NICTA University Planning Committee (2016)
- Funding negotiations with DFAT, DSTG, ... (2016)
- ANU-URC: Member of ANU=University Research Committee (2015)
- CECS-EXEC: Monthly meeting of CECS executive (2015)
- CECS-DEAN: Biweekly meeting of (Assoc.)Dirs&Deans CECS/RSCS/RSEng (2015)

- RSCS-SAC: Member of Strategic Appointment Committee (2015-2016)
- RSCS-ERA: Chair of RSCS ERA 2015 FOR08I&CS committee (2014-2015) (classify publications, compose explanatory statement, determine co-author and co-PI statistics, ...)
- ANU-ERA: Member of ANU ERA 2015 Working Group (2015)
- RSCS-QS: Preparing documents for QS ranking (2015)
- CECS-ARC: Support of CECS ARC DP&DECRA Peer Review Workshop (2012&2013&2015&2016)
- CECS-RC: Member/Chair of CECS Research Committee (2014-2016)
- CECS-HDR: Member of CECS HDR scholarship committee (2014-2017) (intl&local PhD student application ranking)
- CECS-TA: Member of dean's teaching award committee (2014&2015)
- ANU-RWG: Member of the ANU Repository Working Group (2013)
- CASS-SSC: Member of the RSSS/CASS@ANU Search/Selection Committee (2011&2013)
- CECS-AB: Member of the CECS@ANU Advisory board (2008)
- CECS-PC: Member of the CECS@ANU Promotions' Committee (2008&2014)
- ANU-CEDAM: Attended one-day CEDAM/ANU course 'Introduction for New Research Supervisors' (2006)

Miscellaneous (mostl. intl.) service

- Invited to speak at many more intl. events (since 2007) most of them declined due to time/distance/financial constraints.
- (Scientific) Advisor for startup company NNAIsense (since 2015)
- Steering Committee member of the Algorithmic Learning Theory (ALT) conference series (2011-2017).
- Founding board member of the AGI Society (since 2011)
- Steering Committee Chair of the European Workshop for Reinforcement Learning (EWRL) Series (since 2011).
- Member of the Editorial Board of the Machine Learning Journal (2011–2014)
- Editor of the Journal of Artificial General Intelligence (since 2009).
- Steering committee member of the AGI conference series (since 2008)
- First Editor of Scholarpedia (since 2007)
- Examiner of a couple of PhD theses (since 2007)
- Chair and sponsor of the 50'000€ Prize for Compressing Human Knowledge (since 2006)
- Expert assessor for the Australian Research Council (ARC) (since 2003).
- Moderator of the Algorithmic Information Theory mailing list (since 2002).
- Reviewer of journals (since 2002) IEEE (TPAMI, TIT, TSP, SMC, TEC), Elsevier (TCS, I&C, JCSS, IPL, IJAR, SIMPAT),
  Springer (AMAI, MLJ, Algorithmica, M&M, ToCS, Erkenntnis, Synthese),

Others (JACM, JMLR, JAIR, JAGI, JCS, Episteme, JoP, JBSB, FI, SS, PIP, Entropy, Algorithms, ...).
- Reviewer of conferences (since 2001) ISIT, UAI, AGI, STACS, NIPS, COLT, ALT, IJCAI, ECML, ICANN, AusAI, CATS, CiE, Benelearn, ACC, ...

Outreach / Interviews / Press

Public outreach

- Society in the wake of Super-Intelligent Machines. *NewScientist (2017), Sydney (Invited public lecture and panel discussion. About 200 attendees)*
- The Future of Artificial Intelligence and Humanity. *National Science Week (2017), Melbourne (Invited public lecture and panel discussion. 150+ attendees)*
- Future of AI. *Symposium: Monash Future Thinkers - Australia in 2050 (2016), Melbourne (Invited public lecture. About 200 attendees)*
- Foundations of Intelligent Agents. *Singularity Summit (2009), New York (Invited public lecture. About 800 participants)*
- On Science, Fiction, and Future Reality. *Guest lecture to high-school students at The American School of Switzerland (TASIS 2005 & 2006), Lugano.*

Interviews

- Five short interviews at IJCAI. *(21 Aug 2017) Adam Ford interviewing MH.*
- Live Radio Midnight in the Desert. *(26 August 2015) Art Bell interviewing MH for 3 hours.*
- Science Magazine: Which movies get AI right? *(17 July 2015) David Shultz interviewing Marcus Hutter and Stuart Russell.*
- ABC24 TV newsline about 'Scientists using AI to advance technology' Very short interview *(31 Jan 2013) Kesha West interviewing Kevin Korb and Frank Farrall and Serge Zhevelyuk and Marcus Hutter.*
- Australian Singularity Summit: Six short interviews. *(17 Aug 2012) Adam Ford interviewing Marcus Hutter.*
- New Scientist Magazine. Universal intelligence: One test to rule them all. *(13 Sep 2011, Issue#2829, p42-45) Celeste Biever interviewing Hernandez-Orallo and Marcus Hutter and others.*
- SoundProof: Woroni's Podcast Experiment. *AI Part I (3 June 2010) and AI Part II (10 June 2010). Jamie Freestone and Mathew McGann interviewing Marcus Hutter and David Chalmers.*
- ABC Radio National. All In the Mind. The coming of 'The Singularity'...or not? *(19 September 2009) 13:00-13:30. Mike McRae interviewing Nick Bostrom, Marcus Hutter, Noel Sharkey, Nigel Dobson-Keeffe, and Richard K. Morgan.*

- L'Eretico - idee arte pensiero. Filosofia e Scienza. Le Frontiere Dell'Intelligenza. *Daniele Lanzillo interviewing Marcus Hutter (July 2009)*

Press coverage

- That 'theory of everything'? A researcher says it has a lot going for it (2010)
  Work [P10ctoex] featured in Medill Reports Chicago (8 Oct 2010).
- 50'000€ Prize for Compressing Human Knowledge (2006)
  Discussion in the Hutter-Prize mailing list, Yahoo group ai-philosophy, news net groups comp.ai.nat-lang, comp.compression, comp.ai, at Slashdot, in the Online Heise news, KurzweilAI.net news, Accelerating Future page, ebiquity news, AGI mailing list, and many others.
- Measuring the Intelligence of a Machine (2005)
  Work [P05iors] featured in Le Monde de l'intelligence (No.1, Nov/Dec.2005) Also commented on in NewScientist Magazine (13.Aug'05,p27#2512) (Spot the Bots with Brains).
- Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability (2005)
  Reviews of book [P05uaibook]: ACM Reviews (27.Apr 2005,#CR131175), Artificial Intelligence Journal (2006), amazon.com (2004–2008)
- Intelligent Machines that Learn Unaided (2004)
  Ticino Ricerca, Project of the month 9 (2004)
- Diversity Trumps Fitness (2001)
  Paper [P02fuss] featured in the Technology Research News Magazine (2001)
- About Infinity in Computer Science (2001)
  Paper [P02fast] featured in SuperEva
- Universal Artificial Intelligence based on Algorithmic Probability (2001)
  Paper [P01aixi] intensely discussed in the AGI mailing list (2002, 2003-2006) and also in the commp.ai.philosophy newsgroup and ai-philosophy Yahoo group (2005–2006).

Attended Conferences / Workshops / Symposia / ...

(usually presenting a paper, organizing, PC member|chair, and/or invited.)
(see corresponding section for details)
2019: etc ...
2018: AII&OP
2017: IJCAI, AGI, AIS@Japan
2016: ISTAS
2015: ICML, EWRL
2014: ALT, AMPC

2013: MaxEnt, STF, RL@Dagstuhl
2012: CiE, ICML+EWRL, AusSS, CHASS
2011: UAI+IJCAI, ECML, EWRL, ALT, NIPS
2010: AGI, CMGM80
2009: AGI, UAI+COLT, Singularity, AusAI, NIPS
2008: ISBA, CEC, PCAR
2007: COLT, ALT, NIPS
2006: ALT, TAMC, Benelearn, Dagstuhl
2005: AISTATS, COLT, ICML
2004: COLT+ICML, ECML, ALT
2003: COLT+ICML, KI, Dagstuhl
2002: CEC, COLT, NIPS
2001: TAI, EWRL, ECML, NIPS
...

(Invited) Short Research Visits

(usually for a couple of days and accompanied by a talk)

- Universal Bayesian Agents: Theory and Applications
  University of Cambridge (UC 2011), Cambridge
- Universal Artificial Intelligence
  Japanese AGI Research Institute (Dwango 2017), Tokyo
  University of Heidelberg (ML&TCS@UoH 2015), Heidelberg
  Dalle Molle Institute for Artificial Intelligence (IDSIA 2015), Lugano
  AgroParisTech (MIA@APT 2012), Paris
  University of Oxford (FHI@UoO 2012), Oxford
  CSIT Monash University (2010), Melbourne
- Foundations of Intelligent Systems
  Swiss Federal Institute of Technology Zurich (ML@ETH 2011), Zürich
  Max Planck Institute for Intelligent Systems (MPI 2011), Stuttgart
  RMIT University (2010), Melbourne
- Predictive Hypothesis Identification
  National University of Singapore (COMP@NUS 2010), Singapore
- Learning to Predict with MDL & Bayes
  National University of Singapore (DSAP@NUS 2010), Singapore
- Introduction to and Applications of Algorithmic Information Theory
  Australian National University (MSI@ANU 2009), Australia
- Generic Reinforcement Learning Agents
  University of Technology, Sydney (UTS 2010), Sydney
  California Institute of Technology (CALTECH 2009), Pasadena (Los Angeles)
  Max Planck Institute for Biological Cybernetics (MPI 2009), Tübingen
  Swiss Federal Institute of Technology Zurich (ETHZ 2009), Zürich
  University of Technology (TUM 2009), Munich
  Dalle Molle Institute for Artificial Intelligence (IDSIA 2009), Lugano

University of New South Wales (UNSW 2009), Sydney

- On Universal Induction and Intelligent Agents
  Australian National University (MSI 2010 and RSSS 2008), Australia
- Bayes-Optimal Policies in General Environments
  University of Alberta (UA 2007), Edmonton
- On the Philosophical, Statistical, and Computational Foundation of Inductive Inference
  University of Queensland (UQLD 2008), Brisbane
  University of Alberta (UA 2007), Edmonton
- On Universal Prediction and Bayesian Confirmation
  Swiss Federal Institute of Technology Zurich (ETHZ 2006), Znrich
- Bayesian PC-Regression for Detecting Aberrations in DNA of Cancer Cells
  Swiss Federal Institute of Technology Zurich (ETHZ 2006), Zürich
  Dalle Molle Institute for Artificial Intelligence (IDSIA 2010), Lugano
- Bayesian and Universal Induction
  Swiss Federal Institute of Technology Zurich (ETHZ 2006), Znrich
- Universal Prediction: Concepts, Tools and Applications
  Oncology Institute of Southern Switzerland & Dalle Molle Institute for Artificial Intelligence (IOSI/IDSIA 2005), Lugano
- Foundations of Machine Learning = Information + Decision Theory
  University of Alberta (UA 2007), Edmonton
  Australian National University (ANU 2005), Canberra
- Fast/Exact Non-Parametric Bayesian Inference on Infinite Trees
  Queensland University of Technology (QUT 2010), Brisbane.
  University of Sydney (USYD 2005), Sydney
- Optimal Sequential Decisions Based on Algorithmic Probability
  Swiss Federal Institute of Technology Zurich (ETHZ 2004), Znrich
  California Institute of Technology (CALTECH 2003), Pasadena (Los Angeles)
- Bayesian Mutual Information and Robust Feature Selection
  Ludwig-Maximilian Univerity Munich (LMU 2004), Munich
- MDL Predictions based on Kolmogorov Complexity
  Boston University (BU 2003), Boston
- Towards a Universal Theory of Artificial Intelligence based on Algorithmic Probability and Sequential Decisions
  Workshop on Universal Learning Algorithms and Optimal Search (NIPS-2002), Vancouver
  University of Queensland (UQLD 2002), Brisbane
  Monash University (2002), Melbourne
  University of New South Wales (UNSW 2002), Sydney
  Australian National University (ANU 2002), Canberra
  Boston University (BU 2002), Boston
  Centrum voor Wiskunde en Informatica (CWI 2002), Amsterdam
- The Fastest and Shortest Algorithm for All Well-Defined Problems

University of New South Wales (UNSW 2005), Sydney
California Institute of Technology (CALTECH 2003), Pasadena (Los Angeles)
Centrum voor Wiskunde en Informatica (CWI 2001), Amsterdam

- New Error Bounds for Solomonoff Prediction
  University of Technology Munich (TUM 2000), Munich
- A Theory of Universal Artificial Intelligence based on Algorithmic Complexity
  Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA 2000), Lugano
  University of Technology (TUM 2000), Munich
- Instantons in QCD: Theory and Application of the Instanton Liquid Model
  University of Tel Aviv (1995), Tel Aviv
- Instantons and Meson Correlators in QCD
  CERN (1995), Geneve

Lecturing

Full & Shared Courses for Students

- Foundations of Artificial Intelligence
  Winter Semester (2010 & 2012 & 2013 & 2015 & 2017) ANU, Canberra, Lectures
- Theory of Computation
  Summer Semester (2014 & 2016 & 2018) ANU, Canberra, Lectures
- Introduction to Artificial Intelligence (2 weeks)
  Summer Semester (2007 & 2008 & 2009 & 2010 & 2011 & 2013 & 2014 & 2015 & 2016) ANU, Canberra, Lectures
- Information Theory (1 week)
  Winter Semester (2012 & 2013 & 2014 & 2015 & 2016) ANU, Canberra, Lectures
- Guest Lectures in various ANU courses (1 hour) and Summer/Winter Schools
  BIOL3191 (2014 & 2015 & 2016), COMP1130 (2014 & 2015), SCOM6027 (2015), AIOC (2015), COMP3620 (2017 & 2018), COMP2610 (2018), ...
- AI&RL&AIT&Logic Reading Groups
  Every Wednesday 11:00-12:30 (2009–2018), ANU, Organizer
- Reinforcement Learning and Planning under Uncertainty
  Winter Semester (2008) NICTA & ANU, Canberra, Lectures
- Introduction to Statistical Machine Learning (2 weeks)
  Summer Semester (2007 & 2008) NICTA & ANU, Canberra, Lectures
- Combinatorics and Probability (2 weeks)
  Winter Semester (2006) Australian National University, Canberra, Lectures
- Universal Artificial Intelligence: Math. and Phil. Foundations
  Helsinki Graduate School in CS&E (HeCSE 2006), Helsinki, Lectures
- Theory Reading Group
  Every Wednesday 15:00-16:30 (2004-2005), IDSIA, Organizer

- Algorithmic Information Theory and Machine Learning
  Winter Semester (2003), University of Technology Munich, Lectures
- Quantum Electro Dynamics
  Summer Semester (1995), LM-University Munich, Chief Tutor
- Theoretical Mechanics
  Winter Semester (1993), LM-University Munich, Tutor

Short Tutorials for Students

- 2015... *etc.*
- Universal Reinforcement Learning
  Reinforcement Learning Seminar (Dagstuhl 2013), Germany, Tutorial
- One Decade of Universal Artificial Intelligence
  6th Conf. on Artificial General Intelligence (2013), Beijing, Tutorial
- Foundations of Machine Learning [slides,video]
  Machine Learning Summer School (2008), ANU/RSISE/NICTA, Tutorial
- Introduction to Statistical Machine Learning [slides,video]
  Machine Learning Summer School (2008&2009&2010), ANU/RSISE/NICTA, Tutorial
- Foundations of Intelligent Agents
  First Summer School (ACISS'09) at AusAI Conference (2009) Melbourne, Tutorial
- Universal Artificial Intelligence [slides,video]
  Conf. on Artificial General Intelligence (2010), Lugano, Tutorial
  Summer Schools of Logic & Learning (2009), ANU/RSISE/NICTA Canberra, Lecture
- On the Philosophical, Statistical, and Computational Foundations of Inductive Inference and Intelligent Agents
  International Conference on Algorithmic Information Theory (2007), Sendai, Tutorial
- How to Predict with Bayes, MDL, and Experts [slides,video]*International Conf. on Machine Learning (2005), Bonn, Tutorial*
  Machine Learning Summer School (2005), ANU/RSISE/NICTA Canberra, Lectures

Invited Lectures at Conferences & Workshops

- Observer Localization in Multiverse Theories
  Algorithmic Information, Induction and Observers in Physics (AII&OP 2018), Waterloo, Canada
- Society in the wake of Super-Intelligent Machines
  AI and Society symposium (AIS 2017), Tokyo, Japan
- On Thompson Sampling and Asymptotic Optimality

Intl. Joint Conf. on AI (IJCAI 2017), Melbourne, Australia

- Advances in Universal Artificial Intelligence
  Intl. Conf. on Artificial General Intelligence (AGI 2017), Melbourne, Australia
- Unifying Foundations for Intelligent Agents
  Invitational Symposium On Trusted Autonomous Systems (ISTAS 2016), Barossa Valley, Australia
- Universal Reinforcement Learning
  European Workshop on Reinforcement Learning (EWRL 2015), Lille, France
- The Technological Singularity
  Surviving the Apocalypse (RLC@ANU 2014), Canberra, Australia
- Soft Aspects of Hard Intelligence [video]
  Workshop on Computational Creativity, Concept Invention, and General Intelligence (C3GI@IJCAI 2013), Beijing, China
- Uncertainty and Induction in AGI [video]
  Workshop on Probability Theory or Not (ProbOrNot 2013), Beijing, China
- Ingredients of Super-Intelligent Machines
  Conf. on Science Technology Future (STF 2013), RMIT University, Melbourne, Australia
- Observer Localization in Multiverse Theories
  33rd Intl. Workshop on Bayesian Inference and Maximum Entropy (MaxEnt 2013), Canberra, Australia
- The Technological Singularity
  In Canberra Tonight: A Vision of the Future (2013), Shine Dome, ANU, Canberra, Australia
  Conf. on Science Technology Future (STF December 2013), RMIT University, Melbourne, Australia
  Inaugural CHASS National Forum: The Human Dimension (2012), University of Canberra, Australia [video]
- Can Intelligence Explode? [slides+audio]
  University of Alberta (UA 2012 & 2013), Edmonton
- Can Intelligence Explode? *and* Universal AI *and* Participation in three panel discussions *Singularity Summit Australia (AusAI 2012), RMIT University, Melbourne, Australia*
- Foundations of Induction
  PhiMaLe NIPS Workshop (2011), Sierra Nevada, Spain
- Formalizing Intelligence and the Human Knowledge Compression Prize
  JTF Workshop on the Foundational Questions in the Mathematical Sciences (2011), Traunkirchen, Austria
- Foundations of Intelligent Agents [slides,video]
  Singularity Summit (2009), New York
- Foundations of Rational Agents
  Second International Symposium on Practical Cognitive Agents and Robots

(PCAR 2008), University of Technology, Sydney

- The Fastest and Shortest Algorithm for All Well-Defined Problems
  5th Turing Days Conference on Randomness and Complexity (2006), Bilgi University, Istanbul
- On the Foundations of Universal Sequence Prediction
  Symposium on Theory and Applications of Models of Computation (TAMC-2006), Learning Theory Session, Beijing
- Universal Artificial Intelligence
  The National Conference for Computing Students (CompCon 2013), Canberra, Australia [slides,video]
  Workshop Toward a Serious Computational Science of Intelligence (SCSI@AGI 2010), Lugano [slides,video]
  Swiss Mathematical Society, Fall Meeting (SMS 2005), Lugano
- Theoretically Optimal Program Induction and Universal Artificial Intelligence
  Inductive Programming Workshop W1 at (ICML-2005), Bonn
- MDL Predictions based on Kolmogorov Complexity
  Centennial Seminar on Kolmogorov Complexity and Applications (2003), Dagstuhl
- On the Existence and Convergence of Universal Priors
  Workshop on Computability and Randomness (2003), Uni-Heidelberg
- Solomonoff Induction and the Foundations of Occam's, Epicurus', Bayes', and Utility Principles
  Workshop on Foundations of Occam's razor (NIPS-2001), Vancouver
- An effective Procedure for Speeding up Algorithms
  Conference on Mathematical Approaches to Biological Computation (MaBiC-2001), Lavin
  Workshop on Algorithmic Information Theory (TAI-2001), Porquerolles
- Universal Sequential Decisions in Unknown Environments
  Workshop on Universal Learning Algorithms and Optimal Search (NIPS-2002), Vancouver
  5th European Workshop on Reinforcement Learning (EWRL-2001), Utrecht

Talks at Conferences

- Universal Compression of Piecewise i.i.d. Sources
  Data Compression Conference (DCC-2018), Snowbird
- Offline to Online Conversion *and* Extreme State Aggregation beyond MDPs
  25th Intl. Conf. on Algorithmic Learning Theory (ALT-2014), Bled
- A Mathematical Definition of Intelligence
  3rd Australian Mathematical Psychology Conference (AMPC-2014), Canberra
- Unifying Probability and Logic for Learning
  Workshop on Weighted Logics for AI (WL4AI@IJCAI-2013), Beijing

- Ray Solomonoff's Legacy [video]
  Conference on Artificial General Intelligence (AGI-2010), Lugano
- Observer Localization in Multiverse Theories
  Conference in Honor of Murray Gell-Mann's 80th Birthday (CMGM80-2010), Singapore
- Principled Large-Scale POMDP Learning
  Symposium on Partially Observable Reinforcement Learning (PORL-2009) at NIPS, Vancouver
- Feature Markov Decision Processes [video]
  2nd Conf. on Artificial General Intelligence (AGI-2009), Arlington
- Feature Dynamic Bayesian Networks
  2nd Conf. on Artificial General Intelligence (AGI-2009), Washington
- The Loss Rank Principle for Model Selection
  20th Annual Conf. on Learning Theory (COLT-2007), San Diego
- General Discounting versus Average Reward
  16th International Conf. on Algorithmic Learning Theory (ALT-2006), Barcelona
- Universal Learning of Repeated Matrix Games
  Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn-2006), Ghent
- Fast Non-Parametric Bayesian Inference on Infinite Trees
  15th International Conference on Artificial Intelligence and Statistics (AISTATS-2005), Barbados
- Universal Convergence of Semimeasures on Individual Random Sequences
  15th International Conf. on Algorithmic Learning Theory (ALT-2004), Padova
  Kolmogorov Complexity and Applications (Dagstuhl-2006), Germany
- Prediction with Expert Advice by Following the Perturbed Leader for General Weights
  15th International Conf. on Algorithmic Learning Theory (ALT-2004), Padova
- Online Prediction - Bayes versus Experts
  EU PASCAL Workshop (LTBIP-2004), London
- Self-Optimizing and Pareto-Optimal Policies in General Environments based on Bayes-Mixtures
  15th Annual Conference on Computational Learning Theory (COLT-2002), Sydney
- Fitness Uniform Selection to Preserve Genetic Diversity
  Congress on Evolutionary Computation (CEC-2002), Honolulu
  Conference of the European Chapter on Combinatorial Optimization (ECCO-2002), Lugano
- Distribution of Mutual Information
  14th Conference on Neural Information Processing Systems (NIPS-2001),

Vancouver

- General Loss Bounds for Universal Sequence Prediction
  18th International Conference on Machine Learning (ICML-2001), Williamstown
- Towards a Universal Theory of Artificial Intelligence based on Algorithmic Probability and Sequential Decisions
  12th European Conference on Machine Learning (ECML-2001), Freiburg
- Convergence and Error Bounds for Universal Prediction of Nonbinary Sequences
  12th European Conference on Machine Learning (ECML-2001), Freiburg

Poster Presentations

- Universal Bayesian Rational Agents *and* On Universal Prediction and Bayesian confirmation. *33rd Intl. Workshop on Bayesian Inference and Maximum Entropy (MaxEnt 2013), Canberra, Australia*
- Discrete MDL Predicts in Total Variation
  23rd Conference on Neural Information Processing Systems (NIPS 2009), Vancouver
- An Improved Bayesian Method for DNA Copy Number Estimation
  9th ISBA World Meeting (ISBA 2008), Hamilton Island
- Universal Bayesian Solution to the Induction Problem
  9th ISBA World Meeting (ISBA 2008), Hamilton Island
- Temporal Difference Updating without a Learning Rate
  21st Conference on Neural Information Processing Systems (NIPS 2007), Vancouver
- Bayesian Regression of Piecewise Constant Functions
  ISBA 8th International Meeting on Bayesian Statistics (ISBA 2006), Benidorm
- Sequence Prediction based on Monotone Complexity
  16th Annual Conf. on Learning Theory (COLT 2003), Washington, DC

More Lectures

- Artificial Intelligence and Society (Talks&Podium Discussion)
  Regularly: Bruce/Barton&Garran/other Halls & at other Events (ANU 2008-...), Canberra
- On Science, Fiction, and Future Reality
  The American School of Switzerland (TASIS 2005 & 2006), Lugano
- Various Lectures in the Theory Reading Group
  Dalle Molle Institute for Artificial Intelligence (IDSIA 2004-2005), Lugano
- The Pencil Beam Algorithm in RadioTherapy
  Company BrainLAB (1999-2000), Munich

- Parallel Algorithms in Fluid Mechanics
  Ferienakademie, TU-München, Infomatik (1991), Maria Laach
- Many other lectures at employed places (BrainLAB, IDSIA, ANU, ...)

Past & Current Group Members

Past and current PostDocs

2017 - 2018 Badri Vellambi - *Unifying Foundations for Intelligent Agents*, RSCS@ANU
2012 - 2014 Peter Sunehag - *Feature Reinforcement Learning*, RSCS@ANU
2009 - 2012 Peter Sunehag - *From Universal Induction to Intelligent Systems*, RSISE/ANU/NICTA
2008 - 2009 Rakib Ahmed - *Image-based Car Damage Detection*, RSISE@ANU
2005 - 2007 Daniil Ryabko - *Optimal Rational AIXI Agent based on Algorithmic Complexity*, IDSIA
2003 - 2005 Jan Poland - *Optimal Rational Agents in Unknown Environments*, IDSIA

Past and current PhD students

2018 - 2022 Elliot Catt - *Universal Artificial Intelligence*, RSCS@ANU
2015 - 2019 Sultan Majeed - *Extreme State Aggregation*, RSCS@ANU
2015 - 2018 Tom Everitt - *AGI Safety*, RSCS@ANU
2014 - 2016 Jan Leike - *Nonparametric General Reinforcement Learning*, RSCS@ANU
2012 - 2016 Yiyun Shou - *Causal Reasoning with Ambiguous Information*, PSY@ANU
2012 - 2015 Hadi Afshar - *Probabilistic Inference in Piecewise Graphical Models*, RSCS@ANU
2011 - 2014 Mayank Daswani - *Generic Reinforcement Learning Beyond Small MDPs*, RSCS@ANU
2010 - 2014 Di Yang - *Image-based Car Damage Detection*, RSISE@ANU
2010 - 2013 Tor Lattimore - *Theory of General Reinforcement Learning*, RSISE@ANU&NICTA
2009 - 2012 Phuong Nguyen - *Feature Reinforcement Learning Agents*, RSISE@ANU
2009 - 2012 Srimal Jayawardena - *Image-based Car Damage Detection*, RSISE@ANU
2009 - 2011 Matthew Robards - *Continuous-State Reinforcement Learning*, NICTA&ANU
2009 - 2011 Joel Veness - *Approximate Universal AI and Games*, UNSW

2009 - 2010	Ian Wood - *Information-Theoretic Foundations of Inductive Reasoning*, DCS@ANU
2008 - 2011	Nathan Brewer - *Image Processing and Computer Vision*, RSISE@ANU&NICTA
2006 - 2010	Paola Rancoita - *Bayesian Integrative Genomics*, IDSIA/IOSI
2003 - 2007	Shane Legg - *Machine Super Intelligence*, IDSIA (SIAI Award)

Past and current Master's students

2018/19	Michael Cohen, *Grue&Strong Asympt. Optimality*, MCOMP@ANU
2018/19	Tim McMahon, *Non-MDP Abstractions*, MCOMP@ANU
2018	David Quarel, *Injectivity of Solomonoff*, MCOMP@ANU
2017/18	Elliot Catt, *Quantum Universal Prediction*, MMATH@ANU
2016	Suraj Sasikumar, *Exploration in Feature Space for RL*,RSCS@ANU
2016	Boris Repasky, *Q-learning beyond MDPs*, RSCS@ANU
2016	Manlio Valenti, *KC&Comp. of Learning Agents*, UoT&ANU
2016	John Aslanides, *AIXIjs: A Software Demo for GRL*, RSCS@ANU
2016	Jarryd Martin, *Optimism and Death in RL*, RSCS@ANU
2012/13	Tom Everitt, KTH-Stockholm/ANU, *Universal Opt. - FLOUD*
2012/13	Marco Nembrini, ETHZ/ANU, *Sparse KT Estimator - SSDC*
2011-2013	Wen Shao, MPhil, RSCS@ANU, *Text Compression - LASC*
2008/9	Ke Zhang, RSISE/ANU/NICTA, *Outlier Detection - LDOF*
2007	Nathan Brewer, RSISE@ANU, *Dynamic Bayesian Networks*
2001	Daniele Pongan, ETHZ/IDSIA, *Evolutionary Algorithms - FUSS*
SS 1998	Hannes Mahlknecht, BrainLAB, *Voxel/Surface-Library*
WS 1997/8	Andreas Bertagnoll, BrainLAB, *Voxel/Surface-Library*
1994/5	Michael Birkel, LMU, *Particle physics*

Past and current other students (honors or interns or project)

2018/19	Nikhil Babu & Matthew Aitchison, Summer Scholar@ANU, *TMAZE*
2018/19	James Parker, Honours CS@ANU, *EXSAGG*
2018	Michael Cohen, CS project @ANU, *BOMAI*
2018	Samuel Yang-Zhao, Honours MSI@ANU, *RLWLINFA*
2017/18	Maleakhi Agung Wijaya, Summer Scholar @ANU, *PNA*
2017	Owen David Cameron, Honours MSI@ANU, *PIIDKKT*
2017	Mikael Böörs, BMath Thesis @ANU&LU, *2X2GAMES*
2017	Tobias Wängberg, BMath Thesis @ANU&LU, *2X2GAMES*
2016	James Bailie, UGrad Math project@ANU, *SAI*
2016	Sean Lamont, UGrad COMP2550 project@ANU, *EXPDISC*
2015	Matthew Alger, BSc Honours@ANU, UGrad Project, *DIRL*

2015	Daniel Filan, Honours, PhB@ANU, *SPEEDPRIORS*
2015	David Johnston, Grad COMP6470 Project, *EXSAGG*
2014/15	Alexander Mascolo, Honours, PhB@ANU, *DECISION*
2014/15	Trevor Rose, Summer Scholar @ ANU, *EXSAGG*
2014	Xi Li, PhD@PekingUni, 3 months visitor@ANU, *PHILUAI*
2014	Daniel Filan, 2x ASC Project, PhB@ANU, *UAI*
2014	Tom Shafron, ENGN4200 Project, PhB@ANU, *UAI*
2014	Alexander Mascolo, honours, PhB@ANU, *DECISION*
2014	Michael Buck Shlegeris (BSc Science@ANU) and Matthew Alger (BSc honours@ANU), UGrad COMP3740 Project, *RMODINF*
2013/14	Alexander Mascolo, Summer Scholar @ ANU, *TCDISC*
2013	Ian Hon, honours, MSI@ANU, *OPY*
2013	Johannes Kirschner, BSc Thesis, ETHZ/ANU, *RLFA*
2012/13	Ian Hon, Summer Scholar @ ANU, *UNILEARN*
2012/13	Daniel Nolan, Summer Scholar @ ANU, *CIID*
2011	Daniel Visentin, honours, PhB@ANU, *FRL*
2010/11	Jan Melchior, Intern from CE, SoCS@ANU, *ICAR*
2010	Mayank Daswani, honours, CS@ANU, *PHIMDP*
2010	Daniel Visentin, ASC Project, PhB@ANU, *MC-AIXI-CTW*
2010	Alexander O'Neill, honours, CS@ANU, *ADAPCTW*
2010	Samuel Rathmanner, honours, CS@ANU, *UIPHIL*
2009/10	Rachel Bunder, Intern @ ANU from Uni Wollongong, *PHIMDPX*
2009/10	Mayank Daswani, Summer Scholar, CSL@ANU, *PHIMDP*
2009	Tor Lattimore, honours, MSI@ANU, *EVENBITS*
2008	Tor Lattimore, Project, FEIT@ANU, *UIvNFL*
2007	Kassel Hingee, Project, MSI@ANU, *SELECT*
2007	Minh Ngoc Tran, Intern, Vietnam student @ NICTA, *LORPC*
2007	Tiago da Silva, Intern, Brazilian student @ NICTA, *AIXIFOE*
2004	Akshat Kumar, Project, IIT Guwahai, India *FUSSEXP*

Publications

Most articles are available online at http://www.hutter1.net/. Generally accessible articles are marked with a°. Some key publications are highlighted by a *. They include a book [P05uaibook], with recent applications [P11aixictwx], some award-winning papers [P07pcregx, P09phimdpx, P11aixictwx, P16thompgrl, P16selfmod], my physics [P97family] and AI [P01aixi] ideas I'm most proud of, my most offbeat paper [P02fast], my most cited paper [P07iorx], a nice demo of AIXI and related models [P17urlsurexp], a patent [P02uspatent], and my first publication [P87cad]. Publications in top conferences in (theoretical) computer science are of equal rank to journal publications. For a comprehensive list of Conference and Journal reputations see e.g. http://www.arc.gov.au/era/era_journal_list.htm.

Monographs and Edited Books

[P13alttcs] M. Hutter, F. Stephan, V. Vovk, and T. Zeugmann. ALT'10 special issue. *Theoretical Computer Science*, 473:1–3/178, 2013.

[P11ewrlproc] M. Hutter and S. Sanner, editors. *European Workshop on Reinforcement Learning*, volume 7188 of *LNAI*, Athens, 2011. Springer.

[P10altproc] M. Hutter, F. Stephan, V. Vovk, and T. Zeugmann, editors. *Algorithmic Learning Theory*, volume 6331 of *LNAI*, Canberra, 2010. Springer.

[P10agiproc] E. Baum, M. Hutter, and E. Kitzelmann, editors. *Artificial General Intelligence*. Atlantis Press, Lugano, 2010.

[P09alttcs] M. Hutter and R. A. Servedio, editors. ALT'07 special issue. *Theoretical Computer Science*, 410(19):1747–1748/1912, 2009.

[P09agiproc] B. Goertzel, P. Hitzler, and M. Hutter, editors. *Artificial General Intelligence*. Atlantis Press, Arlington, 2009.

[P07altproc] M. Hutter, R. A. Servedio, and E. Takimoto, editors. *Algorithmic Learning Theory*, volume 4754 of *LNAI*, Sendai, 2007. Springer.

[P05uaibook]* M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. 300 pages. Springer, Berlin, 2005. http://www.hutter1.net/ai/uaibook.htm.

Journal papers

[P18align] T. Everitt and M. Hutter. The alignment problem for history-based Bayesian reinforcement learners. *submitted*, ???, 2018. First winner of the AI alignment prize round 2

[P18off2onx] M. Hutter. Tractability of batch to sequential conversion. *Theoretical Computer Science*, 733:71–82, 2018.

[P18aixicplexx] J. Leike and M. Hutter. On the computability of Solomonoff induction and AIXI. *Theoretical Computer Science*, 716:28–49, 2018.

[P16exsaggx]* M. Hutter. Extreme state aggregation beyond Markov decision processes. *Theoretical Computer Science*, 650:73–91, 2016.

[P15ratagentx]* P. Sunehag and M. Hutter. Rationality, optimism and guarantees in general reinforcement learning. *Journal of Machine Learning Research*, 16:1345–1390, 2015.

[P15mnonconvx] T. Lattimore and M. Hutter. On Martin-löf (non)convergence of Solomonoff's universal mixture. *Theoretical Computer Science*, 588:2–15, 2015.

[P15aixiprior] J. Leike and M. Hutter. Bad universal priors and notions of optimality. *Journal of Machine Learning Research, W&CP: COLT*, 40:1244–1259, 2015. Also presented at EWRL'15.

[P14pacmdpx] T. Lattimore and M. Hutter. Near-Optimal PAC bounds for discounted MDPs. *Theoretical Computer Science*, 558:125–143, 2014.

[P14tcdiscx] T. Lattimore and M. Hutter. General time consistent discounting. *Theoretical Computer Science*, 519:140–154, 2014.

[P13pacgrl] T. Lattimore, M. Hutter, and P. Sunehag. The sample-complexity of general reinforcement learning. *Journal of Machine Learning Research, W&CP: ICML*, 28(3):28–36, 2013.

[P13sad] M. Hutter. Sparse adaptive Dirichlet-multinomial-like processes. *Journal of Machine Learning Research, W&CP: COLT*, 30:432–459, 2013.

[P13problogic]* M. Hutter, J.W. Lloyd, K.S. Ng, and W.T.B. Uther. Probabilities on sentences in an expressive logic. *Journal of Applied Logic*, 11:386–420, 2013.

[P12lstphi] M. Daswani, P. Sunehag, and M. Hutter. Feature reinforcement learning using looping suffix trees. *Journal of Machine Learning Research, W&CP*, 24:11–23, 2012.

[P12singularity]° M. Hutter. Can intelligence explode? *Journal of Consciousness Studies*, 19(1-2):143–166, 2012.

[P11uiphil]° S. Rathmanner, M. Hutter. A philosophical treatise of universal induction. *Entropy*, 16(6):1076–1136, 2011.

[P11aixictwx]* J. Veness, K. S. Ng, M. Hutter, W. Uther, and D. Silver. A Monte-Carlo AIXI approximation. *Journal of Artificial Intelligence Research*, 40:95–142, 2011. Honorable Mention for the 2014 IJCAI-JAIR Best Paper Prize.

[P10ctoex]° M. Hutter. A complete theory of everything (will be subjective). *Algorithms*, 3(4):329–350, 2010.

[P10cnlohx] P. M. V. Rancoita, M. Hutter, F. Bertoni, and I. Kwee. An integrated Bayesian analysis of LOH and copy number data. *BMC Bioinformatics*, 11(321) 1–18, 2010.

[P10lorpx] M. Hutter and M. Tran. Model selection with the loss rank principle. *Computational Statistics and Data Analysis*, 54:1288–1306, 2010.

[P09phimdpx]* M. Hutter. Feature reinforcement learning: Part I: Unstructured MDPs. *Journal of Artificial General Intelligence*, 1:3–24, 2009.

[P09aixiopen] M. Hutter. Open problems in universal induction & intelligence. *Algorithms*, 3(2):879–906, 2009.

[P09improbx] A. Piatti, M. Zaffalon, F. Trojani, and M. Hutter. Limits of learning about a categorical latent variable under prior near-ignorance. *International Journal of Approximate Reasoning*, 50(4):597–611, 2009.

[P09idmx]* M. Hutter. Practical robust estimators under the Imprecise Dirichlet Model. *International Journal of Approximate Reasoning*, 50(2):231–242, 2009.

[P09bcna] P. M. V. Rancoita, M. Hutter, F. Bertoni, and I. Kwee. Bayesian DNA copy number analysis. *BMC Bioinformatics*, 10(10):1–19, 2009.

[P08actoptx] D. Ryabko and M. Hutter. On the possibility of learning in reactive environments with arbitrary dependence. *Theoretical Computer Science*, 405(3):274–284, 2008.

[P08pquestx] D. Ryabko and M. Hutter. Predicting non-stationary processes. *Applied Mathematics Letters*, 21(5):477-482, 2008.

[P07iorx]* S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. *Minds & Machines*, 17(4):391–444, 2007.

[P07pcregx]* M. Hutter. Exact Bayesian regression of piecewise constant functions. *Bayesian Analysis*, 2(4):635–664, 2007. Lindley prize for innovative research in Bayesian statistics.

[P07uspx]* M. Hutter. On universal prediction and Bayesian confirmation. *Theoretical Computer Science*, 384(1):33–48, 2007.

[P07mlconvxx] M. Hutter and An. A. Muchnik. On Semimeasures Predicting Martin-Löf Random Sequences. *Theoretical Computer Science*, 382(3):247–261, 2007.

[P07postbndx] A. Chernov, M. Hutter, and J. Schmidhuber. Algorithmic complexity bounds on future prediction errors. *Information and Computation*, 205(2):242–261, 2007.

[P06unipriorx] M. Hutter. On generalized computable universal priors and their convergence. *Theoretical Computer Science*, 364(1):27–41, 2006.

[P06fuo]* M. Hutter and S. Legg. Fitness uniform optimization. *IEEE Transactions on Evolutionary Computation*, 10(5):568–589, 2006.

[P06knapsack]* M. Mastrolilli and M. Hutter. Hybrid rounding techniques for knapsack problems. *Discrete Applied Mathematics*, 154(4):640–649, 2006.

[P06unimdlx] M. Hutter. Sequential predictions based on algorithmic complexity. *Journal of Computer and System Sciences*, 71(1):95–117, 2006.

[P06mdlspeedx] J. Poland and M. Hutter. MDL convergence speed for Bernoulli sequences. *Statistics and Computing*, 16(2):161–175, 2006.

[P05tree] M. Zaffalon and M. Hutter. Robust inference of trees. *Annals of Mathematics and Artificial Intelligence*, 45:215–239, 2005.

[P05expertx]* M. Hutter and J. Poland. Adaptive online prediction by following the perturbed leader. *Journal of Machine Learning Research*, 6:639–660, 2005.

[P05mdl2px]* J. Poland and M. Hutter. Asymptotics of Discrete MDL for Online Prediction. *IEEE Transactions on Information Theory*, 51(11):3780–3795, 2005.

[P05mifs]* M. Hutter and M. Zaffalon. Distribution of mutual information from complete and incomplete data. *Computational Statistics & Data Analysis*, 48(3):633–657, 2005.

[P03optisp]* M. Hutter. Optimality of universal Bayesian prediction for general loss and alphabet. *Journal of Machine Learning Research*, 4:971–997, 2003.

[P03spupper]* M. Hutter. Convergence and loss bounds for Bayesian sequence prediction. *IEEE Transactions on Information Theory*, 49(8):2061–2067, 2003.

[P02fast]* M. Hutter. The fastest and shortest algorithm for all well-defined problems. *International Journal of Foundations of Computer Science*, 13(3):431–443, 2002.

[P01errbnd] M. Hutter. New error bounds for Solomonoff prediction. *Journal of Computer and System Sciences*, 62(4):653–667, 2001.

[P97instanto] M. Hutter. Instantons and meson correlators in QCD. *Zeitschrift für Physik*, C74:131–143, 1997.

[P97family]* A. Blumhofer and M. Hutter. Family structure from periodic solutions of an improved gap equation. *Nuclear Physics*, B484:80–96, 1997. Missing figures in B494 (1997) 485.

[P96eta] M. Hutter. The mass of the η' in self-dual QCD. *Physics Letters*, B367:275–278, 1996.

Papers in refereed international conference proceedings

[P19actagg] S. J. Majeed and M. Hutter. Performance guarantees for homomorphisms beyond markov decision processes. In *Proc. 33rd AAAI Conference on Artificial Intelligence (AAAI'19)*, Honolulu, USA, 2019. AAAI Press.

[P18agisafe] T. Everitt, G. Lea, and M. Hutter. AGI safety literature review. In *Proc. 27th International Joint Conf. on Artificial Intelligence (IJCAI'18)*, pages 5441–5449, Stockholm, Sweden, 2018. IJCAI Review Track.

[P18qnonmdp] S. J. Majeed and M. Hutter. On Q-learning convergence for non-Markov decision processes. In *Proc. 27th International Joint Conf. on Artificial Intelligence (IJCAI'18)*, pages 2546–2552, Stockholm, Sweden, 2018.

[P18convbinctw] B. N. Vellambi and M. Hutter. Convergence of binarized context-tree weighting for estimating distributions of stationary sources. In *Proc. IEEE International Symposium on Information Theory (ISIT'18)*, pages 731–735, Vail, USA, 2018. IEEE.

[P18piidkkt] B. N. Vellambi, O. Cameron, and M. Hutter. Universal compression of piecewise i.i.d. sources. In *Proc. Data Compression Conference (DCC'18)*, pages 267–276, Snowbird, Utah, USA, 2018. IEEE Computer Society.

[P17thompgrls] J. Leike, T. Lattimore, L. Orseau, and M. Hutter. On Thompson sampling and asymptotic optimality. In *Proc. 26th International Joint Conf. on Artificial Intelligence (IJCAI'17)*, pages 4889–4893, Melbourne, Australia, 2017. Best sister conferences paper track.

[P17corruptrl] T. Everitt, V. Krakovna, L. Orseau, M. Hutter, and S. Legg. Reinforcement learning with corrupted reward signal. In *Proc. 26th International Joint Conf. on Artificial Intelligence (IJCAI'17)*, pages 4705–4713, Melbourne, Australia, 2017.

[P17cbefsrl] J. Martin, S. N. Sasikumar, T. Everitt, and M. Hutter. Count-based exploration in feature space for reinforcement learning. In *Proc. 26th International Joint Conf. on Artificial Intelligence (IJCAI'17)*, pages 2471–2478, Melbourne, Australia, 2017.

[P17urlsurexp]* J. Aslanides, J. Leike, and M. Hutter. Universal reinforcement learning algorithms: Survey and experiments. In *Proc. 26th International Joint Conf. on Artificial Intelligence (IJCAI'17)*, pages 1403–1410, Melbourne, Australia, 2017.

[P17offswitch] T. Wängberg, M. Böörs, E. Catt, T. Everitt, and M. Hutter. A game-theoretic analysis of the off-switch game. In *Proc. 10th Conf. on Artificial General Intelligence (AGI'17)*, volume 10414 of *LNAI*, pages 167–177, Melbourne, Australia, 2017. Springer.

[P17expdisc] S. Lamont, J. Aslanides, J. Leike, and M. Hutter. Generalised discount functions applied to a Monte-Carlo AIμ implementation. In *Proc. 16th Conf. on Autonomous Agents and MultiAgent Systems (AAMAS'17)*, pages 1589–1591, Sao Paulo, Brazil, 2017.

[P16aixideath] J. Martin, T. Everitt, and M. Hutter. Death and suicide in universal artificial intelligence. In *Proc. 9th Conf. on Artificial General Intelligence (AGI'16)*, volume 9782 of *LNAI*, pages 23–32, New York, USA, 2016. Springer.

[P16wirehead] T. Everitt and M. Hutter. Avoiding wireheading with value reinforcement learning. In *Proc. 9th Conf. on Artificial General Intelligence (AGI'16)*, volume 9782 of *LNAI*, pages 12–22, New York, USA, 2016. Springer.

[P16selfmod]* T. Everitt, D. Filan, M. Daswani, and M. Hutter. Self-modification of policy and utility function in rational agents. In *Proc. 9th Conf. on Artificial General Intelligence (AGI'16)*, volume 9782 of *LNAI*, pages 1–11, New York, USA, 2016. Springer. Kurzweil Prize for Best AGI Paper.

[P16vacrecog] B. Fernando, P. Anderson, M. Hutter, and S. Gould. Discriminative hierarchical rank pooling for activity recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'16)*, pages 1924–1932, Las Vegas, NV, USA, 2016. IEEE.

[P16thompgrl]* J. Leike, T. Lattimore, L. Orseau, and M. Hutter. Thompson sampling is asymptotically optimal in general environments. In *Proc. 32nd International Conf. on Uncertainty in Artificial Intelligence (UAI'16)*, pages 417–426, New Jersey, USA, 2016. AUAI Press. Best student paper.

[P16speedprior] D. Filan, J. Leike, and M. Hutter. Loss bounds and time complexity for speed priors. In *Proc. 19th International Conf. on Artificial Intelligence and Statistics (AISTATS'16)*, volume 51, pages 1394–1402, Cadiz, Spain, 2016. Microtome.

[P15sikscplex] J. Leike and M. Hutter. On the computability of Solomonoff induction and knowledge-seeking. In *Proc. 26th International Conf. on Algorithmic Learning Theory (ALT'15)*, volume 9355 of *LNAI*, pages 364–378, Banff, Canada, 2015. Springer.

[P15solraven] J. Leike and M. Hutter. Solomonoff induction violates Nicod's criterion. In *Proc. 26th International Conf. on Algorithmic Learning Theory (ALT'15)*, volume 9355 of *LNAI*, pages 349–363, Banff, Canada, 2015. Springer. Also presented at CCR: http://math.uni-heidelberg.de/logic/conferences/ccr2015/.

[P15seqdts] T. Everitt, J. Leike, and M. Hutter. Sequential extensions of causal and evidential decision theory. In *Proc. 4th International Conf. on Algorithmic Decision Theory (ADT'15)*, volume 9346 of *LNAI*, pages 205–221, Lexington, USA, 2015. Springer.

[P15learncnf] J. Veness, M. Hutter, L. Orseau, and M. Bellemare. Online learning of k-CNF boolean functions. In *Proc. 24th International Joint Conf. on Artificial Intelligence (IJCAI'15)*, pages 3865–3873, Buenos Aires, Argentina, 2015. AAAI Press.

[P15agproblaws] P. Sunehag and M. Hutter. Using localization and factorization to reduce the complexity of reinforcement learning. In *Proc. 8th Conf. on Artificial General Intelligence (AGI'15)*, volume 9205 of *LNAI*, pages 177–186, Berlin, Germany, 2015. Springer.

[P15aixicplex] J. Leike and M. Hutter. On the computability of aixi. In *Proc. 31st International Conf. on Uncertainty in Artificial Intelligence (UAI'15)*, pages 464–473, Amsterdam, Netherlands, 2015. AUAI Press.

[P15cnc] J. Veness, M. Bellemare, M. Hutter, A. Chua, and G. Desjardins. Compress and control. In *Proc. 29th AAAI Conference on Artificial Intelligence (AAAI'15)*, pages 3016–3023, Austin, USA, 2015. AAAI Press.

[P14martosc] J. Leike and M. Hutter. Indefinitely oscillating martingales. In *Proc. 25th International Conf. on Algorithmic Learning Theory (ALT'14)*, volume 8776 of *LNAI*, pages 321–335, Bled, Slovenia, 2014. Springer.

[P14off2on] M. Hutter. Offline to online conversion. In *Proc. 25th International Conf. on Algorithmic Learning Theory (ALT'14)*, volume 8776 of *LNAI*, pages 230–244, Bled, Slovenia, 2014. Springer.

[P14exsagg] M. Hutter. Extreme state aggregation beyond MDPs. In *Proc. 25th International Conf. on Algorithmic Learning Theory (ALT'14)*, volume 8776 of *LNAI*, pages 185–199, Bled, Slovenia, 2014. Springer.

[P14pacbayes] T. Lattimore and M. Hutter. Bayesian reinforcement learning with exploration. In *Proc. 25th International Conf. on Algorithmic Learning Theory (ALT'14)*, volume 8776 of *LNAI*, pages 170–184, Bled, Slovenia, 2014. Springer.

[P14learnutm] P. Sunehag and M. Hutter. Intelligence as inference or forcing Occam on the world. In *Proc. 7th Conf. on Artificial General Intelligence (AGI'14)*, volume 8598 of *LNAI*, pages 186–195, Quebec City, Canada, 2014. Springer.

[P14optcog] P. Sunehag and M. Hutter. A dual process theory of optimistic cognition. In *Proc. 36th Annual Meeting of the Cognitive Science Society (CogSci'14)*, pages 2949–2954, Quebec City, Canada, 2014. Curran Associates.

[P14floud] T. Everitt, T. Lattimore, and M. Hutter. Free lunch for optimisation under the universal distribution. In *Proc. 2014 Congress on Evolutionary Computation (CEC'14)*, pages 167–174, Beijing, China, 2014. IEEE.

[P13ksaprob]* L. Orseau, T. Lattimore, and M. Hutter. Universal knowledge-seeking agents for stochastic environments. In *Proc. 24th International Conf. on Algorithmic Learning Theory (ALT'13)*, volume 8139 of *LNAI*, pages 158–172, Singapore, 2013. Springer, Berlin.

[P13ccbayessp] T. Lattimore, M. Hutter, and P. Sunehag. Concentration and confidence for discrete bayesian sequence predictors. In *Proc. 24th International Conf. on Algorithmic Learning Theory (ALT'13)*, volume 8139 of *LNAI*, pages 324–338, Singapore, 2013. Springer, Berlin.

[P13agscilaws] P. Sunehag and M. Hutter. Learning agents with evolving hypothesis classes. In *Proc. 6th Conf. on Artificial General Intelligence (AGI'13)*, volume 7999 of *LNAI*, pages 150–159. Springer, Heidelberg, 2013.

[P13mnonconv] T. Lattimore and M. Hutter. On Martin-löf convergence of Solomonoff's mixture. In *Proc. 10th Annual Conference on Theory and Applications of Models of Computation (TAMC'13)*, volume 7876 of *LNCS*, pages 212–223, Hong Kong, China, 2013. Springer.

[P12aixiens] J. Veness, P. Sunehag, and M. Hutter. On ensemble techniques for AIXI approximation. In *Proc. 5th Conf. on Artificial General Intelligence (AGI'12)*, volume 7716 of *LNAI*, pages 341–351. Springer, Heidelberg, 2012.

[P12aixiopt] P. Sunehag and M. Hutter. Optimistic AIXI. In *Proc. 5th Conf. on Artificial General Intelligence (AGI'12)*, volume 7716 of *LNAI*, pages 312–321. Springer, Heidelberg, 2012.

[P12pacmdp] T. Lattimore and M. Hutter. PAC bounds for discounted MDPs. In *Proc. 23rd International Conf. on Algorithmic Learning Theory (ALT'12)*, volume 7568 of *LNAI*, pages 320–334, Lyon, France, 2012. Springer, Berlin.

[P12ctmrl] P. Nguyen, P. Sunehag, and M. Hutter. Context tree maximizing reinforcement learning. In *Proc. 26th AAAI Conference on Artificial Intelligence (AAAI'12)*, pages 1075–1082, Toronto, 2012. AAAI Press.

[P12ctswitch] J. Veness, K. S. Ng, M. Hutter, and M. Bowling. Context tree switching. In *Proc. Data Compression Conference (DCC'12)*, pages 327–336, Snowbird, Utah, 2012. IEEE Computer Society.

[P12adapctw]* A. O'Neill, M. Hutter, W. Shao, and P. Sunehag. Adaptive context tree weighting. In *Proc. Data Compression Conference (DCC'12)*, pages 317–326, Snowbird, Utah, 2012. IEEE Computer Society.

[P11evenbits] T. Lattimore, M. Hutter, and V. Gavane. Universal prediction of selected bits. In *Proc. 22nd International Conf. on Algorithmic Learning Theory (ALT'11)*, volume 6925 of *LNAI*, pages 262–276, Espoo, Finland, 2011. Springer, Berlin.

[P11asyoptag] T. Lattimore and M. Hutter. Asymptotically optimal agents. In *Proc. 22nd International Conf. on Algorithmic Learning Theory (ALT'11)*, volume 6925 of *LNAI*, pages 368–382, Espoo, Finland, 2011. Springer, Berlin.

[P11tcdisc]* T. Lattimore and M. Hutter. Time consistent discounting. In *Proc. 22nd International Conf. on Algorithmic Learning Theory (ALT'11)*, volume 6925 of *LNAI*, pages 383–397, Espoo, Finland, 2011. Springer, Berlin.

[P11aixiaxiom] P. Sunehag and M. Hutter. Axioms for rational reinforcement learning. In *Proc. 22nd International Conf. on Algorithmic Learning Theory (ALT'11)*, volume 6925 of *LNAI*, pages 338–352, Espoo, Finland, 2011. Springer, Berlin.

[P10phimp] P. Sunehag and M. Hutter. Consistency of feature Markov processes. In *Proc. 21st International Conf. on Algorithmic Learning Theory (ALT'10)*, volume 6331 of *LNAI*, pages 360–374, Canberra, 2010. Springer, Berlin.

[P10aixictw] J. Veness, K. S. Ng, M. Hutter, and D. Silver. Reinforcement learning via AIXI approximation. In *Proc. 24th AAAI Conference on Artificial Intelligence (AAAI'10)*, pages 605–611, Atlanta, 2010. AAAI Press.

[P09mdltvp]* M. Hutter. Discrete MDL predicts in total variation. In *Advances in Neural Information Processing Systems 22, (NIPS'09)*, pages 817–825, Cambridge, MA, 2009. MIT Press.

[P09phimdp] M. Hutter. Feature Markov decision processes. In *Proc. 2nd Conf. on Artificial General Intelligence (AGI'09)*, volume 8, pages 61–66. Atlantis Press, 2009. First runner-up for the Kurzweil best paper prize.

[P09phidbn] M. Hutter. Feature dynamic Bayesian networks. In *Proc. 2nd Conf. on Artificial General Intelligence (AGI'09)*, volume 8, pages 67–73. Atlantis Press, 2009.

[P08select] K. Hingee and M. Hutter. Equivalence of probabilistic tournament and polynomial ranking selection. In *Proc. 2008 Congress on Evolutionary Computation (CEC'08)*, pages 564–571, Hongkong, 2008. IEEE.

[P07qlearn] M. Hutter and S. Legg. Temporal difference updating without a learning rate. In *Advances in Neural Information Processing Systems 20 (NIPS'07)*, pages 705–712, Cambridge, MA, 2007. MIT Press.

[P07pcreg] M. Hutter. Bayesian regression of piecewise constant functions. In *Proc. 8th International Meeting on Bayesian Statistics (ISBA'06)*, pages 607–612, Benidorm, 2007. Oxford University Press. Lindley prize for innovative research in Bayesian statistics.

[P07pquest] D. Ryabko and M. Hutter. On sequence prediction for arbitrary measures. In *Proc. IEEE International Symposium on Information Theory (ISIT'07)*, pages 2346–2350, Nice, France, 2007. IEEE.

[P07lorp]* M. Hutter. The loss rank principle for model selection. In *Proc. 20th Annual Conf. on Learning Theory (COLT'07)*, volume 4539 of *LNAI*, pages 589–603, San Diego, 2007. Springer, Berlin.

[P07improb] A. Piatti, M. Zaffalon, F. Trojani, and M. Hutter. Learning about a categorical latent variable under prior near-ignorance. In *Proc. 5th International Symposium on Imprecise Probability: Theories and Applications (ISIPTA'07)*, pages 357–364, Prague, 2007.

[P06discount] M. Hutter. General discounting versus average reward. In *Proc. 17th International Conf. on Algorithmic Learning Theory (ALT'06)*, volume 4264 of *LNAI*, pages 244–258, Barcelona, 2006. Springer, Berlin.

[P06actopt] D. Ryabko and M. Hutter. Asymptotic learnability of reinforcement problems with arbitrary dependence. In *Proc. 17th International Conf. on Algorithmic Learning Theory (ALT'06)*, volume 4264 of *LNAI*, pages 334–347, Barcelona, 2006. Springer, Berlin.

[P06usp] M. Hutter. On the foundations of universal sequence prediction. In *Proc. 3rd Annual Conference on Theory and Applications of Models of Computation (TAMC'06)*, pages 408–420, Beijing, 2006.

[P06robot] V. Zhumatiy and F. Gomez and M. Hutter and J. Schmidhuber. Metric state space reinforcement learning for a vision-capable mobile robot. In *Proc. 9th International Conference on Intelligent Autonomous Systems (IAS'06)*, 272–281, Tokyo, 2006.

[P05postbnd] A. Chernov and M. Hutter. Monotone conditional complexity bounds on future prediction errors. In *Proc. International Conf. on Algorithmic Learning Theory (ALT'05)*, pages 414–428, Singapore, 2005.

[P05actexp2] J. Poland and M. Hutter. Defensive universal learning with experts. In *Proc. International Conf. on Algorithmic Learning Theory (ALT'05)*, pages 356–370, Singapore, 2005.

[P05fuds] S. Legg and M. Hutter. Fitness uniform deletion for robust optimization. In *Proc. Genetic and Evolutionary Computation Conference (GECCO'05)*, pages 1271–1278, Washington, OR, 2005.

[P05bayestree] M. Hutter. Fast non-parametric Bayesian inference on infinite trees. In *Proc. 15th International Conference on Artificial Intelligence and Statistics (AISTATS'05)*, pages 144–151, Barbados, 2005.

[P04mlconvx] M. Hutter and An. A. Muchnik. Universal convergence of semimeasures on individual random sequences. In *Proc. 15th International Conf. on Algorithmic Learning Theory (ALT'04)*, volume 3244 of *LNAI*, pages 234–248, Padova, 2004. Springer, Berlin.

[P04expert] M. Hutter and J. Poland. Prediction with expert advice by following the perturbed leader for general weights. In *Proc. 15th International Conf. on Algorithmic Learning Theory (ALT'04)*, volume 3244 of *LNAI*, pages 279–293, Padova, 2004. Springer, Berlin.

[P04mdlspeed] J. Poland and M. Hutter. On the convergence speed of MDL predictions for Bernoulli sequences. In *Proc. 15th International Conf. on Algorithmic Learning Theory (ALT'04)*, volume 3244 of *LNAI*, pages 294–308, Padova, 2004. Springer, Berlin.

[P04mdl2p] J. Poland and M. Hutter. Convergence of discrete MDL for sequential prediction. In *Proc. 17th Annual Conf. on Learning Theory (COLT'04)*, Lecture Notes in Artificial Intelligence, pages 300–314, Berlin, 2004. Springer.

[P04fussexp] S. Legg, M. Hutter, and A. Kumar. Tournament versus fitness uniform selection. In *Proceedings of the 2004 Congress on Evolutionary Computation (CEC'04)*, pages 564–571. IEEE, 2004.

[P03unimdl] M. Hutter. Sequence prediction based on monotone complexity. In *Proceedings of the 16th Annual Conference on Learning Theory (COLT'03)*, Lecture Notes in Artificial Intelligence, pages 506–521, Berlin, 2003. Springer.

[P03unipriors] M. Hutter. On the existence and convergence of computable universal priors. In R. Gavaldá, K. P. Jantke, and E. Takimoto, editors, *Proceedings of the 14th International Conference on Algorithmic Learning Theory (ALT'03)*, volume 2842 of *LNAI*, pages 298–312, Berlin, 2003. Springer.

[P03mimiss] M. Hutter and M. Zaffalon. Bayesian treatment of incomplete discrete data applied to mutual information and feature selection. In R. Kruse A. Günter and B. Neumann, editors, *Proceedings of the 26th German Conference on Artificial Intelligence (KI'03)*, volume 2821 of *Lecture Notes in Computer Science*, pages 396–406, Heidelberg, 2003. Springer.

[P03idm]* M. Hutter. Robust estimators under the Imprecise Dirichlet Model. In J.-M. Bernard, T. Seidenfeld, and M. Zaffalon, editors, *Proceedings of the 3nd International Symposium on Imprecise Probabilities and Their Application (ISIPTA'03)*, volume 18 of *Proceedings in Informatics*, pages 274–289, Canada, 2003. Carleton Scientific.

[P02feature] M. Zaffalon and M. Hutter. Robust feature selection by mutual information distributions. In A. Darwiche and N. Friedman, editors, *Proceedings of the 18th International Conference on Uncertainty in Artificial Intelligence (UAI'02)*, pages 577–584, San Francisco, CA., 2002. Morgan Kaufmann.

[P02selfopt]* M. Hutter. Self-optimizing and Pareto-optimal policies in general environments based on Bayes-mixtures. In *Proceedings of the 15th Annual Conference on Computational Learning Theory (COLT'02)*, Lecture Notes in Artificial Intelligence, pages 364–379, Sydney, Australia, 2002. Springer.

[P01xentropy] M. Hutter. Distribution of mutual information. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14 (NIPS'01)*, pages 399–406, Cambridge, MA, 2002. MIT Press.

[P02fuss] M. Hutter. Fitness uniform selection to preserve genetic diversity. In X. Yao, editor, *Proceedings of the 2002 Congress on Evolutionary Computation (CEC'02)*, pages 783–788. IEEE, 2002.

[P01market] I. Kwee, M. Hutter, and J. Schmidhuber. Market-based reinforcement learning in partially observable worlds. *Proceedings of the International Conference on Artificial Neural Networks (ICANN'01)*, pages 865–873, 2001.

[P01loss] M. Hutter. General loss bounds for universal sequence prediction. *Proceedings of the 18th International Conference on Machine Learning (ICML'01)*, pages 210–217, 2001.

[P01alpha] M. Hutter. Convergence and error bounds of universal prediction for general alphabet. *Proceedings of the 12th Eurpean Conference on Machine Learning (ECML'01)*, pages 239–250, 2001.

[P01aixi]* M. Hutter. Towards a universal theory of artificial intelligence based on algorithmic probability and sequential decisions. *Proceedings of the 12th Eurpean Conference on Machine Learning (ECML'01)*, pages 226–238, 2001.

Book chapters and extended encyclopedic entries

[P18uaitas]° T. Everitt and M. Hutter. Universal artificial intelligence: Practical agents and fundamental challenges. In Hussein A. Abbass, Jason Scholz, and Darryn J. Reid, editors, *Foundations of Trusted Autonomy*, chapter 2, pages 15–46. Springer, 2018.

[P17unilearn]° M. Hutter. Universal learning theory. In C. Sammut and G. Webb, editors, *Encyclopedia of Machine Learning*, pages 1295–1304. Springer, 2nd edition, 2017.

[P15aitcog] P. Sunehag and M. Hutter. Algorithmic complexity. In James D. Wright, editor, *International Encyclopedia of the Social & Behavioral Sciences*, volume 1, pages 534–538. Elsevier, 2nd edition, 2015.

[P12ctoe2]° M. Hutter. The subjective computable universe. In *A Computable Universe: Understanding and Exploring Nature as Computation*, pages 399–416. World Scientific, 2012.

[P12uaigentle]° M. Hutter. One decade of universal artificial intelligence. In *Theoretical Foundations of Artificial General Intelligence*, pages 67–88. Atlantis Press, 2012.

[P11randai]° M. Hutter. Algorithmic randomness as foundation of inductive reasoning and artificial intelligence. In *Randomness through Computation*, pages 159–169, World Scientific, 2011.

[P11unilearn]° M. Hutter. Universal learning theory. In C. Sammut and G. Webb, editors, *Encyclopedia of Machine Learning*, pages 1001–1008. Springer, 2011.

[P08kolmo]° M. Hutter. Algorithmic complexity. *Scholarpedia*, 3(1):2573, 2008.

[P07aixigentle] M. Hutter. Universal algorithmic intelligence: A mathematical top→down approach. In *Artificial General Intelligence*, pages 227–290. Springer, Berlin, 2007.

[P07algprob]° M. Hutter, S. Legg, and P. M. B. Vitányi. Algorithmic probability. *Scholarpedia*, 2(8):2572, 2007.

[P07idefs]° S. Legg and M. Hutter. A collection of definitions of intelligence. In *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms*, volume 157 of *Frontiers in Artificial Intelligence and Applications*, pages 17–24, Amsterdam, NL, 2007. IOS Press.

[P07ait]° M. Hutter. Algorithmic information theory: a brief non-technical guide to the field. *Scholarpedia*, 2(3):2519, 2007.

Other refereed publications

[P15metasearch1] T. Everitt and M. Hutter. Analytical results on the BFS vs. DFS algorithm selection problem. Part i: Tree search. In *Proc. 28th Australasian Joint Conference on Artificial Intelligence (AusAI'15)*, volume 9457 of *LNAI*, pages 157–165, Canberra, Australia, 2015. Springer.

[P15metasearch2] T. Everitt and M. Hutter. Analytical results on the BFS vs. DFS algorithm selection problem. Part ii: Graph search. In *Proc. 28th Australasian Joint Conference on Artificial Intelligence (AusAI'15)*, volume 9457 of *LNAI*, pages 166–178, Canberra, Australia, 2015. Springer.

[P14rladvice] M. Daswani, P. Sunehag, and M. Hutter. Reinforcement learning with value advice. In *Proc. 6th Asian Conf. on Machine Learning (ACML'14)*, volume 39, pages 299–314, Canberra, Australia, 2014. JMLR.

[P14reflect] D. Yang, S. Jayawardena, S. Gould, and M. Hutter. Reflective features detection and hierarchical reflections separation in image sequences. In *Proc. 16th International Conf. on Digital Image Computing: Techniques and Applications (DICTA'14)*, pages 1–7, Wollongong, Australia, 2014. IEEE Xplore.

[P14epipolar] S. Jayawardena, S. Gould, H. Li, M. Hutter, and R. Hartley. Reliable point correspondences in scenes dominated by highly reflective and largely homogeneous surfaces. In *Proc. 12th Asian Conf. on Computer Vision – Workshop (RoLoD@ACCV'14) Part I*, volume 9008 of *LNCS*, pages 659–674, Singapore, 2014. Springer.

[P13rlqh] M. Daswani, P. Sunehag, and M. Hutter. Q-learning for history-based reinforcement learning. In *Proc. 5th Asian Conf. on Machine Learning (ACML'13)*, volume 29, pages 213–228, Canberra, Australia, 2013. JMLR.

[P13problogics] M. Hutter, J.W. Lloyd, K.S. Ng, and W.T.B. Uther. Unifying probability and logic for learning. In *Proc. 2nd Workshop on Weighted Logics for AI (WL4AI'13)*, pages 65–72, Beijing, China, 2013.

[P12windowkt] P. Sunehag, W. Shao, and M. Hutter. Coding of non-stationary sources as a foundation for detecting change points and outliers in binary time-series. In *Proc. 10th Australasian Data Mining Conference (AusDM'12)*, volume 134, pages 79–84, Sydney, 2012. Australian Computer Society.

[P12optopt] P. Sunehag and M. Hutter. Optimistic agents are asymptotically optimal. In *Proc. 25th Australasian Joint Conference on Artificial Intelligence (AusAI'12)*, volume 7691 of *LNAI*, pages 15–26, Sydney, 2012. Springer.

[P12watershed] D. Yang and S. Gould and M. Hutter. A noise tolerant watershed transformation with viscous force for seeded image segmentation. In *Proc. 11th*

Asian Conf. on Computer Vision (ACCV'12), volume 7724 of *LNCS*, pages 775–789, Daejeon, Korea, 2012. Springer.

[P11uivnfl] T. Lattimore and M. Hutter. No free lunch versus Occam's razor in supervised learning. In *Proc. Solomonoff 85th Memorial Conference*, volume 7070 of *LNAI*, pages 223–235, Melbourne, Australia, 2011. Springer.

[P11aixiaxiom2] P. Sunehag and M. Hutter. Principles of Solomonoff induction and AIXI. In *Proc. Solomonoff 85th Memorial Conference*, volume 7070 of *LNAI*, pages 386–398, Melbourne, Australia, 2011. Springer.

[P11unipreq] I. Wood, P. Sunehag, and M. Hutter. (Non-)equivalence of universal priors. In *Proc. Solomonoff 85th Memorial Conference*, volume 7070 of *LNAI*, pages 417–425, Melbourne, Australia, 2011. Springer.

[P11frlexp] P. Nguyen, P. Sunehag, and M. Hutter. Feature reinforcement learning in practice. In *Proc. 9th European Workshop on Reinforcement Learning (EWRL-9)*, volume 7188 of *LNAI*, pages 66–77. Springer, September 2011.

[P11segm3d] S. Jayawardena, D. Yang, and M. Hutter. 3d model assisted image segmentation. In *Proc. 13th International Conf. on Digital Image Computing: Techniques and Applications (DICTA'11)*, pages 51–58, Noosa Heads, Australia, 2011. IEEE Xplore.

[P11losspose] S. Jayawardena, M. Hutter, and N. Brewer. A novel illumination-invariant loss for monocular 3d pose estimation. In *Proc. 13th International Conf. on Digital Image Computing: Techniques and Applications (DICTA'11)*, pages 37–44, Noosa Heads, Australia, 2011. IEEE Xplore.

[P09wheel] M. Hutter and N. Brewer. Matching 2-d ellipses to 3-d circles with application to vehicle pose estimation. In *Proc. 25th Conf. Image and Vision Computing New Zealand (IVCNZ'09)*, page 153–158, Wellington, New Zealand, 2009. IEEE Xplore.

[P09mbpcrcode] P.M.V. Rancoita and M. Hutter. mBPCR: A package for DNA copy number profile estimation. *BioConductor – Open Source Software for BioInformatics*, 0.99:1–25, 2009.

[P09cnloh] P.M.V. Rancoita, M. Hutter, F. Bertoni, and I. Kwee. Bayesian joint estimation of CN and LOH aberrations. In *Proc. 3rd International Workshop on Practical Applications of Computational Biology & Bioinformatics (IW-PACBB'09)*, volume 5518 of *LNCS*, pages 1109–1117, Spain, 2009. Springer.

[P09ldof] K. Zhang, M. Hutter, and W. Jin. A new local distance-based outlier detection approach for scattered real-world data. In *Proc. 13th Pacific-Asia Conf. on Knowledge Discovery and Data Mining (PAKDD'09)*, volume 5467 of *LNAI*, pages 813–822, Bangkok, 2009. Springer.

[P07intest] S. Legg and M. Hutter. Tests of machine intelligence. In *50 Years of Artificial Intelligence*, volume 4850 of *LNAI*, pages 232–242, Monte Verita, Switzerland, 2007.

[P06ior] S. Legg and M. Hutter. A formal measure of machine intelligence In *Proc. 15th Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn'06)*, pages 73–80, Ghent, 2006.

[P06iors2] S. Legg and M. Hutter. A formal definition of intelligence for artificial systems. In *Proc. 50th Anniversary Summit of Artificial Intelligence*, pages 197–198, Monte Verita, Switzerland, 2006.

[P06aixifoe] J. Poland and M. Hutter Universal learning of repeated matrix games In *Proc. 15th Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn'06)*, pages 7–14, Ghent, 2006.

[P05iors] S. Legg and M. Hutter. A universal measure of intelligence for artificial agents. In *Proc. 21st International Joint Conf. on Artificial Intelligence (IJCAI'05)*, pages 1509–1510, Edinburgh, 2005.

[P05actexp] J. Poland and M. Hutter. Master algorithms for active experts problems based on increasing loss values. In *Proc. 14th Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn'05)*, pages 59–66, Enschede, 2005.

[P05mdlreg] J. Poland and M. Hutter. Strong asymptotic assertions for discrete MDL in regression and classification. In *Annual Machine Learning Conference of Belgium and The Netherlands (Benelearn'05)*, pages 67–72, Enschede, 2005.

[P03mlconv] M. Hutter. An open problem regarding the convergence of universal a priori probability. In *Proceedings of the 16th Annual Conference on Learning Theory (COLT'03)*, volume 2777 of *LNAI*, pages 738–740, Berlin, 2003. Springer.

[P01grep] I. Kwee, M. Hutter, and J. Schmidhuber. Gradient-based reinforcement planning in policy-search methods. In *Proceedings of the 5th European Workshop on Reinforcement Learning (EWRL-5)*, volume 27 of *Cognitieve Kunstmatige Intelligentie*, pages 27–29, 2001.

[P01decision] M. Hutter. Universal sequential decisions in unknown environments. In *Proceedings of the 5th European Workshop on Reinforcement Learning (EWRL-5)*, volume 27 of *Cognitieve Kunstmatige Intelligentie*, pages 25–26, 2001.

Theses and other student work

[P03habil]* M. Hutter. Optimal sequential decisions based on algorithmic probability, 288 pages, Habilitation, 2003.

[P96thesis]* M. Hutter. *Instantons in QCD: Theory and application of the instanton liquid model*. PhD thesis, Faculty for Theoretical Physics, LMU Munich, 1996. Translation from the German original.

[P96diss] M. Hutter. *Instantonen in der QCD: Theorie und Anwendungen des Instanton-Flüssigkeit-Modells*. PhD thesis, Fakultät für Theoretische Physik, LMU München, 1996.

[P91cfs] M. Hutter. Implementierung eines Klassifizierungs-Systems. Master's thesis, Theoretische Informatik, TU München, 1992. 72 pages with C listing, in German.

[P90fopra] M. Hutter. A reinforcement learning Hebb net. Fortgeschrittenenpraktikum, Theoretische Informatik, TU München, 1990. 30 pages with Pascal listing, in German.

[P90faka] M. Hutter. Parallele Algorithmen in der Strömungsmechanik. Ferienakademie: Numerische Methoden der Strömungsmechanik, Universität Erlangen-Nürnberg & Technische Universität München, 1990. 10 pages, in German.

Other publications

[P14frlabs] M. Daswani, P. Sunehag, and M. Hutter. Feature reinforcement learning: State of the art. In *Proc. Workshops at the 28th AAAI Conference on Artificial Intelligence: Sequential Decision Making with Big Data*, pages 2–5, Quebec City, Canada, 2014. AAAI Press.

[P14ktoptdif] T. Alpcan, T. Everitt, and M. Hutter. Can we measure the difficulty of an optimization problem? In *IEEE Information Theory Workshop*, pages 356–360, Hobart, Australia, 2014. IEEE Press.

[P13uai4lay]° M. Hutter. To create a super-intelligent machine, start with an equation. *The Conversation*, November(29):1–5, 2013.

[P13ewrlabs] P. Auer, M. Hutter, and L. Orseau, editors. *Reinforcement Learning*, volume 3:8, Dagstuhl, Germany, 2013. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik.

[P12ssdc] J. Veness and M. Hutter. Sparse Sequential Dirichlet Coding. Technical Report arXiv:1206.3618, UoA and ANU, Canada and Australia, 2012.

[P10lorpc] M.-N. Tran and M. Hutter. Model Selection by Loss Rank for Classification and Unsupervised Learning. Technical Report arXiv:1011.1379, NUS and ANU, Singapore and Australia, 2010.

[P10pdpx] W. Buntine and M. Hutter. A Bayesian review of the Poisson-Dirichlet process. Technical Report arXiv:1007.0296, NICTA and ANU, Australia, 2010.

[P10ctoe] M. Hutter. Observer localization in multiverse theories. *Proceedings of the Conference in Honour of Murray Gell-mann's 80th Birthday*, pages 638–645. World Scientific, 2010.

[P10altintro] M. Hutter, F. Stephan, V. Vovk, and T. Zeugmann. Algorithmic learning theory 2010: Editors' introduction. In *Proc. 21st International Conf. on Algorithmic Learning Theory (ALT'10)*, volume 6331 of *LNAI*, pages 1–10, Canberra, 2010. Springer, Berlin.

[P09bayestreex] M. Hutter. Exact non-parametric Bayesian inference on infinite trees. Technical Report arXiv:0903.5342, NICTA and ANU, Australia, 2009.

[P08phi] M. Hutter. Predictive Hypothesis Identification. Presented at 9th Valencia /ISBA 2010 Meeting, arXiv:0809.1270, 2008.

[P07altintro] M. Hutter, R. A. Servedio, and E. Takimoto. Algorithmic Learning Theory 2007: Editors' introduction. In *Proc. 18th International Conf. on Algorithmic Learning Theory (ALT'07)*, volume 4754 of *LNAI*, pages 1–8, Sendai, 2007. Springer, Berlin.

[P06kcdagabs] M. Hutter, W. Merkle, and P. M. B. Vitányi, editors. *Kolmogorov Complexity and Applications*, number 06051 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2006. IBFI. http://drops.dagstuhl.de/portals/06051.

[P04bayespea] M. Hutter. Online prediction – Bayes versus experts. Technical report, July 2004. Presented at the *EU PASCAL Workshop on Learning Theoretic and Bayesian Inductive Principles (LTBIP'04)*, 4 pages.

[P04mdp] S. Legg and M. Hutter. Ergodic MDPs admit self-optimising policies. Technical Report IDSIA-21-04, IDSIA, 2004.

[P04env] S. Legg and M. Hutter. A taxonomy for abstract environments. Technical Report IDSIA-20-04, IDSIA, 2004.

[P00speed] M. Hutter. An effective procedure for speeding up algorithms. *Presented at the 3rd Workshop on Algorithmic Information Theory (TAI'01)*, 10 pages, 2001.

[P00kcunai] M. Hutter. A theory of universal artificial intelligence based on algorithmic complexity. Technical Report arXiv:cs.AI/0004001, 62 pages, 2000.

[P95spin] M. Hutter. Proton spin in the instanton background. Technical Report LMU-95-15, Theoretische Physik, LMU München, 1995. 15 pages.

[P95prop] M. Hutter. Gauge invariant quark propagator in the instanton background. Technical Report LMU-95-03, Theoretische Physik, LMU München, 1995. 15 pages.

[P93gluon] M. Hutter. Gluon mass from instantons. Technical Report LMU-93-18, Theoretische Physik, LMU München, 1993. 13 pages.

[P87cad]❄ M. Hutter. Fantastische 3D-Graphik mit dem CPC-Giga-CAD. *7. Schneider Sonderheft, Happy Computer, Sonderheft 16*, pages 41–92, 1987.

Patents

[P02uspatent]* M. Hutter. System and method for analysing and displaying two- or three-dimensional sets of data. BrainLAB*, US patent*, number US2002041701, pages 1–15, 2002.
http://l2.espacenet.com/espacenet/bnsviewer?CY=ep&LG=en&
DB=EPD&PN=US2002041701&ID=US2002041701A1+I+.

[P01eupatent] M. Hutter. Stufenfreie Darstellung von zwei- oder dreidimensionalen datensätzen durch krümmungsminimierende Verschiebung von Pixelwerten. BrainLAB*, EU patent*, number EP1184812, pages 1–19, 2001.
http://l2.espacenet.com/espacenet/bnsviewer?CY=ep&LG=en&
DB=EPD&PN=EP1184812&ID=EP+++1184812A1+I+.