Curriculum Vitae — Wouter M. Koolen

1. Personal details

Title(s), initial(s), surname: prof dr W.M. Koolen
Web site: `http://wouterkoolen.info/`
Google Scholar profile: `https://scholar.google.com/citations?user=34JTfUcAAAAJ`

2. Current Positions

Senior Researcher (0.8fte)
Machine Learning Group
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands

Full Professor of Mathematical Machine Learning (0.2fte)
Statistics Group, Mathematics of Operations Research Department
Faculty of Electrical Engineering, Mathematics and Computer Science
University of Twente

3. Education

PhD Computer Science, cum laude
Institute for Logic, Language and Computation (ILLC), University of Amsterdam
Date of PhD award: 27/01/2011
Supervisors: prof. P. M. B. Vitányi and prof. P. D. Grünwald
Thesis title: Combining Strategies Efficiently: High-quality Decisions from Conflicting Advice

MSc Logic, cum laude
Institute for Logic, Language and Computation (ILLC), University of Amsterdam
Date of MSc award: 21/12/2006
Main subject: Computer Science (Machine Learning)
Supervisor: prof. P. D. Grünwald

4. Research summary

Studying the foundations of machine learning as a sequential decision-making problem, I developed elegant characterisations of the exact **boundary of learnability** for problems in non-uniform/non-stationary learning and investment (NeurIPS 2013, TCS 2014, JMLR 2016). I extended the toolbox for learning in **large structured domains** with efficient methods for succinctly describable polytopes, (quantum) density matrices, (spectral) directions, and (non-parametric) isotonic regression (NeurIPS 2011, 2017; COLT 2010, 2013, 2016). I discovered that **exact minimax optimal** algorithms for linear regression and time series can be implemented in linear (and not exponential) time (NeurIPS 2014, 2015; COLT 2015). In these problems NP-hardness often looms nearby — as an obstacle to some version of the problem formulation (TCS 2014), or as a barrier that can be circumvented (NeurIPS 2012). *The overall picture that emerges is a rich collection of deep mathematical tools and flexible algorithmic techniques brought forth by addressing in detail a spectrum of natural tasks.*

I also developed principled adaptive algorithms with **data-dependent individual-sequence guarantees** and demonstrated that these subsume and strengthen performance guarantees derived earlier from stochastic assumptions (NeurIPS 2011, 2014, 2016, 2016, JMLR 2014, COLT 2015). Concurrently, I fostered the development of learning and optimisation communities around these concepts by organising NeurIPS workshops in 2013 and 2015, and through invited talks at Lorentz workshops in 2014 and 2016. I also released open-source (MIT license) implementations of the core algorithms MetaGrad and Squint. *Following my lead, the theory community working on online learning has shifted focus towards refined adaptive algorithms.*

From 2015 I started working on fundamental **interactive learning** problems, including the pure exploration formalisation of the Monte Carlo Tree Search problem. I see the latter as a complex yet structured intermediate problem between sequential testing (including classical best arm identification) and full reinforcement learning. This problem poses a challenge, as our tools were simply not ready to address it fully in the instance-optimal sense. I proposed and analysed new algorithms for the depth two case, and conjectured a path to instance-optimality (COLT 2016). I subsequently investigated strategies for game trees of arbitrary depth, which experimentally improved the state of the art sample complexity by an astounding factor 15 (NeurIPS spotlight 2017). I isolated and studied the simple yet fundamental problem of learning how the minimum mean compares to a given threshold (NeurIPS 2018). Our new curiously asymmetric Murphy Sampling rule does match the lower bound asymptotically and works well in practise. In (JMLR 2021) I develop the sharpest-known exponential-family deviation inequalities for self-normalised sums, and introduce the **pure exploration problem rank**, which I show governs the lower-order constants in the stopping rule. I develop new theory and algorithms for pure exploration problems with **multiple correct answers** (NeurIPS 2019). Finally, I initiated a new line of algorithms based on iterative saddle point solvers, for which I can show the first **instance-optimal finite-confidence guarantees** in pure exploration problems (NeurIPS 2019) and **instance-optimal finite-time regret bounds** in reward maximisation problems (ICML 2020). *Taking stock, we see that research on pure exploration problems has gained serious momentum. The community is picking this up, and is developing e.g. extensions to* linear contextual bandits *and* reward-free exploration.

5. Work experience since completing my PhD

Position	Period	FTE	Type	Institution
Postdoc	01/02/2011 – 31/01/2013	1	Fixed term	Royal Holloway, UK
Researcher	01/02/2013 – 31/08/2013	1	Fixed term	CWI
Postdoc	01/09/2013 – 31/07/2015	1	Fixed term	QUT, Brisbane, AU
Visiting scholar	01/09/2013 – 31/07/2015			UC Berkeley
Researcher	15/07/2015 – 15/04/2017	1	Fixed term	CWI
Senior researcher	15/04/2017 – 30/06/2018	1	Tenure track	CWI
Senior researcher	01/07/2018 – 30/05/2022	1	Permanent	CWI
Senior researcher	01/06/2022 – present	0.8	Permanent	CWI
Full Professor	01/06/2022 – present	0.2	Permanent	University of Twente

6. Academic staff supervised

	Name	Role
Postdocs (1)		
Sep 2018 – Sep 2019	R. Degenne	Supervisor
PhD Students (5)		
Started Sep 2022	É. Crepon (ENS Lyon)	Co-promotor
Started Sep 2021	T. Lardy (Leiden University)	Co-promotor
PhD Jul 2023	M. Pérez-Ortiz (CWI)	Co-promotor
PhD Sep 2021	Z. Mhammedi (ANU, Canberra)	Promotor
PhD Jan 2021	R. de Heide (Leiden University)	Co-promotor
PhD Interns (5)		
Nov 2021–Feb 2022	C. Fiegel (ENS Paris-Saclay)	Supervisor (4 months)
Sep 2021–Sep 2022	É. Crepon (ENS Paris-Saclay)	Supervisor (12 months)
Feb–Aug 2020	J-L. Rebuffi (ENS Lyon)	Supervisor (7 months)
May–Jul 2019	H. Shao (CUHK)	Supervisor (3 months)
Jan–Feb 2019	Z. Mhammedi (ANU, Canberra)	Co-supervisor (2 months)
Master Students (12)		
Started Feb 2024	T. Huitema (University of Twente)	Supervisor
MSc Jun 2023	F. Tuininga (University of Twente)	Supervisor
MSc Feb 2023	M. Meijer (UvA)	Supervisor
MSc Aug 2022	F. Gili (UvA)	Supervisor
MSc Aug 2022	J. Kowalska (UvA)	Supervisor
MSc Dec 2020	T. Schiet (TU Delft)	Supervisor
MSc Nov 2020	G. Jin (Leiden University)	Supervisor
May–Jun 2019	A. Simoes (ENS Lyon)	Internship host
MSc Sep 2018	F. Girotti (U. Milano)	Supervisor
MSc Oct 2017	M. Volkert (Leiden University)	Supervisor
MSc Oct 2017	M. Li (Leiden University)	Supervisor
MSc Dec 2016	J. Dubbeldam (Leiden University)	Supervisor

7. International activities

International study visits (selected)

- 21 – 22 Mar 2022: 4TUNE project meeting hosted by A. Taylor (INRIA Paris) with R. Degenne (INRIA Lille) and P. Gaillard (INRIA Grenoble).
- 29 – 30 Nov 2021: A. Taylor (INRIA Paris).
- 20 – 21 Sep 2021: A. Taylor (INRIA Paris).
- 8 – 9 Sep 2021: A. Garivier (ENS Lyon).
- 3 – 11 Jan 2020: S. Juneja (TIFR Mumbai).
- 15 – 19 Oct 2018: A. Garivier (ENS Lyon). *Includes one-day MCTS theory-practise meeting*.
- 14 – 18 May 2018: E. Kaufmann (INRIA Lille)
- 16 – 30 Mar 2018: A. Malek and A. Rakhlin (MIT), O. Shamir (Microsoft Research Cambridge)
- 26 – 29 Nov 2017: A. Garivier and S. Gerschinovitz (IMT, Toulouse)
- 9 – 14 Nov 2017: V. Vovk (Royal Holloway), T. Lattimore and Cs. Szépesvari (Deepmind)
- 30 Oct – 3 Nov 2017: ANR SPADRO workshop on Bandits in Sète, France (invitation only).
- 22 May – 2 Jun 2017: A. Garivier and S. Gerschinovitz (IMT, Toulouse)
- 3 – 14 Apr 2017: A. Garivier and S. Gerschinovitz (IMT, Toulouse)

International visitors (selected)

- 1 – 5 Nov 2021: G. Neu (Barcelona)
- 2 – 18 Oct 2019: S. Gerchinovitz (Toulouse)
- 12 – 15 Feb 2019: E. Kaufmann (INRIA Lille)
- 22 May – 31 Jul 2018: Z. Mhammedi (ANU, Canberra)
- 22 – 24 Apr 2018: C. Monteleoni (UC Boulder)
- 9 – 21 Apr 2018: D. Ostrovskii (SIERRA, INRIA Paris)
- 18 – 20 Mar 2018: E. Kaufmann and B. Guedj (INRIA Lille).
- 27 Feb – 2 Mar 2018: E. Kaufmann (INRIA Lille).
- 23 – 24 Nov 2017: Cs. Szepesvári (Deepmind).
- 18 – 22 Sep 2017: E. Kaufmann (INRIA Lille).

Organisation of Scientific Meetings

2023 **Organiser**. Research Semester Programme on Machine Learning Theory at CWI.
2022 **Organiser**. Complex Feedback in Online Learning workshop at ICML.
2022 **Organiser**. A/B Testing Knowledge Transfer Meeting at CWI
2020 **"Virtual" Core Team**. Conference on Learning Theory (COLT), online (422p)
2018 **Open Problems Chair**. Conference on Learning Theory (COLT), Stockholm (ca. 200p)
2017 **Co-organiser**. CWI Machine Learning Lectures, Amsterdam (ca. 236p)
2017 **Local Arrangements Chair**. Conference on Learning Theory (COLT), Amsterdam (165p, €70K)
2015 **Organiser**, NeurIPS Workshop *Learning Faster From Easy Data II*, Montreal (ca. 75p)
2013 **Organiser**, NeurIPS Workshop *Large Scale Matrix Analysis and Inference*, Lake Tahoe (ca. 80p)

2013 **Organiser**, NeurIPS Workshop *Learning Faster From Easy Data*, Lake Tahoe (ca. 60p)

Editorial Board

'19–'22 **Editor** SIAM Journal on Mathematics of Data Science (SIMODS)
2023– **Editor** Journal of Machine Learning Research (JMLR)

Conference Program Committees

Area Chair NeurIPS '18–'20, ICML 2019
PC Member COLT '15–'21, ICML '16–'18, ALT '12, '15–'19, IJCAI-ML '15

Invited Presentations

I gave plenary presentations of my accepted papers at NeurIPS (2017, 2016, 2012), COLT (2015, 2013, 2010, 2008) and ALT (2012, 2011, 2010).

Workshop Presentations

- 17/11/2021. *A/B/n Testing with Control in the Presence of Subpopulations*. Theory of Reinforcement Learning Workshop at the Simons Institute, UC Berkeley.
  21/10/2021. *A/B/n Testing with Control in the Presence of Subpopulations*. Google Learning Theory Workshop, online.
  26/10/2020. *Pure Exploration Problems*. Mathematics of Online Decision Making Workshop at the Simons Institute, UC Berkeley.
  04/03/2020. *Exploration and Exploitation in Structured Stochastic Bandits* . ELLIS Workshop on Interactive Learning and Interventional Representations, online.
- 03/01/2020. *Exploration and Exploitation in Structured Stochastic Bandits* . Workshop on Learning Theory 2, Mumbai, India.
- 25/09/2019. *Matching Regret Lower Bounds in Structured Stochastic Bandits*. Lancaster/Deepmind Multi Armed Bandit Workshop, London.
- 19/03/2019. *Self-Play in the age of Deep Learning*. Young European Statisticians Workshop, EURANDOM Eindhoven.
- 19/11/2018. *A Pure Exploration Perspective on Monte Carlo Tree Search*. Belgium Netherlands Workshop on Reinforcement Learning (BeNeRL 2018), Delft.
- 26/09/2018. *Sequential Test for the Lowest Mean: From Thompson to Murphy Sampling*. Second CWI-Inria Workshop, Paris.
- 11/09/2018. *Sequential Test for the Lowest Mean: From Thompson to Murphy Sampling*. Optimization and Learning CIMI Workshop, Toulouse.
- 04/04/2018. *Bandit Algorithms for Pure Exploration: Best Arm Identification and Game Tree Search*. Dutch Mathematical Congress (NMC) session on machine learning.
- 31/10/2017. *Towards Information-Theoretically Optimal Monte-Carlo Tree Search*. SPADRO Workshop on Bandits, Sète, France.
- 20/10/2017. *The Design of Adaptive Online Learning Methods*. Online Learning Workshop, Chaire ML for Big Data, Télécom ParisTech.
- 09/11/2016. *MetaGrad: Multiple Learning Rates in Online Learning*. Lorentz workshop on Theoretical Foundations for Learning from Easy Data, Leiden

- 01/04/2016. *MetaGrad: Faster Convergence Without Curvature in Online Convex Optimization*. Learning and Optimization workshop, DALI meeting, Sestri Levante, Italy

Industry Engagement

- 10/12/2020. *The Pure Exploration Renaissance*. Booking.com research, Amsterdam.
  15/09/2017. *Machine Learning Deep Dive*. KLM Operations Decision Support AMS/XX, Schiphol.

Tutorials

- 31/08/2020. *Online Learning and Bandits* (2h tutorial). Theory of Reinforcement Learning Boot Camp, Simons Institute, UC Berkeley
- 20/09/2016. *Online Learning: Robustness and Adaptivity* (3h tutorial). ELC fall school on Theoretical Computer Science, Shonan, Japan

Seminars at Leading Institutions

- 28/01/2022. *A/B/n Testing with Control in the Presence of Subpopulations*. SCOOL Seminar, INRIA Lille.
- 22/03/2019. *Pure Exploration with Multiple Correct Answers*. Google Brain, Zurich.
- 23/02/2018. *Bandit Algorithms for Best Arm Identification and Game Tree Search*. Machine Learning and Statistics for Structures, Thematic Seminar, University of Amsterdam.
- 06/04/2017. *Combining Adversarial Guarantees and Stochastic Fast Rates in Online Learning*. Paul Sabatier University, Toulouse
- 02/03/2017. *Maximin Action Identification*. Adobe Research, San Jose
- 27/02/2017. *Combining Adversarial Guarantees and Stochastic Fast Rates in Online Learning*. Simons Institute, UC Berkeley
- 03/10/2016. *Second-order Quantile Methods for Experts and Combinatorial Games*. Kyushu University, Fukuoka, Japan
- 30/06/2016. *Combining Adversarial Guarantees and Stochastic Fast Rates in Online Learning*. Seminar at CSAIL, MIT
- 15/04/2016. *MetaGrad: Multiple Learning Rates in Online Learning*. Seminar at SequeL, Inria Lille
- 24/02/2016. *MetaGrad: Multiple Learning Rates in Online Learning* Seminar at DIKU, University of Copenhagen
- 28/08/2015. *Second-order Quantile Methods for Experts and Combinatorial Games* Seminar at CLRC, Royal Holloway
- 11/05/2015. *Second-order Quantile Methods for Experts and Combinatorial Games* Machine Learning Seminar, UC Santa Cruz
- 06/11/2014. *Learning the Learning Rate for Prediction with Expert Advice*. Machine Learning Seminar, UC Berkeley
- 09/05/2014. *The Pareto Regret Frontier*. ML Seminar, University of Alberta, Edmonton
- 10/01/2014. *The Pareto Regret Frontier*. Seminar at NICTA, Sydney

8. Other academic activities

Peer review contributions

Annals of Statistics '20, IEEE Trans. Inf. Th. '20, '13, '11, ITW'20, Journal of the Royal Stat. Soc. '19, Bernoulli '19, STOC '18, Math. Prog. '17, JMLR '10–'18, Algorithmica '13, '14, NeurIPS '13–'16, ICML '15, COLT '13, '14, STACS '17, ICTS '17, CSL '15, ESA '13

Commissions of trust

2023 **PhD Thesis Opponent**, *K. Ariu*, KTH Stockholm
2023 **PhD Thesis Examiner**, *E. de Montbrun*, University of Toulouse
'20–'24 **Academic Consortium Partner** for NWO ENW-GROOT project Optimization for and with Machine Learning (OPTIMAL).
2020 **PhD Thesis Examiner**, *D. van der Hoeven*, Mathematical Institute, Leiden University
2019 **PhD Thesis Examiner**, *J. Witteveen*, ILLC, Universiteit van Amsterdam
2018 **MSc Committee aggregate member**, *6 students*, Mathematics, University of Milan
2018 **PhD Thesis Examiner**, *T. Sterkenburg*, Rijksuniversiteit Groningen
2017 **PhD Thesis Examiner**, *C. Vernade*, Télécom Paristech, Paris
2017 **PhD Thesis Examiner**, *P. Kamalaruban*, Australian National University, Canberra
2016 **MSc Thesis Examiner**, *D. van der Hoeven*, Leiden University
2012 **PhD Qualifying Committee**, *J. Nie*, UC Santa Cruz

Teaching

2024 **Lecturer**, *Graphical Models and Causality*, University of Twente
2024 **Lecturer**, *Machine Learning Theory*, MasterMath (149p)
2023 **Lecturer**, *Graphical Models and Causality*, University of Twente (10p)
2023 **Lecturer**, *Machine Learning Theory*, MasterMath (103p)
2022 **Lecturer**, *Machine Learning Theory*, MasterMath (136p)
2021 **Guest Lecturer**, Lecture on "Pure Exploration". *Reinforcement Learning* course, MSc, UvA
2021 **Lecturer**, *Machine Learning Theory*, MasterMath (145p)
2019 **Lecturer**, *Machine Learning Theory*, MasterMath (157p)
2018 **Lecturer**, *Machine Learning Theory*, MasterMath (127p)
2017 **Lecturer**, *Machine Learning Theory*, MasterMath (145p)
2016 **Lecturer**, ELC fall school on Theoretical Computer Science in Shonan, Japan (ca. 50p)
2016 **Lecturer**, *Game-theoretic Probability*, MSc Logic project, ILLC (6p)
2014 **Co-lecturer**, *Statistical Learning Theory* CS 281B / Stat 241B, graduate level, UC Berkeley (ca. 30p)
2012 **Advanced Topics** lectures *Computer architecture: The Simplest Processor* and *The Design of a Computer Opponent for Scrabble*, BSc, Royal Holloway, London (ca. 20p)
'07–'09 **Teaching Assistant**, *Kolmogorov Complexity*, MSc, University of Amsterdam (ca. 15p)
'05,'06 **Guest Lecturer**, "Procedure Calls" *Architectuur en Computerorganisatie*, BSc, UvA (ca. 25p)

Institutional Responsibilities

2024 **Reading Group Organiser**, *E-values*, CWI

2022 **Reading Group Organiser**, *Active Testing*, CWI
2018 **Reading Group Organiser**, *Pure Exploration*, CWI
2017 **Reading Group Organiser**, *Concentration Inequalities*, CWI
2016–now **Seminar Organiser**, *Machine Learning*, CWI

Outreach and Popularization

2020 **Long-form interview** on the Machine Learning Street Talk youtube channel.
2020 COLT **open problem** *Fast and optimal online portfolio selection*
2020 *Tidnabbil* **software** toolkit for structured multi-armed bandit problems
2018 Nieuw Archief voor Wiskunde **special issue** on Machine Learning contribution *Sub-Gaussians in Game-Theoretic Probability*
2016–now *Adversarial Intelligence* **blog**. 10595 views by at least 359 unique visitors in 2019 (source: AWStats)
2016 ERCIM News **special issue** on Machine Learning contribution *Robust and Adaptive Methods for Sequential Decision Making*
2016 *MetaGrad* **software** for online optimisation, MIT license
2015 COLT **open problem** *Online sabotaged shortest path*
2015 *Squint* **software** for adaptive sequential prediction, MIT license
2014 COLT **open problem** *Shifting experts on easy data*
2002–2010 *SIM-PL* digital circuit **simulator**, GPL license. In use 2002-now at University of Amsterdam

9. Scholarships, grants and prizes

Applicant

2015 NWO Veni 639.021.439, €247,443
2013 Queensland University of Technology Vice-Chancellor's Research Fellowship , €242,530
2011 NWO Rubicon 680-50-1010, €147,697

Co-applicant

2020 Inria International Program Associate Team 4TUNE. One of four co-PIs, €25,000
2018 Inria International Program Associate Team 6PAC. One of four co-PIs, €24,800

Honors

2020 Elected **Scholar** in the European Laboratory for Learning and Intelligent Systems (Ellis), member of *Interactive Learning and Interventional Representations* program.
2020 **Invited long-term participant** of the *Theory of Reinforcement Learning* program 2020 at the Simons Institute for the Theory of Computing, UC Berkeley
2017 **NeurIPS spotlight presentation** (3.5% of 3240 submissions) for [R16]
2017 Elected **promising young computer scientist** to the *Heidelberg Laureate Forum* 2017 (200 of 600+ applications).
2017 Elected **outstanding early career researcher** to the inaugural class 2017 of the *ACM Future of Computing Academy* (45 of 300+ applications)

2017 **Invited long-term participant** of the *Foundations of Machine Learning* program 2017 at the Simons Institute for the Theory of Computing, UC Berkeley
2016 **NeurIPS full oral presentation** (1.9% of 2500+ submissions) for [R19]
2012 **NeurIPS spotlight presentation** (5% of 1467 submissions) for [R37]

10. Output

10a. Output Indicators

In computer science, refereed publications at prestigious international conferences are considered at least as important as journal publications. They give more visibility within the community. *For high-impact conferences, the referee process is tough (often with a feedback and rebuttal phase), yet much faster than for journal submissions.*

10b. Publications

The top machine learning theory conference is COLT (Conference on Learning Theory), followed by Algorithmic Learning Theory (ALT). The two top machine learning conferences are NeurIPS (Neural Information Processing Systems) and ICML (International Conference on Machine Learning).

The top machine learning journals are the open access *Journal of Machine Learning Research* and Springer's *Machine Learning*. Related quality journals are *IEEE Transactions on Information Theory* and *Theoretical Computer Science*.

Of my 48 refereed publications, 32 articles are published at the top conferences COLT (11), ALT (4), NeurIPS (16) and ICML (1), and 8 articles are published in top journals JMLR (4), TCS (3) and IEEE Tr.IT (1).

11. Top Publications

1. W. M. Koolen, M. K. Warmuth and J. Kivinen. "Hedging Structured Concepts". In: *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*. ed. by A. T. Kalai and M. Mohri. June 2010, pp. 93–105 **(O)**.
   We show that existing methods for prediction in combinatorial domains suffer from the range factor problem. We develop new methods based on relative entropy regularisation that are optimal. This work forms the basis for our later generalisation of adaptive methods from the expert setting to general combinatorial domains [R26].

2. D. Adamskiy, W. M. Koolen, A. Chernov and V. Vovk. "A Closer Look at Adaptive Regret". In: *Proceedings of the 23rd International Conference on Algorithmic Learning Theory (ALT)*. ed. by N. Bshouty, G. Stoltz, N. Vayatis and T. Zeugmann. LNAI 7568. Springer, Oct. 2012, pp. 290–304.
   We show that the optimal algorithm for adaptive regret is in fact the existing Fixed Share algorithm; thus unifying two existing lines of work on tracking/shifting and on adaptive regret. We also show that the three distinct design methodologies of prior mixing, specialists and restarts all give rise to the same algorithm. This complete picture provides the basis for generalisations to a variety of

complex non-stationary prediction tasks in subsequent literature, including adaptive optimisation [X3], automata [o] and matrix learning [o].

3. W. M. Koolen and S. de Rooij. "Universal Codes from Switching Strategies". In: *IEEE Transactions on Information Theory* 59.11 (Nov. 2013), pp. 7168–7185.
We develop a graphical framework for the design of efficient prediction methods, and apply it to learning in non-stationary environments. Our foundation supported applications to complex hierarchical models including partition and context trees [o, o].

4. S. de Rooij, T. van Erven, P. Grünwald and W. M. Koolen. "Follow the Leader If You Can, Hedge If You Must". In: *Journal of Machine Learning Research* 15 (Apr. 2014), pp. 1281–1316 **(O)**.
We show how the benefits of the adaptive Follow The Leader strategy can be combined with worst-case safety at negligible overhead. This is one of the papers driving the Easy Data *movement, as it showed how adaptivity can be a consequence of data-dependent individual-sequence bounds. It stimulated my research into both stronger stochastic adaptivity proofs [R18] and more broadly adaptive algorithms [R26].*

5. W. M. Koolen and T. van Erven. "Second-order Quantile Methods for Experts and Combinatorial Games". In: *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*. ed. by P. Grünwald, E. Hazan and S. Kale. June 2015, pp. 1155–1175 **(O)**.
We develop advanced adaptive learning methods that are both worst-case safe and adaptive to easy data, by exploiting both low variances and large quantiles. See [S3] in 12e. for use of our open-source implementation. Our methods for the expert setting fuelled the subsequent significant generalisation to online convex optimisation [R19].

12. Output

Publications marked **O** appear in open-access venues.

12a. Refereed articles (48)

[R1] S. Agrawal, W. M. Koolen and S. Juneja. "Optimal Best-Arm Identification Methods for Tail-Risk Measures". In: *Advances in Neural Information Processing Systems (NeurIPS) 34*. Dec. 2021 **(O)**.

[R2] Y. Russac, C. Katsimerou, D. Bohle, O. Cappé, A. Garivier and W. M. Koolen. "A/B/n Testing with Control in the Presence of Subpopulations". In: *Advances in Neural Information Processing Systems (NeurIPS) 34*. Dec. 2021 **(O)**.

[R3] E. Kaufmann and W. M. Koolen. "Mixture Martingales Revisited with Applications to Sequential Tests and Confidence Intervals". In: *Journal of Machine Learning Research* 22.246 (Nov. 2021), pp. 1–44 **(O)**.

[R4] W. M. Koolen and P. Grünwald. "Log-optimal anytime-valid E-values". In: *International Journal of Approximate Reasoning* (Sept. 2021) **(O)**.

[R5] S. Agrawal, S. Juneja and W. M. Koolen. "Regret Minimization in Heavy-Tailed Bandits". In: *Proceedings of the 34th Annual Conference on Learning Theory (COLT)*. Aug. 2021, pp. 26–62 **(O)**.

[R6] S. Sachs, T. van Erven, W. M. Koolen and W. Kotłowski. "Robust Online Convex Optimization in the Presence of Outliers". In: *Proceedings of the 34th Annual Conference on Learning Theory (COLT)*. Aug. 2021, pp. 4174–4194 **(O)**.

[R7] A. Ramdas, J. Ruf, M. Larsson and W. M. Koolen. "Testing exchangeability: fork-convexity, supermartingales and e-processes". In: *International Journal of Approximate Reasoning* (July 2021) **(O)**.

[R8] T. van Erven, W. M. Koolen and D. van der Hoeven. "MetaGrad: Adaptation using Multiple Learning Rates in Online Learning". In: *Journal of Machine Learning Research* 22.161 (July 2021), pp. 1–61 **(O)**.

[R9] R. Degenne, H. Shao and W. M. Koolen. "Structure Adaptive Algorithms for Stochastic Bandits". In: *Proceedings of the 37th International Conference on Machine Learning (ICML)*. July 2020 **(O)**.

[R10] Z. Mhammedi and W. M. Koolen. "Lipschitz and Comparator-Norm Adaptivity in Online Learning". In: *Proceedings of the 33rd Annual Conference on Learning Theory (COLT)*. July 2020, pp. 2858–2887 **(O)**.

[R11] R. Degenne and W. M. Koolen. "Pure Exploration with Multiple Correct Answers". In: *Advances in Neural Information Processing Systems (NeurIPS) 32*. Dec. 2019, pp. 14591–14600 **(O)**.

[R12] R. Degenne, W. M. Koolen and P. Ménard. "Non-Asymptotic Pure Exploration by Solving Games". In: *Advances in Neural Information Processing Systems (NeurIPS) 32*. Dec. 2019, pp. 14492–14501 **(O)**.

[R13] T. van Ommen, W. M. Koolen and P. Grünwald. "Efficient Algorithms for Minimax Decisions under Tree-Structured Incompleteness". In: *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*. Sept. 2019, pp. 336–347 **(O)**.

[R14] Z. Mhammedi, W. M. Koolen and T. van Erven. "Lipschitz Adaptivity with Multiple Learning Rates in Online Learning". In: *Proceedings of the 32nd Annual Conference on Learning Theory (COLT)*. June 2019, pp. 2490–2511 **(O)**.

[R15] E. Kaufmann, W. M. Koolen and A. Garivier. "Sequential Test for the Lowest Mean: From Thompson to Murphy Sampling". In: *Advances in Neural Information Processing Systems (NeurIPS) 31*. Dec. 2018, pp. 6333–6343 **(O)**.

[R16] E. Kaufmann and W. M. Koolen. "Monte-Carlo Tree Search by Best Arm Identification". In: *Advances in Neural Information Processing Systems (NeurIPS) 30*. Dec. 2017, pp. 4904–4913 **(O)**.

[R17] W. Kotłowski, W. M. Koolen and A. Malek. "Random Permutation Online Isotonic Regression". In: *Advances in Neural Information Processing Systems (NeurIPS) 30*. Dec. 2017, pp. 4183–4192 **(O)**.

[R18] W. M. Koolen, P. Grünwald and T. van Erven. "Combining Adversarial Guarantees and Stochastic Fast Rates in Online Learning". In: *Advances in Neural Information Processing Systems (NeurIPS) 29*. Dec. 2016, pp. 4457–4465 **(O)**.

[R19] T. van Erven and W. M. Koolen. "MetaGrad: Multiple Learning Rates in Online Learning". In: *Advances in Neural Information Processing Systems (NeurIPS) 29*. Dec. 2016, pp. 3666–3674 **(O)**.

[R20] A. Garivier, E. Kaufmann and W. M. Koolen. "Maximin Action Identification: A New Bandit Framework for Games". In: *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*. June 2016, pp. 1028–1050 **(O)**.

[R21] W. Kotłowski, W. M. Koolen and A. Malek. "Online Isotonic Regression". In: *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*. June 2016, pp. 1165–1189 **(O)**.

[R22] T. van Ommen, W. M. Koolen, T. E. Feenstra and P. Grünwald. "Robust probability updating". In: *International Journal of Approximate Reasoning* 74 (Apr. 2016), pp. 30–57.

[R23] D. Adamskiy, W. M. Koolen, A. Chernov and V. Vovk. "A Closer Look at Adaptive Regret". In: *Journal of Machine Learning Research* 17.23 (Apr. 2016), pp. 1–21 **(O)**.

[R24] W. M. Koolen, A. Malek, P. L. Bartlett and Y. Abbasi-Yadkori. "Minimax Time Series Prediction". In: *Advances in Neural Information Processing Systems (NeurIPS) 28*. Dec. 2015, pp. 2548–2556 **(O)**.

[R25] P. L. Bartlett, W. M. Koolen, A. Malek, M. K. Warmuth and E. Takimoto. "Minimax Fixed-design Linear Regression". In: *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*. June 2015, pp. 226–239 **(O)**.

[R26] W. M. Koolen and T. van Erven. "Second-order Quantile Methods for Experts and Combinatorial Games". In: *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*. June 2015, pp. 1155–1175 **(O)**.

[R27] W. M. Koolen, A. Malek and P. L. Bartlett. "Efficient Minimax Strategies for Square Loss Games". In: *Advances in Neural Information Processing Systems (NeurIPS) 27*. Dec. 2014, pp. 3230–3238 **(O)**.

[R28] W. M. Koolen, T. van Erven and P. Grünwald. "Learning the Learning Rate for Prediction with Expert Advice". In: *Advances in Neural Information Processing Systems (NeurIPS) 27*. Dec. 2014, pp. 2294–2302 **(O)**.

[R29] W. M. Koolen and V. Vovk. "Buy low, sell high". In: *Theoretical Computer Science* 558.0 (Oct. 2014). The special issue on Algorithmic Learning Theory for ALT 2012, pp. 144–158.

[R30] S. de Rooij, T. van Erven, P. Grünwald and W. M. Koolen. "Follow the Leader If You Can, Hedge If You Must". In: *Journal of Machine Learning Research* 15 (Apr. 2014), pp. 1281–1316 **(O)**.

[R31] M. K. Warmuth, W. M. Koolen and D. P. Helmbold. "Combining Initial Segments of Lists". In: *Theoretical Computer Science* 519 (Jan. 2014). The special issue on Algorithmic Learning Theory for ALT 2011, pp. 29–45 **(O)**.

[R32] W. M. Koolen. "The Pareto Regret Frontier". In: *Advances in Neural Information Processing Systems (NeurIPS) 26*. Dec. 2013, pp. 863–871 **(O)**.

[R33] W. M. Koolen and S. de Rooij. "Universal Codes from Switching Strategies". In: *IEEE Transactions on Information Theory* 59.11 (Nov. 2013), pp. 7168–7185.

[R34] T. Scarfe, W. M. Koolen and Y. Kalnishkan. "A Long-Range Self-similarity Approach to Segmenting DJ Mixed Music Streams". In: *Artificial Intelligence Applications and Innovations*. Vol. 412. IFIP Advances in Information and Communication Technology. Sept. 2013, pp. 235–244 **(O)**.

[R35] W. M. Koolen, J. Nie and M. K. Warmuth. "Learning a set of directions". In: *Proceedings of the 26th Annual Conference on Learning Theory (COLT)*. June 2013 **(O)**.

[R36] W. M. Koolen and S. de Rooij. "Switching Investments". In: *Theoretical Computer Science* 473.0 (Feb. 2013). The special issue on Algorithmic Learning Theory for ALT 2010, pp. 61–76 **(O)**.

[R37] W. M. Koolen, D. Adamskiy and M. K. Warmuth. "Putting Bayes to Sleep". In: *Advances in Neural Information Processing Systems (NeurIPS) 25*. Dec. 2012, pp. 135–143 **(O)**.

[R38] D. Adamskiy, W. M. Koolen, A. Chernov and V. Vovk. "A Closer Look at Adaptive Regret". In: *Proceedings of the 23rd International Conference on Algorithmic Learning Theory (ALT)*. LNAI 7568. Oct. 2012, pp. 290–304.

[R39] W. M. Koolen and V. Vovk. "Buy low, sell high". In: *Proceedings of the 23rd International Conference on Algorithmic Learning Theory (ALT)*. LNAI 7568. Oct. 2012, pp. 335–349.

[R40] W. M. Koolen, W. Kotłowski and M. K. Warmuth. "Learning Eigenvectors for Free". In: *Advances in Neural Information Processing Systems (NeurIPS) 24*. Dec. 2011, pp. 945–953 **(O)**.

[R41] T. van Erven, S. de Rooij, W. M. Koolen and P. Grünwald. "Adaptive Hedge". In: *Advances in Neural Information Processing Systems (NeurIPS) 24*. Dec. 2011, pp. 1656–1664 **(O)**.

[R42] M. K. Warmuth, W. M. Koolen and D. P. Helmbold. "Combining Initial Segments of Lists". In: *Proceedings of the 22nd International Conference on Algorithmic Learning Theory (ALT)*. LNAI 6925. Oct. 2011, pp. 219–233.

[R43] H. Buhrman, P. T. S. van der Gulik, S. M. Kelk, W. M. Koolen and L. Stougie. "Some mathematical refinements concerning error minimization in the genetic code". In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)* 8 (Mar. 2011), pp. 1358–1372.

[R44] W. M. Koolen and S. de Rooij. "Switching Investments". In: *Proceedings of the 21st International Conference on Algorithmic Learning Theory (ALT)*. LNAI 6331. Oct. 2010, pp. 239–254.

[R45] W. M. Koolen, M. K. Warmuth and J. Kivinen. "Hedging Structured Concepts". In: *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*. June 2010, pp. 93–105 **(O)**.

[R46] E. G. Daylight, W. M. Koolen and P. M. B. Vitányi. "Time-bounded incompressibility of compressible strings and sequences". In: *Information Processing Letters (IPL)* 109.18 (Aug. 2009), pp. 1055–1059.

[R47] M. Ziegler and W. M. Koolen. "Kolmogorov Complexity Theory over the Reals". In: *Electronic Notes in Theoretical Computer Science (ENTCS)* 221 (Dec. 2008), pp. 153–169.

[R48] W. M. Koolen and S. de Rooij. "Combining Expert Advice Efficiently". In: *Proceedings of the 21st Annual Conference on Learning Theory (COLT)*. June 2008, pp. 275–286 **(O)**.

12b. Non-refereed articles (6)

The following publications serve to increase the visibility of theoretical machine learning in the Netherlands.

[N1] W. M. Koolen. "Sub-Gaussians in game-theoretic probability". In: *Nieuw Archief voor Wiskunde* 5/19.3 (Sept. 2018).

[N2] W. M. Koolen. "Robust and Adaptive Methods for Sequential Decision Making". In: *ERCIM News, special issue on Machine Learning* 2016.107 (Oct. 2016) **(O)**.

[N3] W. M. Koolen and T. van Erven. "Second-order Quantile Methods for Online Sequential Prediction". In: *Proceedings of the 24th Annual Belgian-Dutch Conference on Machine Learning (BeneLearn)*. June 2015.

[N4] W. M. Koolen and T. van Erven. "Freezing and Sleeping: Tracking Experts that Learn by Evolving Past Posteriors". In: *Proceedings of the 18th Annual Belgian-Dutch Conference on Machine Learning (BeneLearn)*. May 2009, pp. 91–92 **(O)**.

[N5] W. M. Koolen and S. de Rooij. "Combining Expert Advice Efficiently". In: *Proceedings of the 20th Belgian-Dutch Conference on Artificial Intelligence (BNAIC)*. Oct. 2008, pp. 323–324.

[N6] B. Bruidegom and W. M. Koolen. "SIM-PL, auteursomgeving voor digitale componenten". In: *Nationaal Informatica Onderwijs Congres (NIOC)*. Apr. 2004, pp. 14–19 **(O)**.

12c. Working Papers (5)

The following papers are in preparation or under review

[P2] J. Ruf, M. Larsson, W. M. Koolen and A. Ramdas. "A composite generalization of Ville's martingale theorem". In: *Electronic Journal of Probability* 28 (Oct. 2023), pp. 1–21.

[P1] M. Pérez-Ortiz and W. M. Koolen. "Luckiness in Multiscale Online Learning". In: *Advances in Neural Information Processing Systems (NeurIPS) 35*. Dec. 2022.

[P3] Z. Mhammedi and W. M. Koolen. "Near-Optimal Concentration Inequalities for CVaR with Unbounded Random Variables". To be submited to Journal of Machine Learning Research. June 2021.

[P4] A. Ramdas, J. Ruf, M. Larsson and W. M. Koolen. "Admissible anytime-valid sequential inference must rely on nonnegative martingales". In: *ArXiv* (Sept. 2020) **(O)**.

[P5] P. Grünwald, R. de Heide and W. M. Koolen. "Safe Testing". In: *ArXiv* (June 2019) **(O)**.

12d. Open problems (3)

The COLT conference has a formal peer reviewed open problem section. Problems accepted here provide direction to the community and typically spur a range of follow-up work.

[O1] T. van Erven, D. van der Hoeven, W. Kotłowski and W. M. Koolen. "Open Problem: Fast and Optimal Online Portfolio Selection". In: *Proceedings of the 33rd Annual Conference on Learning Theory (COLT)*. July 2020 **(O)**.

[O2] W. M. Koolen, M. K. Warmuth and D. Adamskiy. "Open Problem: Online Sabotaged Shortest Path". In: *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*. June 2015, pp. 1764–1766 **(O)**.

[O3] M. K. Warmuth and W. M. Koolen. "Open Problem: Shifting Experts on Easy Data". In: *Proceedings of the 27th Annual Conference on Learning Theory (COLT)*. June 2014, pp. 1295–1298 **(O)**.

Kale, Lee and Pal [X1] show that [O2] is computationally hard. Sani, Neu and Lazaric [X4] and Luo and Schapire [X3] each solve differently [O3, first half]. [O3, last half] remains open.

12e. Software (4)

A crucial aspect of the development of machine learning methods is to implement the algorithms and make them publicly available.

[S1] Tidnabbil (`https://bitbucket.org/wmkoolen/tidnabbil`, Julia, 64 files, 10K lines, public domain, companion to [R12, R9]) is a library of learning algorithms for structured bandit problems based on iterative saddle point solvers.

[S2] MetaGrad (`https://bitbucket.org/wmkoolen/metagrad`, Matlab/Octave, 7 files, 298 lines, MIT license, companion to [R19]) is an adaptive first-order method for online convex optimisation. *By learning the learning rate, MetaGrad significantly outperforms worst-case regret guarantees in many practical examples. It adapts to the structure of the loss functions even without any curvature.*

[S3] Squint (`https://bitbucket.org/wmkoolen/squint`, Matlab/Octave/Python/C++, 14 files, 760 lines, MIT license, companion to [R26]) is an advanced method for prediction with expert advice. *By learning the learning rate, Squint provides superior regret guarantees in the form of second-order and quantile bounds.*

[S4] SIM-PL (`https://bitbucket.org/wmkoolen/sim-pl`, Java, 1867 files, 278771 lines, GPL license) is a discrete event simulator for teaching the design of digital circuits, ranging from simple gates to pipelined computer processors. *SIM-PL has been in use continually from 2002 in the UvA BSc course Architectuur en Computerorganisatie.*

Orabona and Pal [X2] use [S3] in a comparative experimental evaluation.

13. References

[X1] S. Kale, C. Lee and D. Pal. "Hardness of Online Sleeping Combinatorial Optimization Problems". In: *Advances in Neural Information Processing Systems 29*. 2016, pp. 2181–2189.

[X2] F. Orabona and D. Pal. "Coin Betting and Parameter-Free Online Learning". In: *Advances in Neural Information Processing Systems 29*. 2016, pp. 577–585.

[X3] H. Luo and R. E. Schapire. "Achieving All with No Parameters: AdaNormalHedge". In: *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*. Vol. 40. JMLR Workshop and Conference Proceedings. 2015, pp. 1286–1304.

[X4] A. Sani, G. Neu and A. Lazaric. "Exploiting easy data in online optimization". In: *Advances in Neural Information Processing Systems 27*. 2014, pp. 810–818.