Curriculum Vitae

September 2018

Univ.-Prof. Dr. Peter Auer
Montanuniversität Leoben
Chair for Information Technology
Franz-Josef-Strasse 18, A-8700 Leoben, Austria

Phone: +43 (3842) 402-1500
Email: auer@unileoben.ac.at

Date of Birth: 21 January 1964
Place of Birth: Vienna
Nationality: Austria

Education

1982 – 1987: Undergraduate and Graduate Studies at the Vienna University of Technology, receiving an MSc in mathematics ("Technische Mathematik").

1988 – 1992: PhD studies at the Vienna University of Technology, receiving a PhD in mathematics.

Career History

1988 – 1989 Research assistant at the Institute for Statistics and Probability Theory, Vienna University of Technology, Prof. Pal Révész.

Jan 1990 – Mar 1992 Research assistant ("Universitätsassistent") at the Institute for Applied Computer Science, Vienna University of Technology, Prof. Wilhelm Barth.

Apr 1992 – Jun 1997: Research assistant ("Universitätsassistent") at the Institute for Theoretical Computer Science, Graz University of Technology, Prof. Wolfgang Maass.

Oct 1993 — Jul 1994: Civil service at Lebenshilfe Steiermark.

Mar 1995 – Feb 1996: Research Scholar at the University of California, Santa Cruz, supported by an Erwin-Schrödinger-scholarship (FWF, Austrian Science Foundation).

Jun 1997: Habilitation ("Theoretical Computer Science and Probability Theory"), promotion to Associate Professorship ("Universitätsdozent").

Jun 1997 – Jan 2003: Associate Professor at the Institute for Theoretical Computer Science, Graz University of Technology.

Feb 2003 – now: Full professor for Information Technology at the Montanuniversität Leoben.

University teaching

Since 1988 I have been teaching a large number of courses in the (somewhat overlapping) areas of Theoretical Computer Science, Logic, Algorithms and Data Structures, Machine Learning, Neural Networks, Optimization, Introduction to Programming and Software Engineering. The following are my current courses at the Montanuniversität Leoben.

Algorithms and Datastructures:
Analysis and design of algorithms, sorting and searching, graph algorithms, combinatorial optimization, integer programming, heuristics.

Machine Learning:
Classification and regression, clustering, reinforcement learning, online learning, support vector machines, neural networks, deep learning, Bayesian learning, error estimation, model selection and overfitting.

Lab in Software Engineering:
Requirements specification, object oriented analysis and software design, implementation, testing. Strong emphasis is on the requirements specification. Software design and implementation is optional for extra credit.

Computer Application and Programming and **IT II:**
Control structures, types and variables, object oriented programming, classes and objects, lists, assertions and loop invariants, recursion. Programming language: Java.

Didactics training:
Didaktik 1-3, TU Graz, Montanuniversität Leoben, 2010–2015 (48 hours)
Didaktik-Werkstatt der Steirischen Hochschulkonferenz, 2017–2018 (20 hours)

Supervision of postdoctoral researchers and PhD candidates

- Christian Savu-Krohn (2004-2009), now Senior Data Scientist;
- Christos Dimitrakakis (2009), now Docent at Chalmers University of Technology;
- Po Leung (2008–2011), now Assistant Professor at Macau University of Science and Technology;
- Shiau Hong Lim (2009–2012), now researcher at IBM Singapore;
- Odalric-Ambrym Maillard (2011–2013), now researcher at INRIA Lille;
- Chao-Kai Chiang (2014–2016), now postdoctoral researcher at the University of Southern California;
- Martin Antenreiter (2013-2017), now Senior Scientist at the Montanuniversität Leoben;

- Pratik Gajane (2018–ongoing);
- Ahmad Almasri (2018-ongoing);
- Anton Lettner (2018-ongoing);
- Elias Hagendorfer (2018-ongoing).

Project acquisition and management

DELTA Dynamically Evolving Long-Term Autonomy (2018–2021)
Funded by FWF and Horizon 2020 (CHIST-ERA), responsible for 210k€ out of 670k€. Partners: Universitat Pompeu Fabra (coordinator, Spain), INRIA Lille (France), Université de Liège (Belgium).

aiMotionLab Artificial Intelligence in Motion Laboratory (2018–2019)
Funded by Zukunftsfonds Steiermark, responsible for 40k€ out of 530k€. Partners: FH Joanneum (coordinator), TU Graz.

Green Big Data (2018–2021)
Funded by FFG (Innovationslehrgang), responsible for 15k€ out of 540k€. Coordinator: FH Joanneum, 17 partners.

ReWaste4.0 Recycling and Recovery of Waste 4.0 (2017–2021)
Funded by FFG (K-Project), responsible for 70k€ out of 1300k€. Coordinator: Montanuniversität Leoben, 8 partners.

CompLACS Composing Learning for Artificial Cognitive Systems (2011–2015)
Funded by EU-FP7, responsible for 500k€ out of 4400k€. Partners: University College London (coordinator, UK), University of Bristol (UK), Royal Holloway University London (UK), Radboud Universiteit Nijmegen (Netherlands), TU Berlin (Germany), INRIA Lille (France).

PinView Personal Information Navigator Adapting Through Viewing (2008–2011)
Funded by EU-FP7, responsible for 360k€ out of 2550k€. Partners: Helsinki University of Technology (coordinator), University College London (UK), University of Southampton (UK), XRCE (France), cellum gmbh (Austria).

Cognitive Vision (2003–2009)
National Research Network funded by the FWF, responsible for 390k€ out of 3610k€. Partners: ACIN (coordinator, TU Wien), PRIP (TU Wien), IGI (TU Graz), EMT (TU Graz), ICG (TU Graz), Joanneum Research.

PASCAL Pattern Analysis, Statistical Modelling and Computational Learning
(2003–2008) Network of Excellence funded by EU-FP6, responsible for 180k€ out of 5440k€. More than 50 partners.

LAVA Learning for Adaptable Visual Assistants (2002–2005)
Funded by EU-FP5, responsible for 260k€ funding out of a total 1620k€. Partners: Xerox Research Center Europe (France, coordinator), INRIA Grenoble (France), Royal Holloway University London (UK), Lund University (Sweden), IDIAP (Switzerland), CNRS (France), Institut National Polytechnique de Grenoble (France).

Professional Duties

President of the Association for Computational Learning (ACL), 2012-2018.
The ACL organizes the yearly COLT conference on computational learning theory.

Action Editor for the *Journal of Machine Learning Research*.

Editorial board member of the *Machine Learning Journal*.

Position statement

I have been working in several areas of machine learning, mostly on the analysis of machine learning problems, the design of corresponding algorithms, and their theoretical justification. The initial focus of my work was on the founding principles of machine learning: How many training examples are necessary (and sufficient) to guarantee a certain performance of the learned predictor, and which learning algorithms are capable of efficiently utilizing these training examples? Soon I have become interested in on-line learning, where training examples arrive one at a time. Of particular interest I find learning scenarios where the learning algorithm can influence the information it receives. This often leads to an exploitation-exploration dilemma, for example in multi-armed bandit problems and reinforcement learning. Another line of my research has been the application of machine learning in computer vision.

The focus of my current research is on learning agents that efficiently explore their environment and learn (nearly) optimal strategies for their environment. A natural model for such learning is reinforcement learning, but current theoretical analysis of on-line reinforcement learning is insufficient for reasonably large environments. The main question is about which properties of an environment allow for its efficient exploration and for finding good strategies.

In a more general perspective, while the mathematical underpinning of machine learning is relatively well established, it is often too pessimistic as it provides mostly worst-case performance bounds. I strongly believe that in order to advance machine learning (its application *and* its theory), we need to work on concrete practical learning problems to understand the properties of real-world data and how their regularities can be utilized. I also believe that further investigating interactive learning modes — where the learning algorithm has some control about the data it will receive — is crucial in addressing the bottleneck of supervised and labeled data.

My primary goals in teaching are to instill curiosity about how things work, what is possible, and what are good or optimal solutions to given problems. This includes intuition

about the mechanisms and objects studied, formal and empirical tools to verify this intuition, and analytical and algorithmic problem solving skills. A message I want to convey is that the right view on a problem, the right question, the right abstraction, or just the right notation, often makes a problem quite transparent and much easier to solve. Thus the main task in problem solving is to find the right angle. Here experience and a good toolset are a great help. Therefore, I try to introduce appropriate tools but also to allow for enough time for experimenting with this tools.

List of Publications

Peter Auer September 2018

Articles in Peer-Reviewed Journals

[j29] Monireh Dabaghchian, Amir Alipour-Fanid, Kai Zeng, Qingsi Wang, and Peter Auer. Online learning with randomized feedback graphs for optimal PUE attacks in cognitive radio networks. *IEEE/ACM Trans. Networking*, 2018.

[j28] R. Ortner, D. Ryabko, P. Auer, and R. Munos. Regret bounds for restless markov bandits. *Theor. Comput. Sci.*, 558:62–76, 2014.

[j27] Y. Seldin, F. Laviolette, N. Cesa-Bianchi, J. Shawe-Taylor, and P. Auer. PAC-Bayesian inequalities for martingales. *IEEE Trans. Information Theory*, 58(12):7086–7093, 2012.

[j26] C. Savu-Krohn, G. Rantitsch, P. Auer, F. Melcher, and T. Graupner. Geochemical Fingerprinting of Coltan Ores by Machine Learning on Uneven Datasets. *Natural resources research*, 20(3):177–191, 2011.

[j25] T. Jaksch, R. Ortner, and P. Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11:1563–1600, 2010.

[j24] P. Auer and R. Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.

[j23] P. Auer, H. Burgsteiner, and W. Maass. A learning rule for very simple universal approximators consisting of a single layer of perceptrons. *Neural Networks*, 21(5):786–795, 2008.

[j22] P. Auer and R. Ortner. A new PAC bound for intersection-closed concept classes. *Machine Learning*, 66(2-3):151–163, 2007.

[j21] A. Opelt, A. Pinz, M. Fussenegger, and P. Auer. Generic object recognition with boosting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(3):416–431, 2006.

[j20] P. Auer and N. Cesa-Bianchi. A distributed voting scheme to maximize preferences. *RAIRO - Theoretical Informatics and Applications*, 40(2):389 – 403, 2006.

[j19] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397 – 422, 2002.

[j18] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235 – 256, 2002.

[j17] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multi-armed bandit problem. *SIAM J. Comput.*, 32(1):48–77, 2002.

[j16] P. Auer, N. Cesa-Bianchi, and C. Gentile. Adaptive and self-confident on-line learning algorithms. *J. Comput. Syst. Sci.*, 64(1):48–75, 2002.

[j15] K. Andrews, W. Kienreich, V. Sabol, J. Becker, G. Droschl, F. Kappe, M. Granitzer, P. Auer. and K. Tochtermann. The infosky viusal explorer: exploiting hierarchical structure and documents similarities. *Information Visualization*, 1(3-4):166–181, 2002.

[j14] P. Auer and P. M. Long. Structural results about on-line learning models with and without queries. *Machine Learning*, 36(3):147–181, 1999.

[j13] P. Auer, P. M. Long, and A. Srinivasan. Approximating hyper-rectangles: Learning and pseudorandom sets. *J. Comput. Syst. Sci.*, 57(3):376–388, 1998.

[j12] P. Auer. On learning from ambiguous information. *Periodica Polytechnica Electrical Engineering*, 42(1):115–122, 1998.

[j11] P. Auer and N. Cesa-Bianchi. On-line learning with malicious noise and the closure algorithm. *Annals of Mathematics and Artificial Intelligence*, 23(1-2):83–99, 1998.

[j10] P. Auer and M. K. Warmuth. Tracking the best disjunction. *Machine Learning*, 32(2):127–150, 1998.

[j9] J. Kivinen, M. K. Warmuth, and P. Auer. The perceptron algorithm vs. Winnow: linear vs. logarithmic mistake bounds when few input variables are relevant. *Artificial Intelligence*, 97(1-2):325–343, 1997.

[j8] P. Auer. Learning nested differences in the presence of malicious noise. *Theoretical Computer Science*, 185(1):159–175, 1997.

[j7] P. Auer and K. Hornik. The number of points of an empirical or Poisson process covered by unions of sets. *Journal of Multivariate Analysis*, 57(1):37–51, 1996.

[j6] P. Auer and K. Hornik. Limit laws for the maximal and minimal increments of the Poisson process. *Studia Scientiarum Mathematicarum Hungarica*, 31(1-3):1–13, 1996.

[j5] P. Auer, P. M. Long, W. Maass, and G. J. Wöginger. On the complexity of function learning. *Machine Learning*, 18(2-3):187–230, 1995.

[j4] K. Hornik, M. Stinchcombe, H. White, and P. Auer. Degree of approximation results for feedforward networks approximating unknown mappings and their derivatives. *Neural Computation*, 6(6):1262–1275, 1994.

[j3] P. Auer and K. Hornik. On the number of points of a homogeneous Poisson process. *Journal of Multivariate Analysis*, 48(1):115–156, 1994.

[j2] P. Auer, K. Hornik, and P. Révész. Some limit theorems for the homogeneous Poisson process. *Statistics & Probability Letters*, 12(2):91–96, 1991.

[j1] P. Auer. The circle homogeneously covered by random walk on Z^2. *Statistics & Probability Letters*, 9(5):403–407, 1990.

Edited Volumes

[e5] P. Auer, A. Clark, and T. Zeugmann. Special Issue on Algorithmic Learning Theory. *Theor. Comput. Sci.*, 650, 2016.

[e4] P. Auer, A. Clark, T. Zeugmann, and S. Zilles, editors. *Algorithmic Learning Theory - 25th International Conference, ALT 2014*, volume 8776 of *Lecture Notes in Computer Science*. Springer, 2014.

[e3] P. Auer and P. M. Long, editors. Special Issue on Learning Theory 2005. *J. Comput. Syst. Sci.*, 74(8), 2008.

[e2] P. Auer and R. Meir, editors. *Learning Theory, 18th Annual Conference on Learning Theory, COLT 2005*, volume 3559 of *Lecture Notes in Computer Science*. Springer, 2005.

[e1] P. Auer and W. Maass, editors. Special Issue on Computational Learning Theory. *Algorithmica*, 22(1/2), 1998.

Peer-Reviewed Conference and Workshop Articles

[c45] Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best arm with an unknown number of distribution changes. Accepted for *EWRL 2018, 14th European Workshop on Reinforcement Learning*, 2018.

[c44] Pratik Gajane, Ronald Ortner, and Peter Auer. A sliding-window algorithm for Markov decision processes with arbitrarily changing rewards and transitions. In *LLARLA 2018, ICML workshop on Lifelong Learning: A Reinforcement Learning Approach*, 2018.

[c43] P. Auer, C.-K. Chiang, R. Ortner, and M. M. Drugan. Pareto front identification from stochastic bandit feedback. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016*, volume 51 of *JMLR Workshop and Conference Proceedings*, pages 939–947. JMLR.org, 2016.

[c42] P. Auer and C.-K. Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016*, volume 49 of *JMLR Workshop and Conference Proceedings*, pages 116–120. JMLR.org, 2016.

[c41] R. Ortner, D. Ryabko, P. Auer, and R. Munos. Regret bounds for restless Markov bandits. In *Algorithmic Learning Theory - 23rd International Conference, ALT 2012*, volume 7568, pages 214–228. Springer, 2012.

[c40] Y. Seldin, C. Szepesvári, P. Auer, and Y. Abbasi-Yadkori. Evaluation and analysis of the performance of the EXP3 algorithm in stochastic environments. In *Proceedings of the Tenth European Workshop on Reinforcement Learning, EWRL 2012*, volume 24 of *JMLR Proceedings*, pages 103–116. JMLR.org, 2012.

[c39] S. Kalyanakrishnan, A. Tewari, P. Auer, and P. Stone. PAC subset selection in stochastic multi-armed bandits. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, pages 655–662.

[c38] S. H. Lim and P. Auer. Autonomous exploration for navigating in MDPs. In *Proceedings of the 25th Annual Conference on Learning Theory, COLT 2012*, volume 23 of *JMLR Proceedings*, pages 40.1–40.24. JMLR.org, 2012.

[c37] A. Carpentier, A. Lazaric, M. Ghavamzadeh, R. Munos, and P. Auer. Upper-confidence-bound algorithms for active learning in multi-armed bandits. In *Algorithmic Learning Theory - 22nd International Conference, ALT 2011*, volume 6925 of *Lecture Notes in Computer Science*, pages 189–203. Springer, 2011.

[c36] Y. Seldin, P. Auer, F. Laviolette, J. Shawe-Taylor, and R. Ortner. PAC-Bayesian analysis of contextual bandits. In *Advances in Neural Information Processing Systems 24, NIPS 2011*, pages 1683–1691.

[c35] S. H. Lim and P. Auer. Noisy search with comparative feedback. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2011*, pages 445–452. AUAI Press, 2011.

[c34] Z. Hussain, A. P. Leung, K. Pasupa, D. R. Hardoon, P. Auer, and J. Shawe-Taylor. Exploration-exploitation of eye movement enriched multiple feature spaces for content-based image retrieval. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2010*, volume 6321 of *Lecture Notes in Computer Science*, pages 554–569. Springer, 2010.

[c33] J. Prankl, M. Antenreiter, P. Auer, and M. Vincze. Consistent interpretation of image sequences to improve object models on the fly. In *7th International Conference on Computer Vision Systems, ICVS 2009*, volume 5815 of *Lecture Notes in Computer Science*, pages 384–393. Springer, 2009.

[c32] M. Antenreiter, J. Prankl, M. Vincze, and P. Auer. Using a spatio-temporal reasoning system to improve object models on the fly. In *33rd Workshop of the Austrian Association for Pattern Recognition - Visual Learning, ÖAGM 2009*, pages 25–36.

[c31] M. Antenreiter, R. Ortner, and P. Auer. Combining classifiers for improved multilabel image classification. In *Learning from Multi-label Data Workshop at ECML 2009*, pages 16–27.

[c30] P. Auer, T. Jaksch, and R. Ortner. Near-optimal regret bounds for reinforcement learning. In *Advances in Neural Information Processing Systems 21, NIPS 2008*, pages 89–96.

[c29] A. P. Leung and P. Auer. An efficient search algorithm for content-based image retrieval with user feedback. In *Workshops Proceedings of the 8th IEEE International Conference on Data Mining, ICDM 2008*, pages 884–890. IEEE Computer Society, 2008.

[c28] P. Auer, R. Ortner, and C. Szepesvri. Improved rates for the stochastic continuum-armed bandit problem. In *20th Annual Conference on Learning Theory, COLT 2007*, volume 4539 of *Lecture Notes in Artificial Intelligence*, pages 454–468. Springer, 2007.

[c27] P. Auer and R. Ortner. Logarithmic online regret bounds for undiscounted reinforcement learning. In *Advances in Neural Information Processing Systems 19, NIPS 2006*, pages 49–56.

[c26] M. Antenreiter, Ch. Savu-Krohn, and P. Auer. Visual classification of images by learning geometric appearances through boosting. In *IAPR Workshop, ANNPR 2006*, volume 4087 of *Lecture Notes in Computer Science*, pages 233–243. Springer, 2006.

[c25] C. Allenberg, P. Auer, L. Gyrfi, and G. Ottucsk. Hannan consistency in online learning in case of unbounded losses under partial monitoring. In *Algorithmic Learning Theory, ALT 2006*, volume 4246 of *Lecture Notes in Computer Science*, pages 229–243. Springer, 2006.

[c24] Ch. Savu-Krohn and P. Auer. A simple feature extraction for high dimensional image representations. In *Subspace, Latent Structure and Feature Selection*, volume 3940 of *Lecture Notes in Computer Science*, pages 163–172. Springer, 2006.

[c23] P. Auer and R. Ortner. A boosting approach to multiple instance learning. In *15th European Conference on Machine Learning, ECML 2004*, volume 3201 of *Lecture Notes in Artificial Intelligence*, pages 63–74. Springer, 2004.

[c22] P. Auer and R. Ortner. A new PAC-bound for intersection-closed concept classes. In *17th Annual Conference on Learning Theory, COLT 2004*, volume 3120 of *Lecture Notes in Computer Science*, pages 408–414. Springer, 2004.

[c21] M. Fussenegger, A. Opelt, A. Pinz, and P. Auer. Object recognition using segmentation for feature detection. In *28th Workshop of the Austrian Association for Pattern Recognition - Digital Imaging in Media and Education , ÖAGM 2004*, pages 103–110.

[c20] M. Fussenegger, A. Opelt, A. Pinz, and P. Auer. Object recognition using segmentation for feature detection. In *17th International Conference on Pattern Recognition, ICPR 2004*, pages 41–44.

[c19] A. Opelt, M. Fussenegger, A. Pinz, and P. Auer. Weak hypotheses and boosting for generic object detection and recognition. In *8th European Conference on Computer Vision, ECCV 2004*, volume 4022 of *Lecture Notes in Computer Science*, pages 71 – 84. Springer, 2004.

[c18] P. Auer, H. Burgsteiner, and Wolfgang Maass. Reducing communication for distributed learning in neural networks. In *International Conference on Artificial Neural Networks, ICANN 2002*, volume 2415 of *Lecture Notes in Computer Science*, pages 123–128. Springer, 2002.

[c17] P. Auer. Using upper confidence bounds for online learning. In *41th Annual Symposium on Foundations of Computer Science, FOCS 2000*, pages 270–293. IEEE Computer Society, 2000.

[c16] P. Auer. An improved on-line algorithm for learning linear evaluation functions. In *13th Ann. Conf. Computational Learning Theory, COLT 2000*, pages 118–125. Morgan Kaufmann, 2000.

[c15] P. Auer and C. Gentile. Adaptive and self-confident on-line learning algorithms. In *13th Ann. Conf. Computational Learning Theory, COLT 2000*, pages 107–117. Morgan Kaufmann, 2000.

[c14] P. Auer, P. M. Long, and A. Srinivasan. Approximating hyper-rectangles: Learning and pseudo-random sets. In 29th Ann. Symp. Theory of Computing, STOC 1997, pages 314–323. ACM, 1997.

[c13] P. Auer. On learning from multi-instance examples: Empirical evaluation of a theoretical approach. In *Proc. 14th Int. Conf. Machine Learning, ICML 1997*, pages 21–29. Morgan Kaufmann, 1997.

[c12] P. Auer, S. Kwek, W. Maass, and M. Warmuth. Learning of depth two neural nets with constant fan-in at the hidden nodes. In *9th Annual ACM Conference on Computational Learning Theory, COLT 1996*, pages 333–343. ACM, 1996.

[c11] P. Auer, M. Herbster, and M. K. Warmuth. Exponentially many local minima for single neurons. In *Advances in Neural Information Processing Systems 8, NIPS 1995*, pages 316–322. MIT Press, 1996.

[c10] P. Auer, R. C. Holte, and W. Maass. Theory and applications of agnostic PAC-learning with small decision trees. In *12th International Machine Learning Conference, ICML 1995*, pages 21–29. Morgan Kaufmann, 1995.

[c9] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *36th Annual Symposium on Foundations of Computer Science, FOCS 1995*, pages 322–331. IEEE Computer Society, 1995.

[c8] P. Auer and M. K. Warmuth. Tracking the best disjunction. In *36th Annual Symposium on Foundations of Computer Science, FOCS 1995*, pages 312–321. IEEE Computer Society, 1995.

[c7] P. Auer. Learning nested differences in the presence of malicious noise. In *6th International Conference on Algorithmic Learning Theory, ALT 1995*, volume 997 of *Lecture Notes in Computer Science*, pages 123–137. Springer, 1995.

[c6] P. Auer and P. M. Long. Simulating access to hidden information while learning. In *26th Annual ACM Symposium on the Theory of Computing, STOC 1994*, pages 263–272. ACM, 1994.

[c5] P. Auer and N. Cesa-Bianchi. On-line learning with malicious noise and the closure algorithm. In *5th International Workshop on Analogical and Inductive Inference, ALT'94*, volume 872 of *Lecture Notes in Computer Science*, pages 229–247. Springer, 1994.

[c4] P. Auer. On-line learning of rectangles in noisy environments. In *6th Annual ACM Conference on Computational Learning Theory, COLT 1993*, pages 253–261. ACM, 1993.

[c3] P. Auer, P. M. Long, W. Maass, and G. J. Wöginger. On the complexity of function learning. In *6th Annual ACM Conference on Computational Learning Theory, COLT 1993*, pages 392–401. ACM, 1993.

[c2] P. Auer and P. Révész. On the relative frequency of points visited by random walk on $\mathbf{Z}^2$. In *Limit Theorems in Probability and Statistics, Pecs 1990*, volume 57 of *Colloquia Mathematica Societatis Janos Bolyai*, pages 27–33. North-Holland, 1989.

[c1] P. Auer. Some hitting probabilities of random walks on Z^2. In *Limit Theorems in Probability and Statistics, Pecs 1990*, volume 57 of *Colloquia Mathematica Societatis Janos Bolyai*, pages 9–25. North-Holland, 1989.

Book Chapters

[b3] P. Auer. Online learning. In *Encyclopedia of Machine Learning and Data Mining*, pages 929–937. Springer, 2017.

[b2] P. Auer. Learning with malicious noise. In *Encyclopedia of Algorithms*, pages 1086–1089. Springer, 2016.

[b1] P. Auer and A. P. Leung. Relevance feedback models for content-based image retrieval. In L. Weisi et al., editors, *Multimedia Analysis, Processing and Communications*, pages 59–79. Springer, 2011.

Other Conference and Workshop Articles

[o2] P. Auer. Solving string equations with constant restrictions. In *Word Equations and Related Topics*, pages 103–132. Lecture Notes of Computer Science 677, Springer, 1993.

[o1] P. Auer. Unification in the combination of disjoint theories. In *Word Equations and Related Topics*, pages 177–186. Lecture Notes of Computer Science 677, Springer, 1993.