



Universiteit
Leiden
The Netherlands

The many faces of online learning

Hoeven, D. van der

Citation

Hoeven, D. van der. (2021, March 4). *The many faces of online learning*. Retrieved from <https://hdl.handle.net/1887/3147345>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3147345>

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <https://hdl.handle.net/1887/3147345> holds various files of this Leiden University dissertation.

Author: Hoeven, D. van der

Title: The many faces of online learning

Issue Date: 2021-03-4

Bibliography

- Abernethy, J., Hazan, E., and Rakhlin, A. (2008). Competing in the dark: An efficient algorithm for bandit linear optimization. In *Proceedings of the 21st Annual Conference on Learning Theory (COLT)*, pages 263–274.
- Abernethy, J., Hazan, E., and Rakhlin, A. (2012). Interior-point methods for full-information and bandit online learning. *IEEE Trans. Information Theory*, 58(7):4164–4175.
- Abernethy, J. and Rakhlin, A. (2009). An efficient bandit algorithm for \sqrt{T} -regret in online multiclass prediction? In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*.
- Abernethy, J. D., Jung, Y. H., Lee, C., McMillan, A., and Tewari, A. (2019). Online learning via the differential privacy lens. In *Advances in Neural Information Processing Systems 32*, pages 8894–8904.
- Agarwal, A., Dekel, O., and Xiao, L. (2010). Optimal algorithms for online convex optimization with multi-point bandit feedback. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, pages 28–40. Citeseer.
- Agarwal, A., Foster, D. P., Hsu, D. J., Kakade, S. M., and Rakhlin, A. (2011). Stochastic convex optimization with bandit feedback. In *Advances in Neural Information Processing Systems 24*, pages 1035–1043.
- Agarwal, A. and Hazan, E. (2005). Efficient algorithms for online game playing and universal portfolio management. *ECCC, TR06-033*.
- Agarwal, A., Hazan, E., Kale, S., and Schapire, R. E. (2006). Algorithms for portfolio management based on the newton method. In *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, pages 9–16.
- Agarwal, N., Bullins, B., Chen, X., Hazan, E., Singh, K., Zhang, C., and Zhang, Y. (2018). The case for full-matrix adaptive regularization. *arXiv preprint arXiv:1806.02958*.

- Agarwal, N. and Singh, K. (2017). The price of differential privacy for online learning. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 32–40.
- Arora, S., Hazan, E., and Kale, S. (2012). The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The non-stochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77.
- Azoury, K. S. and Warmuth, M. K. (2001). Relative loss bounds for on-line density estimation with the exponential family of distributions. *Machine Learning*, 43(3):211–246.
- Banerjee, A., Merugu, S., Dhillon, I. S., and Ghosh, J. (2005). Clustering with bregman divergences. *The Journal of Machine Learning Research*, 6:1705–1749.
- Bartlett, P. L., Hazan, E., and Rakhlin, A. (2007). Adaptive online gradient descent. In *Advances in Neural Information Processing Systems 20*, pages 65–72.
- Bartlett, P. L., Jordan, M. I., and McAuliffe, J. D. (2006). Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156.
- Bartlett, P. L. and Mendelson, S. (2006). Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334.
- Beck, A. and Teboulle, M. (2003). Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175.
- Berkson, J. (1944). Application of the logistic function to bio-assay. *Journal of the American Statistical Association*, 39(227):357–365.
- Beygelzimer, A., Orabona, F., and Zhang, C. (2017). Efficient online bandit multiclass learning with $o(\sqrt{T})$ regret. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 488–497.
- Bhatia, R. and Davis, C. (2000). A better bound on the variance. *The American Mathematical Monthly*, 107(4):353–357.
- Bilodeau, B., Foster, D. J., and Roy, D. M. (2020). Improved bounds on minimax regret under logarithmic loss via self-concordance. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*.

- Bubeck, S. (2015). Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3–4):231–357.
- Bubeck, S., Cesa-Bianchi, N., and Kakade, S. M. (2012). Towards minimax policies for online linear optimization with bandit feedback. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*, pages 41.1–41.14.
- Bubeck, S., Dekel, O., Koren, T., and Peres, Y. (2015). Bandit convex optimization: \sqrt{T} regret in one dimension. In *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*, pages 266–278.
- Bubeck, S. and Eldan, R. (2015). The entropic barrier: a simple and optimal universal self-concordant barrier. In *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*, pages 279–279.
- Bubeck, S. and Eldan, R. (2016). Multi-scale exploration of convex functions and bandit convex optimization. In *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*, pages 583–589.
- Bubeck, S., Lee, Y. T., and Eldan, R. (2017). Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 72–85. ACM.
- Cesa-Bianchi, N., Conconi, A., and Gentile, C. (2005). A second-order perceptron algorithm. *SIAM Journal on Computing*, 34(3):640–668.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge university press.
- Cesa-Bianchi, N., Mansour, Y., and Stoltz, G. (2007). Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66(2-3):321–352.
- Chang, C.-C. and Lin, C.-J. (2011). Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):27.
- Chen, Y., Lee, C.-W., Luo, H., and Wei, C.-Y. (2019). A new algorithm for non-stationary contextual bandits: Efficient, optimal, and parameter-free. In *Proceedings of the 32nd Annual Conference On Learning Theory (COLT)*, pages 696–726.
- Chen, Z., Xu, Y., Chen, E., and Yang, T. (2018). Sadagrad: Strongly adaptive stochastic gradient methods. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 912–920.

- Chernov, A. and Vovk, V. (2010). Prediction with advice of unknown number of experts. *Uncertainty in Artificial Intelligence*, pages 117–125.
- Chiang, C.-K., Yang, T., Le, C.-J., Mahdavi, M., Lu, C.-J., Jin, R., and Zhu, S. (2012). Online optimization with gradual variations. In *Proc. of the 25th Annual Conference on Learning Theory (COLT)*, pages 6.1–6.20.
- Cover, T. M. (1991). Universal portfolios. *Mathematical Finance*, 1(1):1–29.
- Cover, T. M. and Ordentlich, E. (1996). Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363.
- Cover, T. M. and Thomas, J. A. (1991). *Elements of information theory*. John Wiley & Sons.
- Crammer, K., Kulesza, A., and Dredze, M. (2009). Adaptive regularization of weight vectors. In *Advances in Neural Information Processing Systems 22*, pages 414–422.
- Crammer, K. and Singer, Y. (2001). On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2(Dec):265–292.
- Crammer, K. and Singer, Y. (2003). Ultraconservative online algorithms for multi-class problems. *Journal of Machine Learning Research*, 3(Jan):951–991.
- Csiszár, I. (1975). I-divergence geometry of probability distributions and minimization problems. *The Annals of Probability*, 3(1):146–158.
- Cutkosky, A. (2019). Artificial constraints and hints for unbounded online learning. In *Proceedings of the 32nd Annual Conference on Learning Theory (COLT)*, pages 874–894.
- Cutkosky, A. and Boahen, K. (2017). Online learning without prior information. In *Proceedings of the 30th Annual Conference on Learning Theory (COLT)*, pages 643–677.
- Cutkosky, A. and Orabona, F. (2018). Black-box reductions for parameter-free online learning in banach spaces. In *Proceedings of the 31th Annual Conference on Learning Theory (COLT)*, pages 1493–1529.
- Dani, V., Kakade, S. M., and Hayes, T. P. (2008). The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems 21*, pages 345–352.

- Daniely, A., Gonen, A., and Shalev-Shwartz, S. (2015). Strongly adaptive online learning. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 1405–1411.
- Deswarte, R. (2018). *Linear regression and learning: contributions to regularization and aggregation methods*. PhD thesis, Université Paris-Saclay.
- Dick, T., György, A., and Szepesvári, C. (2014). Online learning in Markov decision processes with changing cost sequences. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 512–520.
- Do, C. B., Le, Q. V., and Foo, C.-S. (2009). Proximal regularization for online and batch learning. In *Proceedings of the 26th Annual International Conf. on Machine Learning (ICML)*, pages 257–264.
- Duchi, J., Hazan, E., and Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159.
- Duchi, J. C., Jordan, M. I., and Wainwright, M. J. (2014). Privacy aware learning. *Journal of the ACM (JACM)*, 61(6):38.
- Dwork, C. and Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407.
- Van Erven, T., Grünwald, P. D., Mehta, N. A., Reid, M. D., and Williamson, R. C. (2015). Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16:1793–1861.
- Fink, M., Shalev-Shwartz, S., Singer, Y., and Ullman, S. (2006). Online multiclass learning by interclass hypothesis sharing. In *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, pages 313–320.
- Flaxman, A. D., Kalai, A. T., Kalai, A. T., and McMahan, H. B. (2005). Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the 16th annual ACM-SIAM symposium on Discrete algorithms*, pages 385–394.
- Foster, D. J., Kale, S., Luo, H., Mohri, M., and Sridharan, K. (2018a). Logistic regression: The importance of being improper. In *Proceedings of the 31st Annual Conference On Learning Theory (COLT)*, pages 167–208.

- Foster, D. J., Kale, S., Mohri, M., and Sridharan, K. (2017). Parameter-free online learning via model selection. In *Advances in Neural Information Processing Systems 30*, pages 6020–6030.
- Foster, D. J. and Krishnamurthy, A. (2018). Contextual bandits with surrogate losses: Margin bounds and efficient algorithms. In *Advances in Neural Information Processing Systems 31*, pages 2621–2632.
- Foster, D. J., Krishnamurthy, A., and Luo, H. (2019). Model selection for contextual bandits. In *Advances in Neural Information Processing Systems 32*, pages 14714–14725.
- Foster, D. J., Rakhlin, A., and Sridharan, K. (2015). Adaptive online learning. In *Advances in Neural Information Processing Systems 28*, pages 3375–3383.
- Foster, D. J., Rakhlin, A., and Sridharan, K. (2018b). Online learning: Sufficient statistics and the burkholder method. In *Proceedings of the 31th Annual Conference on Learning Theory (COLT)*, pages 3028–3064.
- Freund, Y. and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139.
- Freund, Y., Schapire, R. E., Singer, Y., and Warmuth, M. K. (1997). Using and combining predictors that specialize. In *Proc. 29th Annual ACM Symposium on Theory of Computing*, pages 334–343. ACM.
- Ghashami, M., Liberty, E., Phillips, J. M., and Woodruff, D. P. (2016). Frequent directions: Simple and deterministic matrix sketching. *SIAM Journal on Computing*, 45(5):1762–1792.
- Golub, G. H. and Van Loan, C. F. (2012). *Matrix computations*, volume 3. JHU Press.
- Grünwald, P. D. (2007). *The minimum description length principle*. MIT press.
- Hazan, E., Agarwal, A., and Kale, S. (2007). Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192.
- Hazan, E. et al. (2016). Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325.
- Hazan, E. and Kale, S. (2010). Extracting certainty from uncertainty: Regret bounded by variation in costs. *Machine learning*, 80(2-3):165–188.

- Hazan, E. and Kale, S. (2011). Newtron: an efficient bandit algorithm for online multiclass prediction. In *Advances in Neural Information Processing Systems 24*, pages 891–899.
- Hazan, E. and Kale, S. (2015). An online portfolio selection algorithm with regret logarithmic in price variation. *Mathematical Finance*, 25(2):288–310.
- Hazan, E., Koren, T., and Levy, K. Y. (2014). Logistic regression: Tight bounds for stochastic and online optimization. In *Proceedings of the 27th Annual Conference on Learning Theory (COLT)*, pages 197–209.
- Hazan, E. and Levy, K. (2014). Bandit convex optimization: Towards tight bounds. In *Advances in Neural Information Processing Systems 27*, pages 784–792.
- Helmbold, D. P., Schapire, R. E., Singer, Y., and Warmuth, M. K. (1998). On-line portfolio selection using multiplicative updates. *Mathematical Finance*, 8(4):325–347.
- Helmbold, D. P. and Warmuth, M. K. (2009). Learning permutations with Exponential Weights. *Journal of Machine Learning Research*, 10:1705–1736.
- Hiriart-Urruty, J.-B. (2006). A note on the legendre-fenchel transform of convex composite functions. In *Nonsmooth Mechanics and Analysis*, pages 35–46. Springer.
- Ihara, S. (1993). *Information Theory for Continuous Systems*, volume 2. World Scientific.
- Jain, P., Kothari, P., and Thakurta, A. (2012). Differentially private online learning. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*, pages 24–1.
- Jain, P. and Thakurta, A. G. (2014). (near) dimension independent risk bounds for differentially private learning. In *Proceedings of the 31th International Conference on Machine Learning (ICML)*, pages 476–484.
- Jun, K.-S. and Orabona, F. (2019). Parameter-free online convex optimization with sub-exponential noise. *Proceedings of the 32nd Annual Conference on Learning Theory (COLT)*, pages 1802–1823.
- Jézéquel, R., Gaillard, P., and Rudi, A. (2020). Efficient improper learning for online logistic regression. In Abernethy, J. and Agarwal, S., editors, *Proceedings of 33rd Annual Conference on Learning Theory (COLT)*, volume 125 of *Proceedings of Machine Learning Research*, pages 2085–2108. PMLR.

- Kakade, S. M., Shalev-Shwartz, S., and Tewari, A. (2008). Efficient bandit algorithms for online multiclass prediction. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, pages 440–447.
- Kalai, A. and Vempala, S. (2002). Efficient algorithms for universal portfolios. *Journal of Machine Learning Research*, 3(Nov):423–440.
- Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., and Smith, A. (2011). What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826.
- Kempka, M., Kotowski, W., and Warmuth, M. K. (2019). Adaptive scale-invariant online algorithms for learning linear models. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 3321–3330.
- Kivinen, J. and Warmuth, M. K. (1997). Exponentiated Gradient versus Gradient Descent for linear predictors. *Information and Computation*, 132(1):1–63.
- Kleinberg, R. and Leighton, T. (2003). The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 594–605. IEEE.
- Kleinberg, R. D. (2005). Nearly tight bounds for the continuum-armed bandit problem. In *Advances in Neural Information Processing Systems 18*, pages 697–704.
- Koolen, W. M. (2015). The relative entropy bound for Squint. *Blog August 13*: http://blog.wouterkoolen.info/Squint_PAC/post.html.
- Koolen, W. M. (2016). Exploiting curvature using Exponential Weights. *Blog September 6*: <http://blog.wouterkoolen.info/EW4Quadratic/post.html>.
- Koolen, W. M., Van Erven, T., and Grünwald, P. D. (2014). Learning the learning rate for prediction with expert advice. In *Advances in Neural Information Processing Systems 27*, pages 2294–2302.
- Koolen, W. M., Grünwald, P., and Van Erven, T. (2016). Combining adversarial guarantees and stochastic fast rates in online learning. In *Advances in Neural Information Processing Systems 29*, pages 4457–4465.

- Koolen, W. M. and Van Erven, T. (2015). Second-order quantile methods for experts and combinatorial games. In *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*, pages 1155–1175.
- Kotłowski, W. (2017). Scale-invariant unconstrained online learning. In *Proceedings of the 28th International Conference on Algorithmic Learning Theory (ALT)*, pages 412–433.
- Krichevsky, R. and Trofimov, V. (1981). The performance of universal encoding. *IEEE Transactions on Information Theory*, 27(2):199–207.
- Langford, J. and Zhang, T. (2008). The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in Neural Information Processing Systems 21*, pages 817–824.
- Lattimore, T. (2015). The pareto regret frontier for bandits. In *Advances in Neural Information Processing Systems 28*, pages 208–216.
- Lattimore, T. and Szepesvári, C. (2018). *Bandit algorithms*. Cambridge University Press.
- Li, B. and Hoi, S. C. (2014). Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)*, 46(3):1–36.
- Littlestone, N. and Warmuth, M. K. (1994). The Weighted Majority algorithm. *Information and Computation*, 108(2):212–261.
- Luo, H., Agarwal, A., Cesa-Bianchi, N., and Langford, J. (2017). Efficient second order online learning by sketching. *ArXiv preprint: arXiv:1602.02202*.
- Luo, H., Wei, C.-Y., and Zheng, K. (2018). Efficient online portfolio with logarithmic regret. In *Advances in Neural Information Processing Systems 31*, pages 8235–8245.
- Luo, L., Chen, C., Zhang, Z., Li, W.-J., and Zhang, T. (2019). Robust frequent directions with application in online learning. *Journal of Machine Learning Research*, 20(45):1–41.
- McMahan, B. and Streeter, M. (2012). No-regret algorithms for unconstrained online convex optimization. In *Advances in Neural Information Processing Systems 25*, pages 2402–2410.
- McMahan, H. B. and Orabona, F. (2014). Unconstrained online linear learning in hilbert spaces: Minimax algorithms and normal approximations. In *Proceedings of the 27th Annual Conference on Learning Theory (COLT)*, pages 1020–1039.

- McMahan, H. B. and Streeter, M. (2010). Adaptive bound optimization for online convex optimization. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, pages 244–256.
- Mhammedi, Z. and Koolen, W. M. (2020). Lipschitz and comparator-norm adaptivity in online learning. In *Proceedings of the 33rd Annual Conference on Learning Theory (COLT)*, pages 2858–2887.
- Mhammedi, Z., Koolen, W. M., and Van Erven, T. (2019). Lipschitz adaptivity with multiple learning rates in online learning. In Beygelzimer, A. and Hsu, D., editors, *Proceedings of the 32nd Annual Conference on Learning Theory (COLT)*, pages 2490–2511.
- Mikolov, T., Chen, K., Corrado, G. S., and Dean, J. (2013). Efficient estimation of word representations in vector space. International Conference on Learning Representations. Arxiv.org/abs/1301.3781.
- Narayanan, H. and Rakhlin, A. (2017). Efficient sampling from time-varying log-concave distributions. *The Journal of Machine Learning Research*, 18(1):4017–4045.
- Nemirovski, A. (2004). Lecture notes: Interior point polynomial time methods in convex programming. *Spring Semester*.
- Nesterov, Y. (2009). Primal-dual subgradient methods for convex problems. *Mathematical programming*, 120(1):221–259.
- Nesterov, Y. and Nemirovskii, A. (1994). *Interior-point polynomial algorithms in convex programming*. SIAM.
- Neu, G. and Zhivotovskiy, N. (2020). Fast rates for online prediction with abstention. In *Proceedings of the 33rd Annual Conference on Learning Theory (COLT)*, pages 3030–3048.
- Neuteboom, T. W. H. (2020). Modifying Squint for prediction with expert advice in a changing environment. *Bachelor Thesis*. To appear at <https://www.universiteitleiden.nl/en/science/mathematics/education/theses#bachelor-theses-mathematics>.
- Nielsen, F. and Nock, R. (2010). Entropies and cross-entropies of exponential families. In *Image Processing (ICIP), 2010 17th IEEE International Conference on*, pages 3621–3624. IEEE.

- Orabona, F. (2014). Simultaneous model selection and optimization through parameter-free stochastic learning. In *Advances in Neural Information Processing Systems 27*, pages 1116–1124.
- Orabona, F., Cesa-Bianchi, N., and Gentile, C. (2012). Beyond logarithmic bounds in online learning. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 823–831.
- Orabona, F., Crammer, K., and Cesa-Bianchi, N. (2015a). A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435.
- Orabona, F., Crammer, K., and Cesa-Bianchi, N. (2015b). A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435.
- Orabona, F. and Pál, D. (2016). Coin betting and parameter-free online learning. In *Advances in Neural Information Processing Systems 29*, pages 577–585.
- Orabona, F. and Pál, D. (2018). Scale-free online learning. *Theoretical Computer Science*, 716:50–69.
- Orabona, F. and Tommasi, T. (2017). Training deep networks without learning rates through coin betting. In *Advances in Neural Information Processing Systems 30*, pages 2160–2170.
- Orseau, L., Lattimore, T., and Legg, S. (2017). Soft-bayes: Prod for mixtures of experts with log-loss. In *Proceedings of the 28th International Conference on Algorithmic Learning Theory (ALT)*, pages 372–399.
- Rennie, J. D. and Srebro, N. (2005). Loss functions for preference levels: Regression with discrete ordered labels. In *Proceedings of the IJCAI multidisciplinary workshop on advances in preference handling*, volume 1, pages 180–186.
- De Rooij, S., Van Erven, T., Grünwald, P. D., and Koolen, W. M. (2014). Follow the Leader if you can, Hedge if you must. *Journal of Machine Learning Research*, 15:1281–1316.
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386.
- Saha, A. and Tewari, A. (2011). Improved regret guarantees for online smooth convex optimization with bandit feedback. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 636–642.

- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117.
- Shalev-Shwartz, S. (2011). Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194.
- Shalev-Shwartz, S., Singer, Y., Srebro, N., and Cotter, A. (2011). Pegasos: Primal estimated sub-gradient solver for SVM. *Mathematical Programming*, 127(1):3–30.
- Shamir, G. I. (2020). Logistic regression regret: What’s the catch? In *Proceedings of the 33rd Annual Conference on Learning Theory (COLT)*, pages 3296–3319.
- Song, S., Chaudhuri, K., and Sarwate, A. (2015). Learning from data with heterogeneous noise using sgd. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 894–902.
- Srebro, N., Sridharan, K., and Tewari, A. (2010). Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems 23*, pages 2199–2207.
- Steinhardt, J. and Liang, P. (2014). Adaptivity and optimism: An improved exponentiated gradient algorithm. In *Proceedings of the 31st Annual International Conf. on Machine Learning (ICML)*, pages 1593–1601.
- Thakurta, A. G. and Smith, A. (2013). (nearly) optimal algorithms for private online learning in full-information and bandit settings. In *Advances in Neural Information Processing Systems 26*, pages 2733–2741.
- Van der Hoeven, D. (2019). User-specified local differential privacy in unconstrained adaptive online learning. In *Advances in Neural Information Processing Systems 32*, pages 14103–14112.
- Van der Hoeven, D. (2020). Exploiting the surrogate gap in online multiclass classification. *To Appear in Advances in Neural Information Processing Systems 33*.
- Van der Hoeven, D., Cutkosky, A., and Luo, H. (2020). Comparator-adaptive convex bandits. *To Appear in Advances in Neural Information Processing Systems 33*.
- Van der Hoeven, D., Van Erven, T., and Kotłowski, W. (2018). The many faces of exponential weights in online learning. In *Proceedings of the 31st Annual Conference on Learning Theory (COLT)*, pages 2067–2092.

- van Erven, T. and Koolen, W. M. (2016). Metagrad: Multiple learning rates in online learning. In *Advances in Neural Information Processing Systems 29*, pages 3666–3674.
- Van Erven, T., Koolen, W. M., and Van der Hoeven, D. (2020a). Metagrad: Universal adaptation using multiple learning rates in online learning. *Manuscript in preparation*.
- Van Erven, T., Van der Hoeven, D., Kotłowski, W., and Koolen, W. M. (2020b). Open problem: Fast and optimal online portfolio selection. In *Proceedings of the 33rd Annual Conference on Learning Theory (COLT)*, pages 3864–3869.
- Vovk, V. (1990). Aggregating strategies. *Proceedings of 3rd Annual Workshop on Computational Learning Theory (COLT)*, page pages 371–383.
- Vovk, V. and Zhdanov, F. (2009). Prediction with expert advice for the brier game. *Journal of Machine Learning Research*, 10(Nov):2445–2471.
- Vovk, V. G. (2001). Competitive on-line statistics. *International Statistical Review*, 69(2):213–248.
- Wasserman, L. and Zhou, S. (2010). A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389.
- Xiao, L. (2010). Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11:2543–2596.
- Xie, Q. and Barron, A. R. (2000). Asymptotic minimax regret for data compression, gambling, and prediction. *IEEE Transactions on Information Theory*, 46(2):431–445.
- Zhang, L., Wang, G., Tu, W., and Zhou, Z. (2019). Dual adaptivity: A universal algorithm for minimizing the adaptive regret of convex functions. *CoRR*, abs/1906.10851.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML)*, pages 928–936.
- Zinkevich, M. (2004). *Theoretical Guarantees for Algorithms in Multi-Agent Settings*. PhD thesis, Carnegie Mellon University.