



















## REFERENCES

- [1] Peter Auer, Ronald Ortner, and Csaba Szepesvári. 2007. Improved rates for the stochastic continuum-armed bandit problem. In *Proceedings of the International Conference on Machine Learning (ICML)*. Springer, 454–468.
- [2] Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5, 1 (2012), 1–122.
- [3] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. 2009. Pure exploration in multi-armed bandits problems. In *Proceedings of the International conference on Algorithmic learning theory (ALT)*. 23–37.
- [4] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. 2011. X-armed bandits. *Journal of Machine Learning Research* 12, May (2011), 1655–1695.
- [5] Adam D Bull. 2011. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research* 12, Oct (2011), 2879–2904.
- [6] Roberto Calandra, André Seyfarth, Jan Peters, and Marc Peter Deisenroth. 2014. An experimental comparison of Bayesian optimization for bipedal locomotion. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1951–1958.
- [7] Wei Chen, Yajun Wang, and Yang Yuan. 2013. Combinatorial multi-armed bandit: General framework and applications. In *Proceedings of the International Conference on Machine Learning (ICML)*. 151–159.
- [8] Mung Chiang, Prashanth Hande, Tian Lan, Chee Wei Tan, et al. 2008. Power control in wireless cellular networks. *Foundations and Trends® in Networking* 2, 4 (2008), 381–533.
- [9] Sayak Ray Chowdhury and Aditya Gopalan. 2017. On kernelized multi-armed bandits. In *Proceedings of the International Conference on Machine Learning (ICML)*. 844–853.
- [10] Carlo D’Eramo, Alessandro Nuara, Matteo Pirota, and Marcello Restelli. 2017. Estimating the maximum expected value in continuous reinforcement learning problems. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*. 1840–1846.
- [11] Carlo D’Eramo, Marcello Restelli, and Alessandro Nuara. 2016. Estimating maximum expected value through gaussian approximation. In *Proceedings of the International Conference on Machine Learning (ICML)*. 1032–1040.
- [12] Subhashish Ghosal and Anindya Roy. 2006. Posterior consistency of Gaussian process prior for nonparametric binary regression. *The Annals of Statistics* 34, 5 (2006), 2413–2429.
- [13] Steffen Grünewälder, Jean-Yves Audibert, Manfred Opper, and John Shawe-Taylor. 2010. Regret bounds for Gaussian process bandit problems. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*. 273–280.
- [14] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. 2012. Thompson sampling: An asymptotically optimal finite-time analysis. In *Proceedings of the International conference on Algorithmic learning theory (ALT)*. 199–213.
- [15] Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. 2008. Multi-armed bandits in metric spaces. In *Proceedings of the annual ACM symposium on Theory of computing (STOC)*. ACM, 681–690.
- [16] Andreas Krause and Cheng S Ong. 2011. Contextual gaussian process bandit optimization. In *Proceedings of the neural information processing systems conference (NIPS)*. 2447–2455.
- [17] Harold J Kushner. 1964. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering* 86, 1 (1964), 97–106.
- [18] Daniel J Lizotte, Tao Wang, Michael H Bowling, and Dale Schuurmans. 2007. Automatic Gait Optimization with Gaussian Process Regression. In *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*, Vol. 7. 944–949.
- [19] Jonas Mockus. 1975. On Bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference*. Springer, 400–404.
- [20] Alessandro Nuara, Nicola Sosio, Francesco Trovò, Maria Chiara Zaccardi, Nicola Gatti, and Marcello Restelli. 2019. Dealing with Interdependencies and Uncertainty in Multi-Channel Advertising Campaigns Optimization. In *Proceedings of the ACM World Wide Web Conference (WWW)*. 1376–1386.
- [21] Alessandro Nuara, Francesco Trovò, Nicola Gatti, and Marcello Restelli. 2018. A Combinatorial-Bandit Algorithm for the Online Joint Bid/Budget Optimization of Pay-per-Click Advertising Campaigns. In *Proceedings of the Conference on Artificial Intelligence (AAAI)*. 2379–2386.
- [22] C. E. Rasmussen and C. K. Williams. 2006. *Gaussian Processes for Machine Learning*. Vol. 1. MIT press Cambridge.
- [23] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *Proceedings of the neural information processing systems conference (NIPS)*. 2951–2959.
- [24] N. Srinivas, A. Krause, S.M. Kakade, and M. Seeger. 2010. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the International Conference on Machine Learning (ICML)*. 1015–1022.
- [25] Francesco Trovò, Stefano Paladino, Marcello Restelli, and Nicola Gatti. 2016. Budgeted multi-armed bandit in continuous action space. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*. IOS Press, 560–568.
- [26] Francesco Trovò, Stefano Paladino, Marcello Restelli, and Nicola Gatti. 2018. Improving multi-armed bandit algorithms in online pricing settings. *International Journal of Approximate Reasoning* 98 (2018), 196–235.
- [27] Zi Wang and Stefanie Jegelka. 2017. Max-value entropy search for efficient Bayesian optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*. JMLR. org, 3627–3635.
- [28] Zi Wang, Bolei Zhou, and Stefanie Jegelka. 2016. Optimization as estimation with Gaussian processes in bandit settings. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*. 1022–1031.