Prior-mean-RObust Bayesian Optimization (PROBO)

Julian Rodemann[1]

[1]Department of Statistics, LMU Munich
julian@stat.uni-muenchen.de

Abstract for Young Statisticians Session (YSS)

Bayesian optimization (BO) with Gaussian processes (GP) as surrogates is used to optimize functions that are expensive to evaluate and lack analytical description, e.g. in hyperparameter-tuning of predictive models [3]. In my presentation at the YSS, I would like to propose Prior-mean-RObust Bayesian Optimization (PROBO). In the first part of my talk, the effect of the GP prior specifications on classical BO's convergence is studied. I find the prior's mean parameters to have the highest influence on convergence among all prior components. In response to this result, I introduce PROBO that aims at rendering BO more robust towards GP prior mean parameter misspecification. This is achieved by explicitly accounting for GP imprecision via a prior near-ignorance model [1] [2]. At the heart of this is a novel acquisition function, the generalized lower confidence bound (GLCB). In the second part of my talk, I will test my approach against classical BO on a real-world problem from material science and observe PROBO to converge faster.

Keywords: Bayesian optimization · Imprecise probabilities · Prior near-ignorance · Model imprecision

Acknowledgments: I warmly thank Prof. Dr. Thomas Augustin for the supervision of my master's thesis, whose results this presentation builds on.

References

1. Benavoli, A., Zaffalon, M.: Prior near ignorance for inferences in the k-parameter exponential family. Statistics **49**(5), 1104–1140 (2015)
2. Mangili, F.: A prior near-ignorance Gaussian process model for nonparametric regression. International Journal of Approximate Reasoning **78**, 153–171 (2016)
3. Snoek, J., Larochelle, H., Adams, R.P.: Practical Bayesian optimization of machine learning algorithms. Advances in Neural Information Processing Systems **25** (2012)