Understanding High-Dimensional Bayesian Optimization

Leonard Papenmeier[1], Matthias Poloczek[2], Luigi Nardi[1,3]

[1] Lund University
[2] Amazon
[3] DBtune

Surrogate model

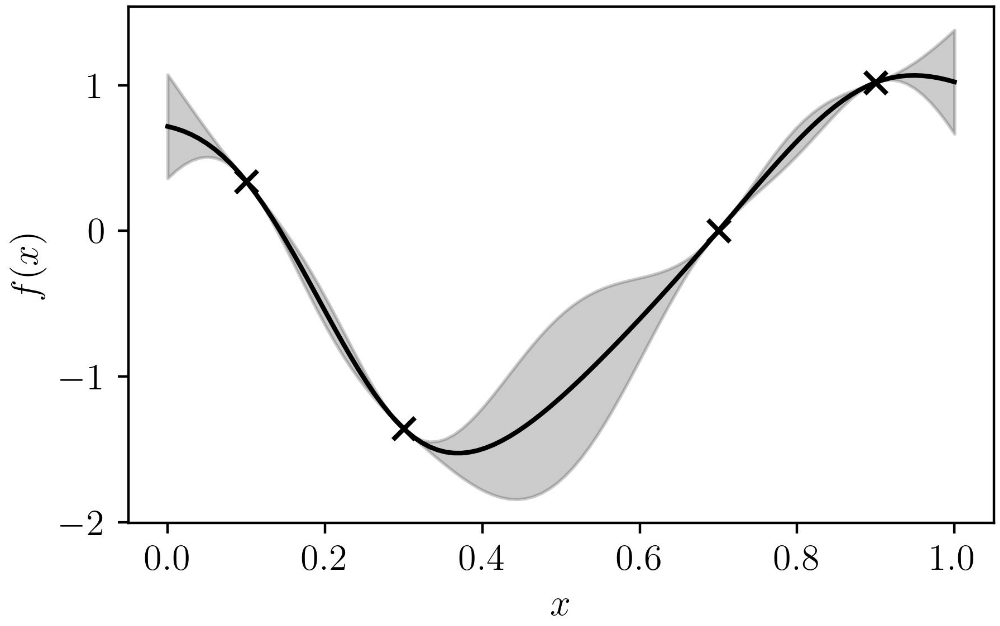

Why is High-Dimensional BO challenging?

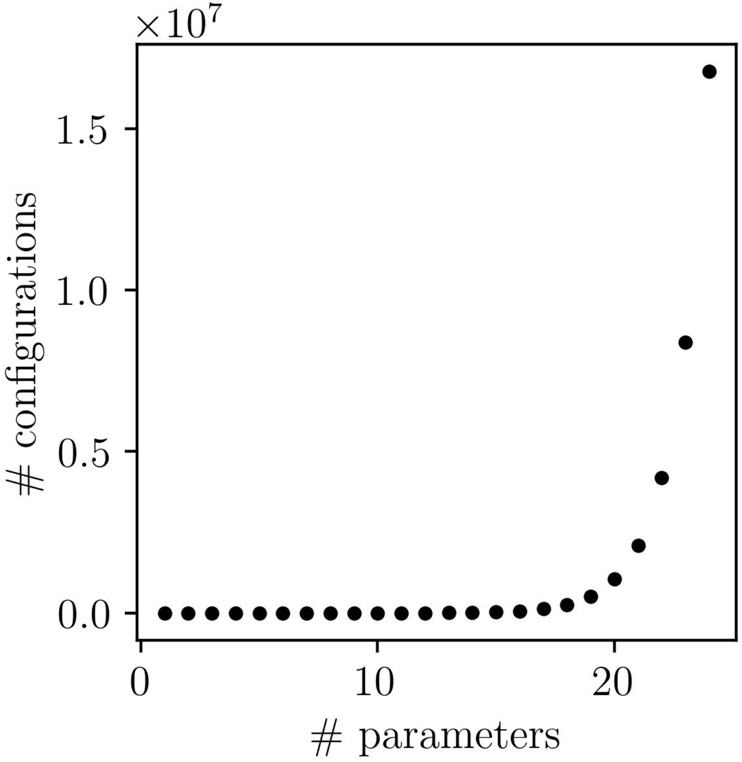

“Vanilla Bayesian Optimization Performs Great in High Dimensions”[1]

“Standard Gaussian Process Can Be Excellent for High-Dimensional Bayesian Optimization”[2]

---

[1] Hvarfner, C., Hellsten, E. O., & Nardi, L. (2024, July). Vanilla Bayesian Optimization Performs Great in High Dimensions. In International Conference on Machine Learning (pp. 20793-20817). PMLR.
[2] Xu, Z., Wang, H., Phillips, J. M., & Zhe, S. (2024). Standard Gaussian Process Can Be Excellent for High-Dimensional Bayesian Optimization. arXiv preprint arXiv:2402.02746.

MLL Gradients

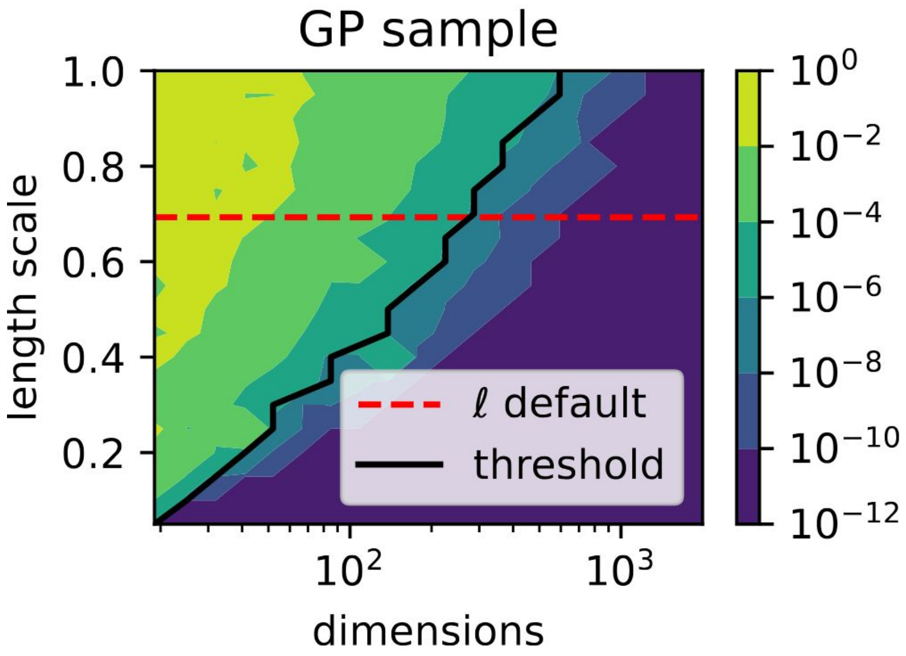

Acquisition Function Maximization Without Local Samples

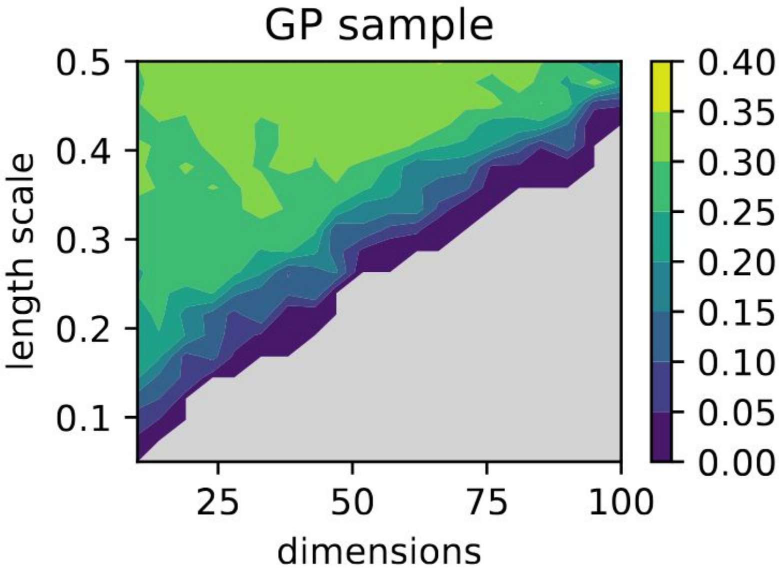

Acquisition Function Maximization With Local Samples

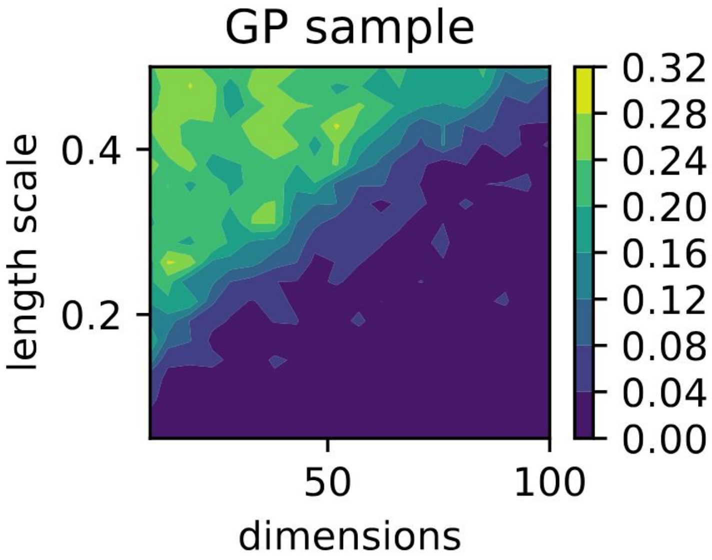

MLE vs MAP

	MLE	MAP
non-vanishing gradients?	✅ (if initialized properly)	✅
unbiased?	✅	❌

MSR = “MLE - scaled with RAASP”

Performance plots

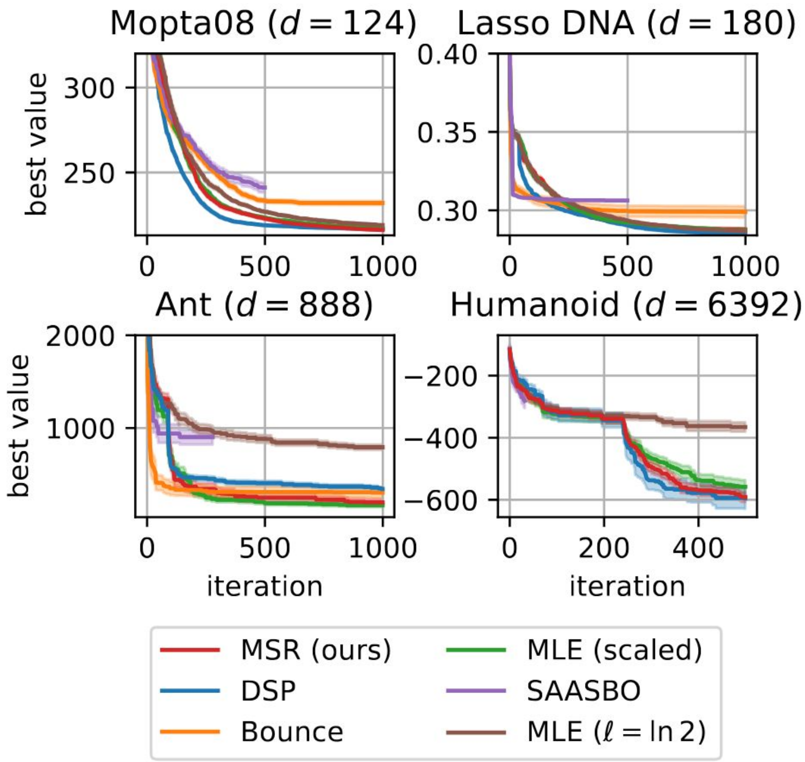

Thank you!

References

[1] David K Duvenaud, Hannes Nickisch, and Carl Rasmussen. Additive Gaussian Processes. Advances in neural information processing systems, 24, 2011.

[2] Kirthevasan Kandasamy, Jeff Schneider, and Barnabás Póczos. High dimensional Bayesian optimisation and bandits via additive models. In International conference on machine learning (ICML), pages 295–304, 2015.

[3] Jacob Gardner, Chuan Guo, Kilian Weinberger, Roman Garnett, and Roger Grosse. Discovering and exploiting additive structure for Bayesian optimization. In International Conference on Artificial Intelligence and Statistics, pages 1311–1319, 2017.

[4] Mickael Binois and Nathan Wycoff. A survey on high-dimensional Gaussian process modeling with application to Bayesian optimization. ACM Transactions on Evolutionary Learning and Optimization, 2(2):1–26, 2022.

[5] Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando de Freitas. Bayesian Optimization in a Billion Dimensions via Random Embeddings. Journal of Artificial Intelligence Research (JAIR), 55:361–387, 2016.

[6] Amin Nayebi, Alexander Munteanu, and Matthias Poloczek. A framework for Bayesian Optimization in Embedded Subspaces. In Proceedings of the 36th International Conference on Machine Learning, volume 97 of Proceedings of Machine Learning Research (PMLR), pages 4752–4761, 09–15 Jun 2019.

[7] Ben Letham, Roberto Calandra, Akshara Rai, and Eytan Bakshy. Re-Examining Linear Embeddings for High-Dimensional Bayesian Optimization. In Advances in Neural Information Processing Systems (NeurIPS), volume 33, pages 1546–1558, 2020.

[8] David Eriksson and Martin Jankowiak. High-dimensional bayesian optimization with sparse axis-aligned subspaces. In Uncertainty in Artificial Intelligence, pages 493–503. PMLR, 2021.

[9] Carl Hvarfner, Erik Orm Hellsten, and Luigi Nardi. Vanilla Bayesian Optimization Performs Great in High Dimensions. In Proceedings of the 41st International Conference on Machine Learning, volume 235 of Proceedings of Machine Learning Research, pages 20793–20817. PMLR, 21–27 Jul 2024.

[10] Zhitong Xu and Shandian Zhe. Standard Gaussian Process is All You Need for High-Dimensional Bayesian Optimization. arXiv preprint arXiv:2402.02746, 2024.

[11] Zhitong Xu, Haitao Wang, Jeff M. Phillips, and Shandian Zhe. Standard Gaussian Process is All You Need for High-Dimensional Bayesian Optimization. In TheThirteenth International Conference on Learning Representations, 2025.