BoGraph: Structured Bayesian Optimization From Logs for Expensive Systems with Many Parameters

Sami Alabed, Eiko Yoneki
University of Cambridge, UK
The Alan Turing Institute, UK

Metrics and logs provide context to auto-tuners leading to better optimization in fewer evaluations

Computer system trace reflects the inner working of the system. BoGraph's pipeline leverage the trace through its pipeline and, using causal learning, finds a structure that contextualizes the model and optimizes the system.

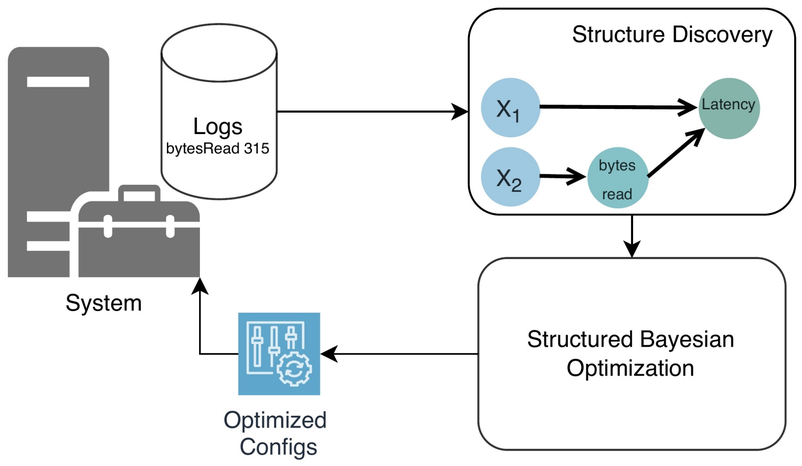

Fig 1. BoGraph ingests the system's logs to build a structure then perform structured Bayesian optimization and return optimized configurations.

BoGraph learns a system structure from logs then optimizes it using structured Bayesian optimization

BoGraph incorporates both expert knowledge and learning from logs to build a structure for structured Bayesian Optimization [1]. The structure is learned using causal structure learning [2] and acts as a blueprint to build a complicated model that reflects the system's behavior.

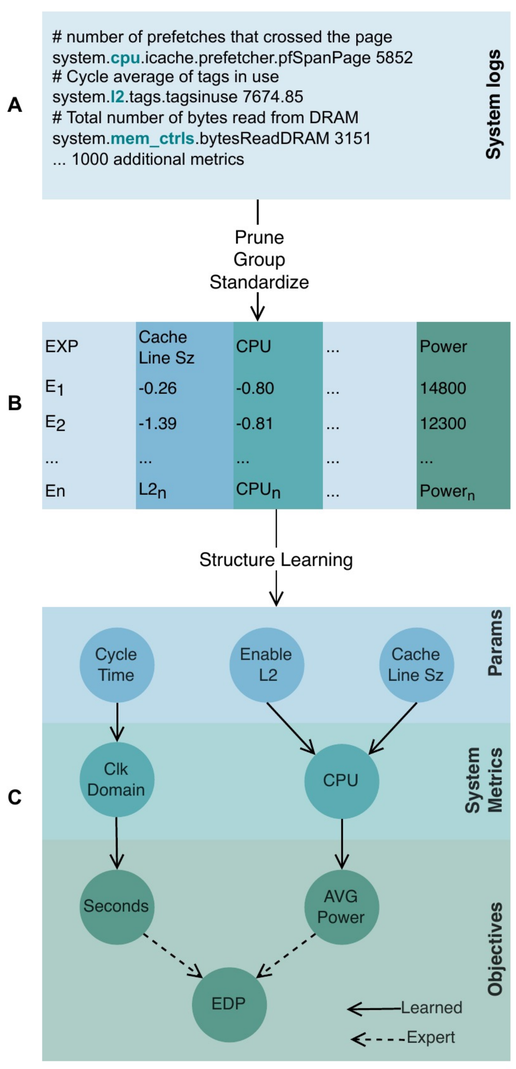

Fig 3. BoGraph pipeline processes the logs (A), performs causal structure discovery (B) and adds any expert knowledge producing a graph representing the statistical association between metrics, objectives, and parameters (C).

Structured Bayesian optimization from logs finds better configurations in fewer evaluations than other methods

BoGraph optimized 20 design parameters of gem5-Aladdin, a system accelerator simulator, improving its energy-delay-product (EDP) over default by $5x-7x$ factors (Fig 4). BoGraph ingested the system stats and found meaningful structure (Fig 3C). Using the structure, BoGraph found optimized configuration in fewer evaluations than other methods (Fig 5).

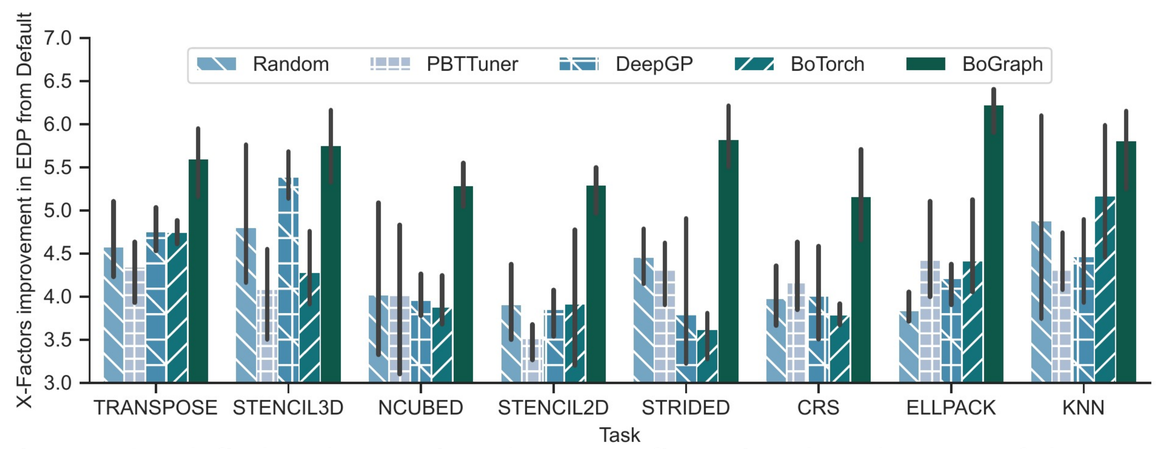

Fig 4. BoGraph finds at least twice a better configuration than the state-of-the-art methods on the MachSuite benchmark.

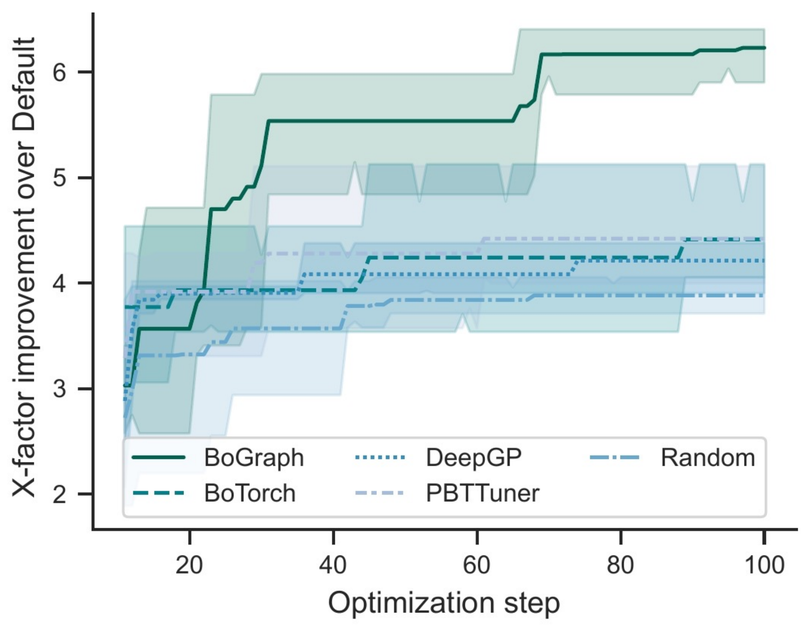

Fig 5. Auto-tuners' convergence rate and improvement over the default. BoGraph consistently finds better optimizations in fewer evaluations.

BoGraph provides optimized configurations to slow to evaluate systems with many parameters

Utilizing a mixture of learned structure and expert provided structure contextualizes the model, allowing BoGraph to find optimal configurations quickly despite many parameters. BoGraph is beneficial to any system that needs tuning and has access to metrics or expert knowledge. Next, we are working on a case study on PostgreSQL.

References

[1] Dalibard, V., Schaarschmidt, M. and Yoneki, E., 2017, April. BOAT: Building auto-tuners with structured Bayesian optimization. In Proceedings of the 26th International Conference on World Wide Web (pp. 479-488).

[2] Zheng, X., Aragam, B., Ravikumar, P.K. and Xing, E.P., 2018. Dags with no tears: Continuous optimization for structure learning. Advances in Neural Information Processing Systems, 31.