Optimizing Neural Networks with Multi-Objective Bayesian Optimization and Augmented Random Search

Workshop on Optimization and Machine Learning, March 2023

Mark Deutel, Georgios Kontes, Christopher Mutschler, and Jürgen Teich

Motivation

Problem

- To enable DNN deployment at the edge, there are often tight resource constraints that can be met by network pruning, weight quantization, and other graph-based and algorithmic optimizations[1] that can be controlled by hyperparameters.
- However, optimizing the assignment of values to these hyperparameters, i.e., solving a multi-objective optimization problem, is expensive because the DNN training and compression pipeline must be executed once for each possible combination of parameters.

Goal

- Use Multi-Objective Bayesian Optimization[2] (MoBOpt) to enable optimized DNN deployment with a limited search budget.
- Investigate and improve MoBOpt strategies such as ParEGO[3], for which we observed that in some situations they fail to reliably find the global optimum.

Contribution

- We propose to use an ensemble of competing local Reinforcement Learning (RL) agents, which we train on different parts of the surrogate model using Augmented Random Search[4] (ARS) to more effectively exploit their encoded knowledge of the target space.
- We show experimentally that we can improve on several existing MoBOpt and evolutionary strategies for two different datasets and DNN architectures.

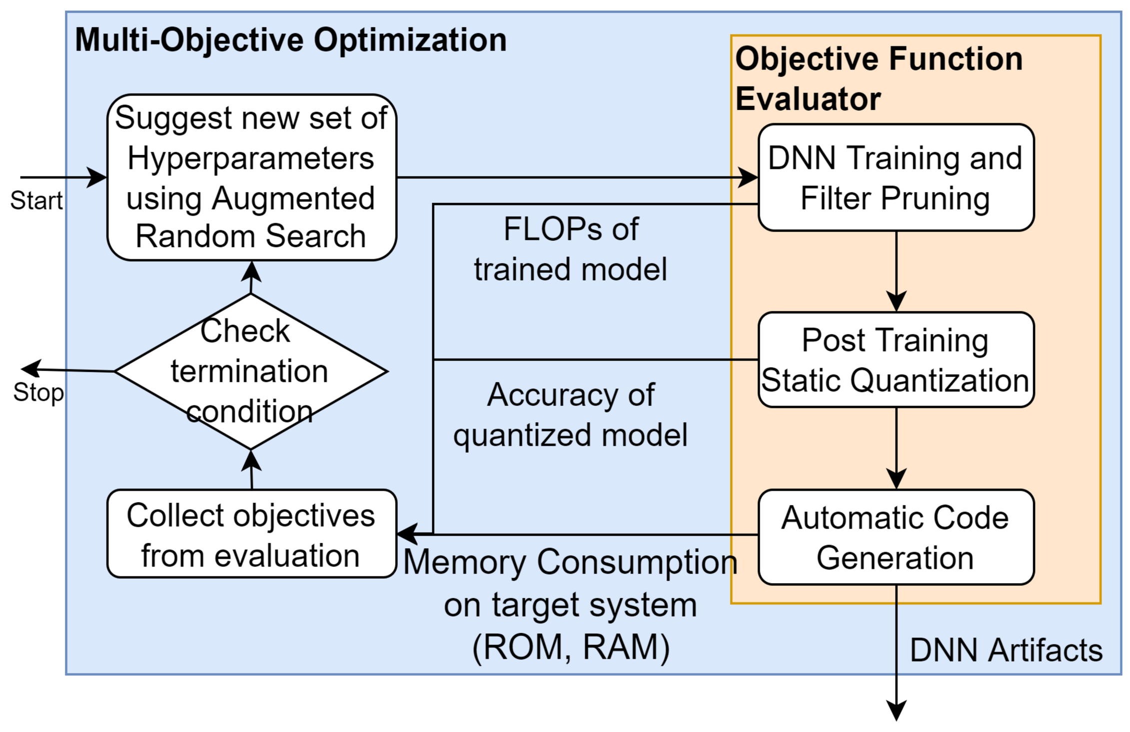

Figure 1. Overview of our RL-based multi-objective Bayesian DNN optimization approach.

Method

- Perform optimization of the objective function $f(x) = [f_1(x), \dots, f_n(x)]$ by sampling sets of hyperparameters until the search budget is exhausted.
- For each sample, perform a nested optimization on the Bayesian surrogate model fitted by maximising the marginal likelihood of previous evaluations of $f(x)$. Use a Latin hypercube during the first samples to obtain an initial prior.
- Train each competing local MLP RL policy π_l on different parts of the search space using the centers of l-means clusters of the Pareto-front as starting points. For each training step i…
  - Sample directions $\varphi_1\ \varphi_2, \dots, \varphi_N \in R^{n \times m}$ with i.i.d. standard normal entries.
  - Perform j ∈ $\{1,2,\dots,N\}$ rollouts over the horizon H from the GP using

$$\pi_{l,j,+}(x) = (\Theta_l + \nu\varphi_j)$$

$$\pi_{l,j,-}(x) = (\Theta_l - \nu\varphi_j)$$

  - Collect rewards $r(\pi_{l,j,+}(x))$ and $r(\pi_{l,j,-}(x))$ by scaling objectives to a single one using Chebyshev scalarization[3] and compute the Expected Improvement[2] (EI).
  - Evaluate EI empirically using Monte Carlo Sampling and the reparametrization trick[5]
  - Sort the N directions by $\max\{r(\pi_{l,j,+}(x)), r(\pi_{l,j,-}(x))\}$ and than update the MLP weights using a top-K selection.

$$\Theta_{l,i+1} = \Theta_{l,i} + \frac{\alpha}{K\sigma_R}\sum_{k=1}^{K}\Big(r(\pi_{l,j,+}(x)) - r(\pi_{l,j,-}(x))\Big)\varphi_k$$

    where σ_R is the standard deviation of the reward used in the update step
- Validate trained policies π_i and suggest hyperparameters proposed by the best performing policy to be evaluated next by the function evaluators (outer optimization loop).

Experiments

Experiment 1: ResNet18, CIFAR10

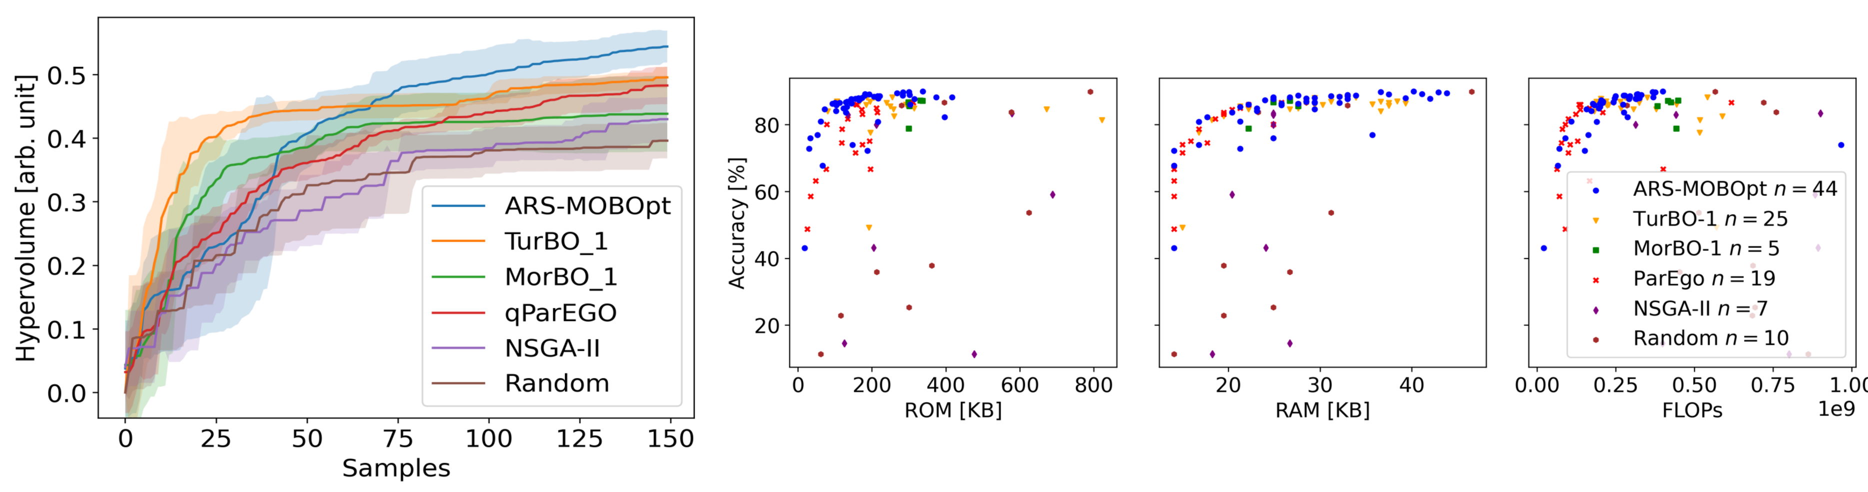

Figure 2. Hypervolume and feasible Pareto fronts of optimizations performed for two problems (ResNet18, CIFAR10 and MobileNetV3, DaLiAc) with a search budget of 150 samples and 5 seeds each.

Experiment 2: Synthetic Problem

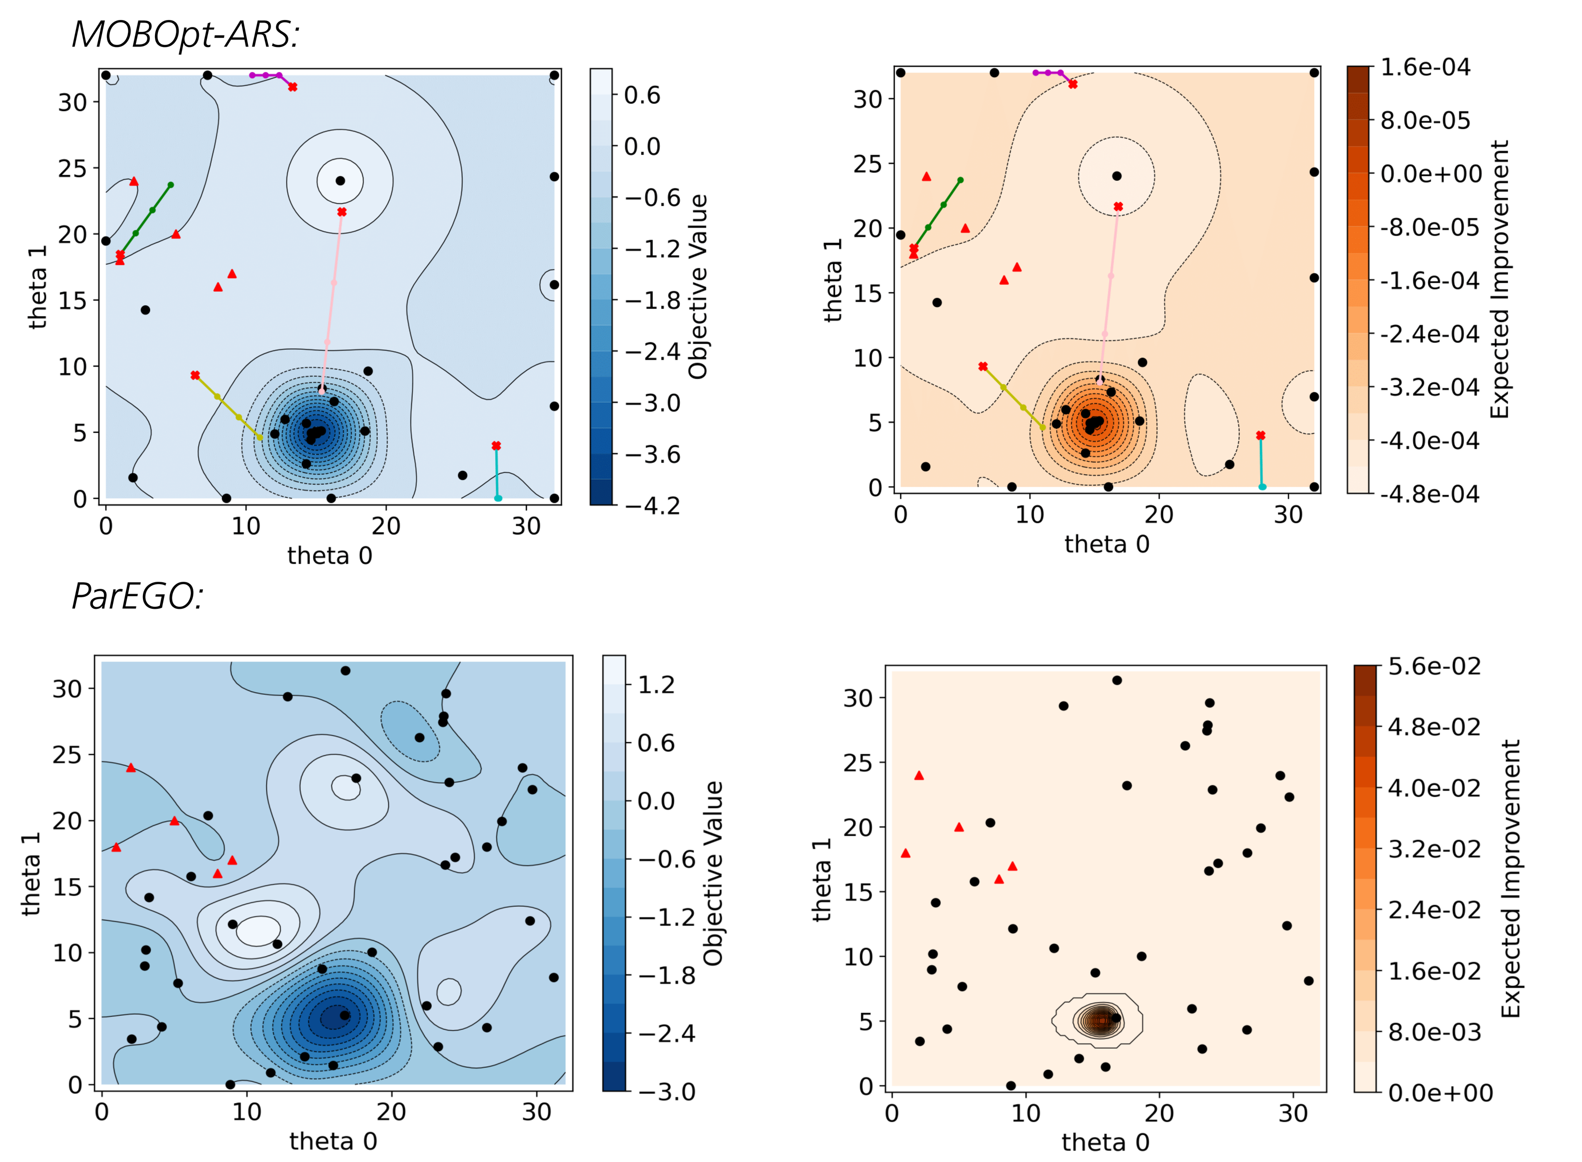

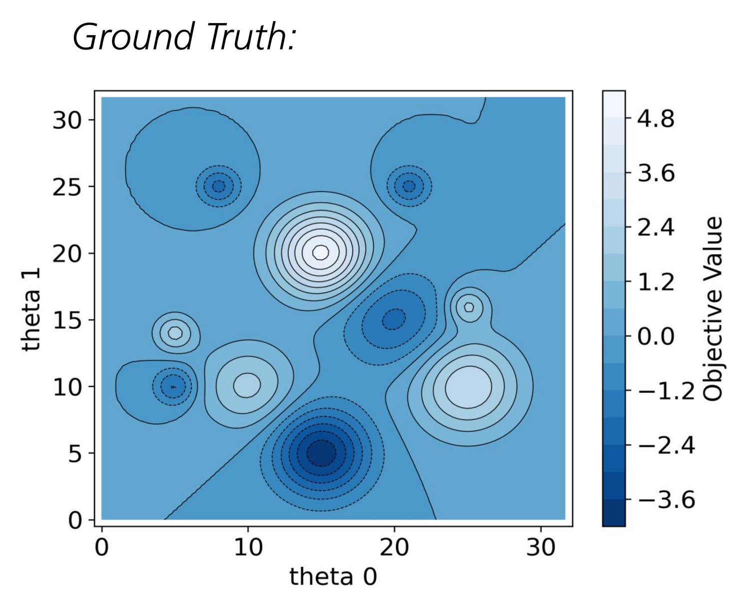

Figure 3. Topography of the objective value and the expected improvement for MOBOpt-ARS and ParEGO after a total of 40 samples, given an initial prior of 5 samples, marked as red triangles. The global minimum can be found at $\theta_0 = 15$ and $\theta_1 = 5$. For MOBOpt-ARS, the rollouts for all competing trained policies are shown as lines with their starting points marked by red crosses.

1. Deutel M., Woller P., Mutschler C., Teich J., Energy-efficient Deployment of Deep Learning Applications on Cortex-M based Microcontrollers using Deep Compression, *MBMV'23*, Freiburg, 23. March 2023 - 24. March 2023
2. Mŏckus, J., On Bayesian methods for seeking the extremum. *Optimization Techniques IFIP Technical Conference*. pp. 400–404. 1975.
3. Knowles, J., ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems, *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 1, pp. 50-66, Feb. 2006
4. Mania, H., Guy, A., and Recht, B., Simple random search of static linear policies is competitive for reinforcement learning, *Advances in Neural Information Processing Systems (NeurIPS),* 2018.
5. Wilson J. T., Moriconi R., Hutter F., Deisenroth M. P.. The Reparameterization Trick for Acquisition Functions. NeurIPS Workshop on Bayesian Optimization, 2017.