

---

# Hyperparameter Transfer Learning with Adaptive Complexity

---

Samuel Horváth  
KAUST

Aaron Klein  
Amazon Web Services

Peter Richtárik  
KAUST

Cédric Archambeau  
Amazon Web Services

## Abstract

Bayesian optimization (BO) is a sample efficient approach to automatically tune the hyperparameters of machine learning models. In practice, one frequently has to solve similar hyperparameter tuning problems sequentially. For example, one might have to tune a type of neural network learned across a series of different classification problems. Recent work on multi-task BO exploits knowledge gained from previous tuning tasks to speed up a new tuning task. However, previous approaches do not account for the fact that BO is a sequential decision making procedure. Hence, there is in general a mismatch between the number of evaluations collected in the current tuning task compared to the number of evaluations accumulated in all previously completed tasks. In this work, we enable multi-task BO to compensate for this mismatch, such that the transfer learning procedure is able to handle different data regimes in a principled way. We propose a new multi-task BO method that learns a set of ordered, non-linear basis functions of increasing complexity via nested drop-out and automatic relevance determination. Experiments on a variety of hyperparameter tuning problems show that our method improves the sample efficiency of recently published multi-task BO methods.

## 1 Introduction

Bayesian Optimization (BO) (Shahriari et al., 2015; Frazier, 2018) is a well-established framework to optimize an expensive black-box function  $f : \mathcal{X} \rightarrow \mathbb{R}$  with

---

Proceedings of the 24<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

a minimal number of function evaluations:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (1)$$

where  $\mathcal{X} \subset \mathbb{R}^D$  denotes the configuration space. We assume  $f(\mathbf{x})$  is only observed through noisy evaluations  $y \sim \text{Normal}(f(\mathbf{x}), \sigma)$ . Intuitively, BO searches for the global optimum  $\mathbf{x}^*$  by evaluating a sequence of promising candidates  $\mathbf{x}_1, \dots, \mathbf{x}_N$ . At each step, candidates are scored according to an exploration-exploitation criterion based on the evaluations collected so far. In the case of hyperparameter optimization (HPO),  $f$  denotes the validation error and  $\mathbf{x}$  is a vector of hyperparameters over which to optimize.

Let  $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$  be a set of evaluated hyperparameter configurations and their associated validation error. BO uses a probabilistic model  $p(f \mid \mathcal{D})$  of the unknown objective function  $f$  to select the next configuration by optimizing an acquisition function  $a : \mathcal{X} \rightarrow \mathbb{R}$  that performs the explore-exploit trade-off. The acquisition function only relies on the model  $p(f \mid \mathcal{D})$  and, hence, is cheap-to-evaluate compared to the objective function. Typical choices for the probabilistic model are Gaussian Processes (GPs) (Jones, 2001; Snoek et al., 2012), random forests (Hutter et al., 2011), tree parzen estimators (Bergstra et al., 2011), and Bayesian neural networks (Snoek et al., 2015; Springenberg et al., 2016; Perrone et al., 2018). Arguably, the most common acquisition function is the expected improvement (EI) (Mockus et al., 1978). EI selects the next candidate to evaluate as the one with the highest expected improvement  $a(\mathbf{x}) = \mathbb{E}[\max\{0, f^* - f(\mathbf{x})\}]$ , where  $f^*$  is the best evaluation observed so far. Numerous other acquisition functions have been proposed (Kushner, 1964; Frazier et al., 2009; Srinivas et al., 2009; Hoffman et al., 2011; Hennig and Schuler, 2012).

In this work, we focus on multi-task BO. We assume that we have already completed  $T$  related HPO tasks  $\{f_1, \dots, f_T\}$  that share the same configuration space  $\mathcal{X}$ . Our goal is to leverage knowledge acquired in previous HPO tasks to accelerate the optimization of a new HPO task  $f_{T+1}$ . In multi-task learning (see, e.g., Argyriou

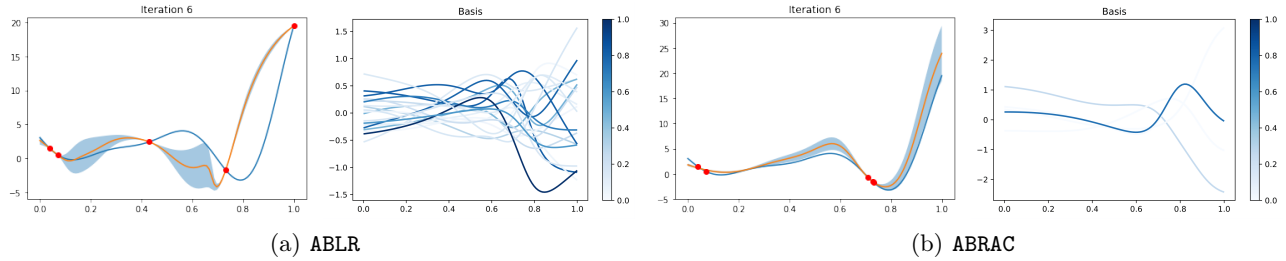


Figure 1: Comparison of the surrogate objectives along with their basis functions obtained by (a) conventional multi-task BO, implemented by ABLR Perrone et al. (2018), and (b) multi-task BO with adaptive complexity, ABRAC (see Section 4), which restricts the number of basis functions that can be activated at transfer. In the first and the third panel the target function is shown in blue, the five evaluations are the red points (which includes the same 3 initial random points for both methods); the surrogate predictive mean is shown in orange with two standard deviations in transparent light blue. The second and fourth panel show the basis functions with their color representing their relative weight used at iteration 6 of the BO algorithm. It can be observed that overfitting occurs in (a) due to the large number of basis functions used, while no overfitting occurs in (b) by restricting that number (note that there are more than 2 basis functions (5) for (b), but they are almost invisible as their relative weights are very close to zero).

et al., 2007) the number of training data available for each task is typically of the same order. Multi-task BO, similar to BO, differs from conventional multi-task learning as it is an instantiation of Bayesian sequential decision making. This means that the training data are collected *sequentially* during the training procedure, which in turns means that the number of data associated to the new task is initially much smaller than the number data that were collected in each previous task. Not accounting for this in practice is problematic as the transfer learning procedure is not adapted to the amount of information available in the target task as illustrated in Figure 1.

Our proposed method achieves transfer learning by learning a set of shared basis functions across multiple tasks. Following Perrone et al. (2018), we use an underlying shared dense neural network to learn the basis functions on top of which we place task-specific Bayesian linear regression head models. We introduce a novel regularization strategy based on nested dropout (Rippel et al., 2014), which we apply to the final layer of the neural network. This enforces the network to learn an ordered set of basis functions of increasing complexity, which in turn allows us to model different data regimes. We further leverage automatic relevance determination (ARD) to automatically determine the number of basis functions to be activated for transfer learning (MacKay et al., 1994).

Next, we summarize our contributions and relate our work to the recent literature on multi-task BO. In Section 2 and 3 we review, respectively, multi-head ABLR for multi-task BO and nested dropout. In Section 4, we describe our new multi-task BO method, which is able

to learn ordered features for a multi-head ABLR model. We present an empirical evaluation in Section 5, showing that the sample efficiency of our proposed approach is superior to conventional multi-task BO. Finally, we provide a conclusion and future directions in Section 6.

**Contributions.** Our contributions can be summarized as follows:

- We use nested dropout to learn an ordered set of features for transfer learning in the context of BO. To the best of our knowledge, no other multi-task learning approach proposed in the literature is able to learn features that take the adaptive complexity of the target task into account.
- We use ARD to automatically determine which basis functions to activate at transfer in a data-driven fashion. Hence, the resulting transfer learning model is able to adapt its capacity to the amount of data available in the target task.
- We show that we can improve the sample efficiency of multi-task BO and avoid overfitting in low data regimes without hurting the transfer learning performance in high data regimes.

Our multi-task BO method scales linearly with respect to the number of data points collected in the target BO task. This is achieved by adopting a two-step procedure: (1) a more expensive offline step consists in extracting knowledge in an ordered fashion from the completed BO tasks; (2) a fast online step exploits this knowledge to solve the new BO task data efficiently. The complexity




Figure 2: Comparison of the average regret and rank on Forrester functions for 5 different methods. Solid lines display average over all the runs of given method, transparent areas depict 2 standard deviations.

of (1) is linear in the evaluations collected across all BO tasks, while the complexity (2) is linear in the data collected in the target BO task only.

**Related Work.** Feurer et al. (2015) defined a set of hand-designed meta-features to describe similarity between tasks. Initial configurations to warm-start Bayesian optimization are then selected based on the euclidean distance of the meta-features of the current task to previous tasks. Instead of manually defined features to describe tasks, Law et al. (2019) proposed a feature extractor based on neural networks that operates directly on the underlying dataset.

Instead of learning features explicitly, Swersky et al. (2013) proposed a discrete kernel to enable GP-based BO to model the correlation across tasks. While this model does not make any strong assumptions, it does not scale well with the number of tasks. In a similar vein, Springenberg et al. (2016) proposed Bayesian neural networks for BO and showed that, by using a linear embedding that is learned together with the other parameters of the neural network, one can model the objective function across tasks. Instead of an additional input to the model, Perrone et al. (2018) considered a multi-head neural network architecture with one head per tasks. To obtain uncertainty estimates, they followed Snoek et al. (2015) by using Bayesian linear regression head models. Salinas et al. (2020) used a semi-parametric Gaussian Copula distribution to map hyperparameter configurations to quantiles of the objective function across different tasks. Most recently, Moss et al. (2020) proposed a new information-theoretic acquisition function for Multi-Task Bayesian optimization which is much faster to compute than previous existing approaches. Compared to our method, none of these transfer learning approaches learn an ordered set of features of variable complexity.

Also related to our approach is the idea to learn low-dimensional embeddings for high dimensional optimization tasks (Moriconi et al., 2019; A.Nayebi et al., 2019).

However, the key difference is that these embeddings are only learned on a single task without any transfer across tasks.

There also has been a large body of work on multi-task feature learning since the seminal work of Argyriou et al. (2007). Most of the published work focusses on learning a sparse set of features and methods that avoid negative transfer in multi-task learning (Guo et al., 2014). More recently, there has been a growing interest in transfer learning in the context of deep learning (Ruder, 2017). To the best of our knowledge, none of these techniques attempt to learn an ordered set of multi-task features.

## 2 Multi-head ABLR for Multi-task Bayesian Optimization

Adaptive Bayesian linear regression (ABLR) is Bayesian linear regression with nonlinear basis functions  $\phi_{\mathbf{z}}(\mathbf{x}) : \mathbb{R}^D \rightarrow \mathbb{R}^d$ , here parameterized by  $\mathbf{z}$  (Bishop, 2006). Perrone et al. (2018) extend this approach to the multi-task setting by maintaining a separate Bayesian linear regression head for each task on top of a *shared* neural network to learn the nonlinear basis functions:

$$\begin{aligned} p(\mathbf{y}_t | \mathbf{w}_t, \mathbf{z}, \beta_t) &= \text{Normal}(\Phi_t(\mathbf{z})\mathbf{w}_t, \beta_t^{-1}\mathbf{I}_{N_t}), \\ p(\mathbf{w}_t | \alpha_t) &= \text{Normal}(\mathbf{0}, \alpha_t^{-1}\mathbf{I}_d), \end{aligned} \quad (2)$$

where  $\mathbf{y}_t \in \mathbb{R}^{N_t}$  is the vector of observed evaluations of  $f_t$ ,  $\mathbf{w}_t \in \mathbb{R}^d$  is the weight vector of the Bayesian linear regression head  $t$ , and  $\Phi_t(\mathbf{z}) = \phi_{\mathbf{z}}(\mathbf{X}_t) \in \mathbb{R}^{N_t \times d}$  the corresponding design matrix with  $\mathbf{X}_t \in \mathbb{R}^{N_t \times D}$  containing the configurations associated to  $\mathbf{y}_t$ . The identity matrices  $\mathbf{I}_{N_t}$  and  $\mathbf{I}_d$  are respectively of size  $N_t \times N_t$  and  $d \times d$ , and the scalars  $\alpha_t > 0$  and  $\beta_t > 0$  are precisions parameters (or inverse variances). The weight vectors  $\{\mathbf{w}_t\}_t$  can be integrated out analytically. The resulting log-marginal likelihood can then be maximized with respect to the network parameters  $\mathbf{z}$ , as well as the task precisions  $\{\alpha_t\}_t$  and  $\{\beta_t\}_t$ .




Figure 3: Comparison of the average regret and rank on the class of Forrester functions for 6 different methods. Solid lines display average over all the runs of given method, transparent areas depict 2 standard deviations.

The shared underlying neural network used in multi-head ABLR makes it a suitable for multi-task BO (Swersky et al., 2013; Springenberg et al., 2016; Perrone et al., 2018). Let  $\mathcal{D} = \{D_t : D_t = \{(\mathbf{x}_{tn}, y_{tn})\}_{n=1}^{N_t}\}_{t=1}^T$  be the data collected while optimizing the set of black-box functions  $\{f_t\}_{t=1}^T$ . The basis functions  $\phi_{\mathbf{z}}(\cdot)$  encode a shared representation based on  $\mathcal{D}$ , while the task-specific head model parameters are learned from the corresponding  $D_t$ . Hence, the optimization of a new task  $f_{T+1}$  can be warm-started by adding a new head model on top of the neural network (that was trained on  $\mathcal{D}$ ) and use it as a surrogate model to optimize:

$$\mathbf{x}_{T+1}^* = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmin}} f_{T+1}(\mathbf{x}), \quad (3)$$

A practical advantage of multi-head ABLR is that the complexity of learning the neural network is linear in the total number of evaluations, unlike multi-output GPs (Nguyen et al., 2014), while the complexity of learning the head parameters is linear in the number of task evaluations. These attractive computational properties are retained by our proposed method, ABRAC, as discussed in more detail in Section 4.

### 3 Nested Dropout for Ordered Representations

Nested dropout is a structured type of dropout originally proposed for autoencoders to learn ordered representations in their bottleneck layer (Rippel et al., 2014). It was also used in later works to learn spectral representations in convolutional neural networks (Rippel et al., 2015) and ordered word representations (Liu et al., 2016).

Assume there are  $d$  neurons in the bottleneck layer. Nested dropout is applied during training by deactivating these neurons according to a discrete distribution  $Q$  over the indices  $1, \dots, d$ . In each training step, one samples  $b \sim Q$  and drops the neurons  $b+1, \dots, d$ . This results in the activation of a different, but nested subset

---

#### Algorithm 1 ABRAC

---

**Input:** number of initial points  $n_0$ , budget  $N$ , feature net  $\phi_{\mathbf{z}}(\cdot) : \mathbb{R}^P \rightarrow \mathbb{R}^d$ , filter  $F^k$ , previous evaluations  $\{(x_{ti}, y_{ti})\}_{i=1}^{N_t}\}_{t=1}^T$ .

Fit  $\phi_{\mathbf{z}}(\cdot)$  using  $\{(x_{ti}, y_{ti})\}_{i=1}^{N_t}\}_{t=1}^T$  –Section 4.1.

Observe  $f_{\text{new}}$  at  $n_0$  randomly selected points  $x_1, x_2, \dots, x_{n_0} \in \mathcal{X}$ .

$\mathcal{C} = \{(x_i, y_i)\}_{i=1}^{n_0}$ , where  $y_i = f(x_i)$ .

Set  $n = n_0$ .

**while**  $n < N$  **do**

Fit probabilistic model  $g$  based on  $\mathcal{C}$  –Section 4.2.

$x_n = \operatorname{argmax}_{x \in \mathcal{X}} A_g(x)$ , where  $A$  is given acquisition function.

Observe  $y_n = f_{\text{new}}(x_n)$ .

Update  $\mathcal{C} = \mathcal{C} \cup \{(x_n, y_n)\}$ .

$n = n + 1$

**end while**

**Output:**  $\hat{x} = \operatorname{argmin}_{i=1,2,\dots,N} f_{\text{new}}(x_i)$

---

of indices at each training step for which, if  $b$  belongs to the subset, then  $1, \dots, b-1$  also belongs to it. Hence, every mini-batch gradient update during training can only touch the neurons in the bottleneck layer that were not dropped, which, as a result, leads to an importance ranking of the representation dimensions. Nested dropout comes with strong theoretical properties for linear autoencoders: the solution obtained when applying nested dropout to the linear autoencoder recovers principal component analysis under an additional orthonormality constraint when  $Q(b=j) > 0, \forall j \in 1, \dots, d$  (see Rippel et al., 2014, for details).

### 4 Multi-head ABLR with Adaptive Complexity (ABRAC)

A deficiency of multi-head ABLR when applied to multi-task BO is that the number of nonlinear basis functions is not adapted to the target task, where the number of observations is typically smaller than in the previous




Figure 4: Comparison of the surrogate objective along with its basis functions after 5 total evaluations obtained by 4 different methods. The target function is shown in blue, the evaluations are the red points (which includes the same 3 initial random points for both methods), and the surrogate predictive mean is shown in orange with two standard deviations in transparent light blue.

tasks, making it prone to overfitting. This issue is due to the direct use of conventional multi-task models in the context of sequential decision making problems such as BO. To solve this issue we propose to learn an ordered set of learned basis functions, which we selectively activate during the transfer learning procedure. Pseudocode is given in Algorithm 1.

To endow multi-head ABLR with ordered representations for transfer learning, we apply nested dropout during training to the last layer of the shared neural network  $\phi_{\mathbf{z}}(\mathbf{x})$ , which operates as the feature extractor. At each training step, we activate a nested subset of the output neurons according to a uniform distribution over the indices:  $Q(b = j) = 1/d, \forall j \in 1, \dots, d$ . Departing from Rippel et al. (2014), who proposed to sample the truncation index according to a geometric distribution, we observed empirically that the uniform distribution performed well without any further modifications, and with the additional benefit that we sample neurons with a high index more often and do not have to set additional hyperparameters.

Next, we introduce multi-head ABLR with adaptive complexity (ABRAC) for multi-task BO and detail the inference procedure. We propose a two-step approach to do transfer learning in the context of Bayesian sequential decision making consisting in (1) an off-line learning of ordered shared representations and (2) an online adaptation of the transfer learning procedure to improve the sample efficiency of the target BO task.

**Probabilistic model.** We make a minor modification of (2), which consists in allowing for a different precision parameter per basis function:

$$p(\mathbf{w}_t | \boldsymbol{\alpha}_t) = \text{Normal}(\mathbf{0}, \text{Diag}(\boldsymbol{\alpha}_t)^{-1}), \quad (4)$$

where  $\alpha_{ti} > 0$  is the precision associated to basis function  $i$ . Training consists in optimizing the log-marginal likelihood wrt the shared parameters  $\mathbf{z}$  and the head

model parameters  $\{\boldsymbol{\alpha}_t, \beta_t\}_t$ :

$$\begin{aligned} \sum_t \ln p(\mathbf{y}_t | \mathbf{z}, \boldsymbol{\alpha}_t, \beta_t) = & \quad (5) \\ \sum_t \ln \text{Normal}(\mathbf{0}, \beta_t^{-1} \mathbf{I}_{N_t} + \boldsymbol{\Phi}_t(\mathbf{z}) \text{Diag}(\boldsymbol{\alpha}_t)^{-1} \boldsymbol{\Phi}_t(\mathbf{z})^\top). \end{aligned}$$

Allowing for different  $\alpha_{ti}$  promotes sparsity by enabling us to automatically determine the relevance of every single component of  $\mathbf{w}_t$  (MacKay et al., 1994). The application of automatic relevance determination (ARD) to Bayesian linear regression was studied in (Tipping, 2001) and (Wipf and Nagarajan, 2008).

#### 4.1 Offline Ordered Representation Learning

To learn ordered representations suitable for transfer learning, we first train a neural network  $\phi_{\mathbf{z}}(\cdot)$  off-line. We minimize the mean square error between neural network predictions and the observations collected during all previous tasks using SGD with momentum (Sutskever et al., 2013). This is different from Perrone et al. (2018), where L-BFGS is used to learn  $\mathbf{z}$  jointly with  $\{\boldsymbol{\alpha}_t, \beta_t\}_t$ , and akin DNGO (Snoek et al., 2015). However, we apply nested dropout to the last layer during training as discussed next.

Before applying a mini-batch update, we first sample a random truncation  $b$  for each data in the mini-batch. Let  $\phi(\mathbf{x})_{b\downarrow}$  be a random truncation of the vector  $\phi(\mathbf{x})$  where the last  $d - b$  elements are masked with 0. For example,  $\phi(\mathbf{x})_{2\downarrow}^\top = [\phi_1(\mathbf{x}), \phi_2(\mathbf{x}), \phi_3(\mathbf{x}), \phi_4(\mathbf{x})]_{2\downarrow}^\top = [\phi_1(\mathbf{x}), \phi_2(\mathbf{x}), 0, 0]^\top$ . If  $b = d$ , then  $\phi(\mathbf{x})_{b\downarrow} = \phi(\mathbf{x})$  and  $(\cdot)_{b\downarrow}$  is the identity operator. Hence, we can view the off-line training procedure as maximizing a likelihood of which the terms correspond to a product of normal distributions:

$$p(y_{tn} | \mathbf{w}_t, \mathbf{z}, \beta_t) = \prod_{i=1}^d \text{Normal}(\phi_{\mathbf{z}}(\mathbf{x}_{tn})_{i\downarrow}^\top \mathbf{w}_t, \beta_t^{-1})^{\delta_{ibtn}}, \quad (6)$$






Figure 5: Comparison of the average regret and rank on the class of parameterized quadratics for 7 different methods. Solid lines display average over all the runs of given method, transparent areas depict 2 standard deviations. The most right plot displays comparison of the speed of ABRAC, different ABLR based methods and GP enriched by subsampling.

where  $\delta_{ij}$  is the Kronecker delta and the truncation  $b_{tn} \sim \text{Uniform}(1, \dots, d)$  is resampled at each iteration of the optimizer.

## 4.2 Online Adaptive Transfer Learning

When tackling a new BO task, we fix the feature extractor  $\phi_{\mathbf{z}}(\mathbf{x})$  and learn the additional head model parameters by maximizing the corresponding term in (5) after each addition of a new observation. Doing so dynamically adjusts the level of sparsity of Bayesian linear regression (BLR) via ARD. Maximizing the log-marginal likelihood of BLR has overall complexity  $\mathcal{O}(d^2 \max\{N_{\text{new}}, d\})$ , which is linear in number of evaluations of the new BO task. See Bishop (2006).

## 5 Experiments

We now present several experiments that show the benefits of our new proposed method, dubbed ABRAC. Firstly, we use the 1-D synthetic Forrester function, which allows us to visually quantify the model fit and uncertainty estimations. Afterwards, we present results on the parametrized quadratic functions by Perrone et al. (2018). Lastly, we benchmark ABRAC on real-world HPO problems containing feed-forward neural networks and support vector machines.

We implemented ABRAC and ABLR in the python framework Emukit (Paley et al., 2018) where the neural network code is based on MXNet (Chen et al., 2015). For L-BFGS, we used the SciPy (Jones et al., 01) implementation. For Gaussian Process (GP) based BO, we used a Matern52 to model the correlation between hyperparameters with automatic relevance determination, and optimize all hyperparameters by empirical Bayes (Rasmussen, 2003). All BO methods used expected improvement as acquisition function (Mockus et al., 1978).

We include BOHAMIANN (Springenberg et al., 2016)

as baseline, but, because of its significant computational overhead (approximately  $100\times$  slower than ABLR (Perrone et al., 2018)) we show its performance only on the first experiment. For all other experiments, we ran BOHAMIANN for 10 iterations and it did not outperform ABRAC in any of those.

To highlight the benefits of using an offline-trained feature extractor with hierarchical basis functions, we include the following two modifications of ABRAC without nested dropout: *ABLR SGD fixed* trains the features extractor by minimizing the squared error with SGD, which can be seen as an efficient multi-task version of DNGO (Snoek et al., 2015) and *ABLR L-BFGS fixed* which minimizes the negative log-likelihood with L-BFGS. Unlike *DNGO*, both versions do not require any contextual information. The main difference to standard *ABLR* is that the neural network  $\phi_{\mathbf{z}}$  is fixed for the new task. Furthermore, to analyse the benefit of using a neural network based feature extractor, we also include another baseline that uses random kitchen sinks as basis functions (ABLR RKS).

For each experiment, we average the performance across tasks in a leave-one-task-out fashion. Each optimization run was initialized with 3 random data points and we reported average performance of each method across 5 repetitions with a different seed. All experiments were conducted on AWS `ml.c4.2xlarge` instances.

**Forrester Function.** For the first experiment, we consider the parameterized class of Forrester functions  $f_t(x|a_t, b_t, c_t) = (a_t x - 2)^2 \sin(b_t x - 4) + c_t$  with the configuration space  $\mathcal{X} = [0, 1]$ . We create new tasks by sampling different coefficients  $a_t \sim \text{Normal}(6, 1)$ ,  $b_t \sim \text{Normal}(12, 4)$ ,  $c_t \sim \text{Uniform}(0, 10)$ . We generated 10 tasks with 20 uniformly distributed data points each. The architecture of the feature extractor  $\phi_{\mathbf{z}}(\mathbf{x})$  consists of a 2-layers fully connected neural network with 50 hidden units and tanh activation and  $d = 20$  output units with linear activation functions.




Figure 6: Comparison of the average rank on the NAS Benchmarks for 5 different methods. Solid lines display average over all the runs of given method, transparent areas depict 2 standard deviations.

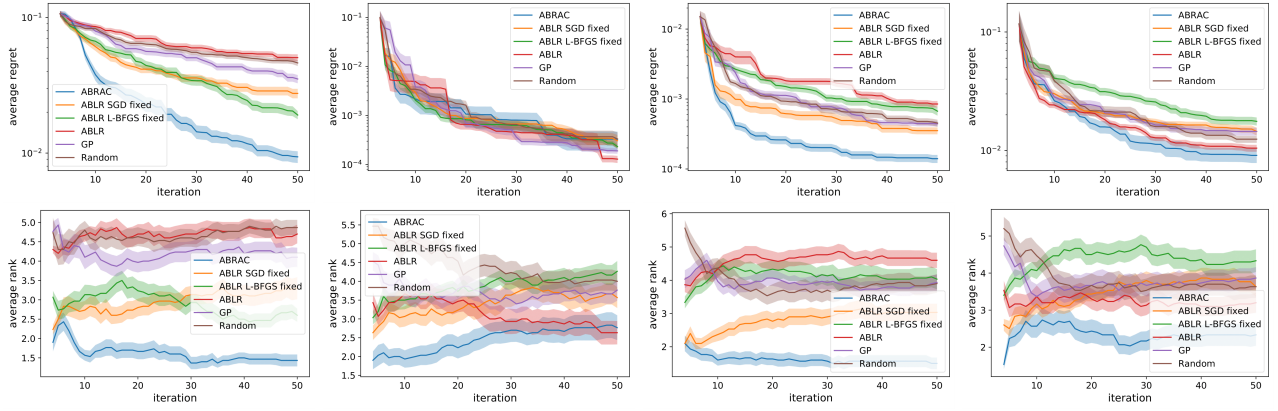


Figure 7: Comparison of the average regret and rank on the NAS Benchmarks for 5 different methods. Each column corresponds to one dataset (from left: **Protein Structure**, **Naval Propulsion**, **Slice Localization**, and **Parkinsons Telemonitoring**). Solid lines display average over all the runs of given method, transparent areas depict 2 standard deviations.

Firstly, we investigate the effects of the two major components of ABRAC: the offline training with nested dropout and the online transfer learning with adaptive basis function selection. To address the former we compare against a modified version of ABLR trained with SGD, dubbed *ABLR SGD adapt*, where we incorporate our adaptive procedure during the optimization. For the latter we include a version of ABRAC which considers all basis functions at every step (ABRAC fixed). Looking at Figure 2, one can see that not only in isolation, each of the two components improves over vanilla ABLR (with SGD training) but in combination, i.e ABRAC, we achieve the best performance.

In Figure 3 we compare against a range of baselines: Multi Task Gaussian process based Bayesian optimization (MT-GP-BO), which is a standard GP defined over the search space and the task-specific contextual information  $(a_t, b_t, c_t)$ , random search (RS), ABLR, BOHAMIANN as well as ABLR SGD fixed and ABLR L-BFGS fixed ascribed above. Similarly to MT-GP-BO, BOHAMIANN is provided with contextual information

$(a_t, b_t, c_t)$  for each task to enable multi-task learning. No other algorithm is given this information. One can see that ABRAC quickly adapts to the new task and achieves already after 5 function evaluations an acceptable regret. In contrast, GP-BO and ABLR require more samples to approach the global optimum. Interestingly, the baselines ABLR SGD fixed and ABLR L-BFGS fixed perform substantially worse, highlighting the importance to nested dropout to learn adaptive basis functions. To obtain a better intuition of our method, we show in Figure 4 the posterior mean and standard deviation of ABRAC and the ABLR variants after two optimization steps. All methods were initialized with the same three data points. Vanilla ABLR is not able to account for the reduced complexity of this low data regime and heavily overfits. On the other hand, ABLR with fixed basis functions explains all data points with observation noise. ABRAC is able to capture the right complexity and provides already in this early stage of optimization a meaningful model. Similar plots can be found in Figure 9 in the Appendix, where we show the same plots but with 5 fixed data




Figure 8: Comparison of the average AUC and rank on the OpenML-SVM Benchmarks for 6 different methods. Solid lines display average over all the runs of given method, transparent areas depict 2 standard deviations.

points among all the tasks. In addition, we also show the extracted basis functions with their relative weights in Figure 1.

### Parameterized Quadratic functions.

In the second experiment, we compare different methods on the parameterized quadratic functions of the form  $f_t(x|a_t, b_t, c_t) = a_t x^\top x + b_t \mathbf{1}^\top x + c_t$  from Perrone et al. (2018), where  $\mathbf{1}$  is an all-ones vector, and the coefficients are sampled uniformly  $a, b, c \sim \text{Uniform}(0.1, 1)$ . The configuration space is defined by  $\mathcal{X} = [-10, 10]^5$ . To train the feature extractor we generated 30 tasks with 100 random data points each.

Due to scalability issues we warm started MT-GP-BO with only 10 randomly selected datasets from each task, i.e. in total 290 data points. In order to distinguish between tasks, we provided MT-GP-BO with contextual information  $(a_t, b_t, c_t)$  for each task. Note that, none of the other methods is provided with contextual information. Figure 5 shows that ABRAC is able to outperform all other methods. The only method that can reach ABRAC’s performance after 45 iterations is ABLR SGD fixed.

In Figure 5 right, we report the optimization overhead of all neural network based methods. One can see that in all scenarios, ABRAC is not just the most efficient method in terms of sampling efficiency, but also in terms of computational overhead where it is more than  $100\times$  faster than ABLR or GPs.

**Tabular Benchmarks.** For the next experiment, we used the tabular benchmark by Klein and Hutter (2019) which simulates the discrete HPO of a 2-layer fully-connected neural network on 4 different regression datasets. Klein and Hutter (2019) performed an exhaustive search of all 62208 possible configurations in the search space per dataset, such that for a single function evaluation one can simply look up the performance in the table instead of training the neural network from scratch. For further details we refer to the original paper. For each task, we generate 1024 random config-

urations which are used for the offline feature learning. Since we have only access to 4 datasets, we increase the number of repetition with a different seed to 10. Since we have no access to contextual information we don’t include MT-GP-BO as a baseline but only GP-BO.

Figure 6 shows that in terms of rank, ABRAC outperforms all other method with significant margin, which do not improve much upon random search in the given budget. We include 4 additional graphs in Figure 7, where each of the figures displays average regret and rank for one of four datasets. One can see that for all the dataset, ABRAC is the best in terms of average rank and the same holds for the average regret except for *Naval Propulsion*. However, for this dataset random search already achieves a good performance and none of the methods significantly improves over it.

### SVM on OpenML datasets.

Lastly, we consider the HPO of support vector machines (SVM) with 2 continuous hyperparameters: the regularization parameter  $C$  (min: 0.000986, max: 998.492437) and the kernel parameter  $\gamma$  (min: 0.000988, max: 913.373845) on OpenML (Vanschoren et al., 2014) classification datasets. We chose the most evaluated flow (flow id 5891) with its 30 most evaluated datasets/tasks with almost one million evaluations in total.

We provide MT-GP-BO with contextual information: the number of features, the number of data points and the percentage of the majority class, and warm start it using 10 random points from each of the 29 related tasks. We compare these methods based on the average rank and average Area Under Curve (AUC). To model AUC on unseen configurations, we use the surrogate modeling approach from Eggensperger et al. (2015) based on a random forest model. Figure 8 confirms the superior behaviour of ABRAC to all other methods while the only method that can reach comparable performance in terms of AUC is ABLR, which considers all data in each step.



## 6 Conclusion

We introduced ABRAC, a probabilistic model for Bayesian optimization with adaptive complexity. It first learns a set of basis function based on neural networks on offline generated data with increasing complexity. During Bayesian optimization, we use Bayesian linear regression to obtain uncertainty estimates and thereby regularizes our basis functions. Compared to Gaussian processes our model scales linearly rather than cubically with the number of datapoints. In a battery of experiments, we show that ABRAC obtains much higher speed ups compared to other baseline methods, such as ABLR or Gaussian processes based Bayesian optimization.

## References

- A.Nayebi, Munteanu, A., and Poloczek, M. (2019). A framework for Bayesian optimization in embedded subspaces. In *Proceedings of the 36th International Conference on Machine Learning (ICML’19)*.
- Argyriou, A., Evgeniou, T., and Pontil, M. (2007). Multi-task feature learning. In *Advances in neural information processing systems*, pages 41–48.
- Bergstra, J. S., Bardenet, R., Bengio, Y., and Kégl, B. (2011). Algorithms for hyper-parameter optimization. In *Advances in neural information processing systems*, pages 2546–2554.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. springer.
- Chen, T., Li, M., Li, Y., Lin, M., Wang, N., Wang, M., Xiao, T., Xu, B., Zhang, C., and Zhang, Z. (2015). Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274*.
- Eggenesperger, K., Hutter, F., Hoos, H., and Leyton-Brown, K. (2015). Efficient benchmarking of hyperparameter optimizers via surrogates. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Feurer, M., Springenberg, J. T., and Hutter, F. (2015). Initializing bayesian hyperparameter optimization via meta-learning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Frazier, P., Powell, W., and Dayanik, S. (2009). The knowledge-gradient policy for correlated normal beliefs. *INFORMS journal on Computing*, 21(4):599–613.
- Frazier, P. I. (2018). A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*.
- Guo, X., Singh, S., Lee, H., Lewis, R. L., and Wang, X. (2014). Deep learning for real-time atari game play using offline monte-carlo tree search planning. In *Advances in neural information processing systems*, pages 3338–3346.
- Hennig, P. and Schuler, C. J. (2012). Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(Jun):1809–1837.
- Hoffman, M. D., Brochu, E., and de Freitas, N. (2011). Portfolio allocation for bayesian optimization. In *UAI*, pages 327–336. Citeseer.
- Hutter, F., Hoos, H. H., and Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization*, pages 507–523. Springer.
- Jones, D. R. (2001). A taxonomy of global optimization methods based on response surfaces. *Journal of global optimization*, 21(4):345–383.
- Jones, E., Oliphant, T., Peterson, P., et al. (2001–). SciPy: Open source scientific tools for Python.
- Klein, A. and Hutter, F. (2019). Tabular benchmarks for joint architecture and hyperparameter optimization. *arXiv:1905.04970 [cs.LG]*.
- Kushner, H. J. (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):97–106.
- Law, H. C., Zhao, P., Chang, L. S., Huang, J., and Sejdinovic, D. (2019). Hyperparameter learning via distributional transfer. In *Proceedings of the 32th International Conference on Advances in Neural Information Processing Systems (NIPS’19)*.
- Liu, A., Xing, C., Feng, Y., and Wang, D. (2016). Learning ordered word representations with  $\gamma$ -decay dropout. In *2016 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*, pages 1–5. IEEE.
- MacKay, D. J. et al. (1994). Bayesian nonlinear modeling for the prediction competition. *ASHRAE transactions*, 100(2):1053–1062.
- Mockus, J., Tiesis, V., and Zilinskas, A. (1978). The application of bayesian methods for seeking the extremum. *Towards global optimization*, 2(117-129):2.
- Moriconi, R., Deisenroth, M. P., and Kumar, K. S. S. (2019). High-dimensional Bayesian optimization using low-dimensional feature spaces. *arXiv:1902.10675 [stat.ML]*.
- Moss, H. B., Leslie, D. S., and Rayson, P. (2020). Mumbo: Multi-task max-value Bayesian optimization. *arXiv:2006.12093 [cs.LG]*.
- Nguyen, T. V., Bonilla, E. V., et al. (2014). Collaborative multi-output gaussian processes. In *UAI*, pages 643–652.

- Paleyes, A., Pullin, M., Mahsereci, M., Lawrence, N., and González, J. (2018). Emukit: Emulation and uncertainty quantification for decision making. <https://github.com/amzn/emukit>.
- Perrone, V., Jenatton, R., Seeger, M. W., and Archambeau, C. (2018). Scalable hyperparameter transfer learning. In *Advances in Neural Information Processing Systems*, pages 6845–6855.
- Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer.
- Rippel, O., Gelbart, M., and Adams, R. (2014). Learning ordered representations with nested dropout. In *International Conference on Machine Learning*, pages 1746–1754.
- Rippel, O., Snoek, J., and Adams, R. P. (2015). Spectral representations for convolutional neural networks. In *Advances in neural information processing systems*, pages 2449–2457.
- Ruder, S. (2017). An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Salinas, D., Shen, H., and Perrone, V. (2020). A quantile-based approach for hyperparameter transfer learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML’20)*.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. (2015). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175.
- Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959.
- Snoek, J., Rippel, O., Swersky, K., Kiros, R., Satish, N., Sundaram, N., Patwary, M., Prabhat, M., and Adams, R. (2015). Scalable bayesian optimization using deep neural networks. In *International conference on machine learning*, pages 2171–2180.
- Springenberg, J. T., Klein, A., Falkner, S., and Hutter, F. (2016). Bayesian optimization with robust bayesian neural networks. In *Advances in Neural Information Processing Systems*, pages 4134–4142.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. (2009). Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*.
- Sutskever, I., Martens, J., Dahl, G., and Hinton, G. (2013). On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147.
- Swersky, K., Snoek, J., and Adams, R. P. (2013). Multi-task bayesian optimization. In *Advances in neural information processing systems*, pages 2004–2012.
- Tipping, M. E. (2001). Sparse bayesian learning and the relevance vector machine. *Journal of machine learning research*, 1(Jun):211–244.
- Vanschoren, J., Van Rijn, J. N., Bischl, B., and Torgo, L. (2014). Openml: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60.
- Wipf, D. W. and Nagarajan, S. (2008). A new view of automatic relevance determination. In *Advances in neural information processing systems*.