
Scalable Bayesian Optimization Using Deep Neural Networks

Jasper Snoek^{*}
Oren Rippe^{†*}
Kevin Swersky[§]
Ryan Kiros[§]
Nadathur Satish[‡]
Narayanan Sundaram[‡]
Md. Mostofa Ali Patwary[‡]
Prabhat^{*}
Ryan P. Adams^{*}

JSNOEK@SEAS.HARVARD.EDU

RIPPEL@MATH.MIT.EDU

KSWERSKY@CS.TORONTO.EDU

RKIROS@CS.TORONTO.EDU

NADATHUR.RAJAGOPALAN.SATISH@INTEL.COM

NARAYANAN.SUNDARAM@INTEL.COM

MOSTOFA.ALI.PATWARY@INTEL.COM

PRABHAT@LBL.GOV

RPA@SEAS.HARVARD.EDU

^{*}Harvard University, School of Engineering and Applied Sciences

[†]Massachusetts Institute of Technology, Department of Mathematics

[§]University of Toronto, Department of Computer Science

[‡]Intel Labs, Parallel Computing Lab

^{*}NERSC, Lawrence Berkeley National Laboratory

Abstract

Bayesian optimization is an effective methodology for the global optimization of functions with expensive evaluations. It relies on querying a distribution over functions defined by a relatively cheap surrogate model. An accurate model for this distribution over functions is critical to the effectiveness of the approach, and is typically fit using Gaussian processes (GPs). However, since GPs scale cubically with the number of observations, it has been challenging to handle objectives whose optimization requires many evaluations, and as such, massively parallelizing the optimization.

In this work, we explore the use of neural networks as an alternative to GPs to model distributions over functions. We show that performing adaptive basis function regression with a neural network as the parametric form performs competitively with state-of-the-art GP-based approaches, but scales linearly with the number of data rather than cubically. This allows us to achieve a previously intractable degree of parallelism, which we apply to large scale hyperparameter optimization, rapidly finding competitive models on benchmark object recognition tasks using convolutional networks, and image caption generation using neural language models.

1. Introduction

Recently, the field of machine learning has seen unprecedented growth due to a new wealth of data, increases in computational power, new algorithms, and a plethora of exciting new applications. As researchers tackle more ambitious problems, the models they use are also becoming more sophisticated. However, the growing complexity of machine learning models inevitably comes with the introduction of additional hyperparameters. These range from design decisions such as the shape of a neural network architecture, to optimization parameters such as learning rates, to regularization hyperparameters such as weight decay. Proper setting of these hyperparameters is critical for performance on difficult problems.

There are many methods for optimizing over hyperparameter settings, ranging from simplistic procedures like grid or random search (Bergstra & Bengio, 2012), to more sophisticated model-based approaches using random forests (Hutter et al., 2011) or Gaussian processes (Snoek et al., 2012). Bayesian optimization is a natural framework for model-based global optimization of noisy, expensive black-box functions. It offers a principled approach to modeling uncertainty, which allows exploration and exploitation to be naturally balanced during the search. Perhaps the most commonly used model for Bayesian optimization is the Gaussian process (GP) due to its simplicity and flexibility in terms of conditioning and inference.

However, a major drawback of GP-based Bayesian optimization is that inference time grows cubically in the number of observations, as it necessitates the inversion of a dense covariance matrix. For problems with a very small

number of hyperparameters, this has not been an issue, as the minimum is often discovered before the cubic scaling renders further evaluations prohibitive. As the complexity of machine learning models grows, however, the size of the search space grows as well, along with the number of hyperparameter configurations that need to be evaluated before a solution of sufficient quality is found. Fortunately, as models have grown in complexity, computation has become significantly more accessible and it is now possible to train many models in parallel. A natural solution to the hyperparameter search problem is to therefore combine large-scale parallelism with a scalable Bayesian optimization method. The cubic scaling of the GP, however, has made it infeasible to pursue this approach.

The goal of this work is to develop a method for scaling Bayesian optimization, while still maintaining its desirable flexibility and characterization of uncertainty. To that end, we propose the use of neural networks to learn an adaptive set of basis functions for Bayesian linear regression. We refer to this approach as Deep Networks for Global Optimization (DNGO). Unlike a standard Gaussian process, DNGO scales linearly with the number of function evaluations—which, in the case of hyperparameter optimization, corresponds to the number of models trained—and is amenable to stochastic gradient training. Although it may seem that we are merely moving the problem of setting the hyperparameters of the model being tuned to setting them for the tuner itself, we show that for a suitable set of design choices it is possible to create a robust, scalable, and effective Bayesian optimization system that generalizes across many global optimization problems.

We demonstrate the effectiveness of DNGO on a number of difficult problems, including benchmark problems for Bayesian optimization, convolutional neural networks for object recognition, and multi-modal neural language models for image caption generation. We find hyperparameter settings that achieve competitive with state-of-the-art results of 6.37% and 27.4% on CIFAR-10 and CIFAR-100 respectively, and BLEU scores of 25.1 and 26.7 on the Microsoft COCO 2014 dataset using a single model and a 3-model ensemble.

2. Background and Related Work

2.1. Bayesian Optimization

Bayesian optimization is a well-established strategy for the global optimization of noisy, expensive black-box functions (Mockus et al., 1978). For an in-depth review, see Lizotte (2008), Brochu et al. (2010) and Osborne et al. (2009). Bayesian optimization relies on the construction of a probabilistic model that defines a distribution over objective functions from the input space to the objective

of interest. Conditioned on a prior over the functional form and a set of N observations of input-target pairs $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$, the relatively cheap posterior over functions is then queried to reason about where to seek the optimum of the expensive function of interest. The promise of a new experiment is quantified using an *acquisition function*, which, applied to the posterior mean and variance, expresses a trade-off between exploration and exploitation. Bayesian optimization proceeds by performing a proxy optimization over this acquisition function in order to determine the next input to evaluate.

Recent innovation has resulted in significant progress in Bayesian optimization, including elegant theoretical results (Srinivas et al., 2010; Bull, 2011; de Freitas et al., 2012), multitask and transfer optimization (Krause & Ong, 2011; Swersky et al., 2013; Bardenet et al., 2013) and the application to diverse tasks such as sensor set selection (Garnett et al., 2010), the tuning of adaptive Monte Carlo (Mahendran et al., 2012) and robotic gait control (Cassandra et al., 2014b).

Typically, GPs have been used to construct the distribution over functions used in Bayesian optimization, due to their flexibility, well-calibrated uncertainty, and analytic properties (Jones, 2001; Osborne et al., 2009). Recent work has sought to improve the performance of the GP approach through accommodating higher dimensional problems (Wang et al., 2013; Djolonga et al., 2013), input non-stationarities (Snoek et al., 2014) and initialization through meta-learning (Feurer et al., 2015). Random forests, which scale linearly with the data, have also been used successfully for algorithm configuration by Hutter et al. (2011) with empirical estimates of model uncertainty.

More specifically, Bayesian optimization seeks to solve the minimization problem

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

where we take \mathcal{X} to be a compact subset of \mathbb{R}^K . In our work, we build upon the standard GP-based approach of Jones (2001) which uses a GP surrogate and the *expected improvement* acquisition function (Mockus et al., 1978). For the surrogate model hyperparameters Θ , let $\sigma^2(\mathbf{x}; \Theta) = \Sigma(\mathbf{x}, \mathbf{x}; \Theta)$ be the marginal predictive variance of the probabilistic model, $\mu(\mathbf{x}; \mathcal{D}, \Theta)$ be the predictive mean, and define

$$\gamma(\mathbf{x}) = \frac{f(\mathbf{x}_{\text{best}}) - \mu(\mathbf{x}; \mathcal{D}, \Theta)}{\sigma(\mathbf{x}; \mathcal{D}, \Theta)}, \quad (2)$$

where $f(\mathbf{x}_{\text{best}})$ is the lowest observed value. The expected improvement criterion is defined as

$$a_{\text{EI}}(\mathbf{x}; \mathcal{D}, \Theta) = \sigma(\mathbf{x}; \mathcal{D}, \Theta) [\gamma(\mathbf{x}) \Phi(\gamma(\mathbf{x})) + \mathcal{N}(\gamma(\mathbf{x}); 0, 1)]. \quad (3)$$

Here $\Phi(\cdot)$ is the cumulative distribution function of a standard normal, and $\mathcal{N}(\cdot; 0, 1)$ is the density of a standard normal. Note that numerous alternate acquisition functions and combinations thereof have been proposed (Kushner, 1964; Srinivas et al., 2010; Hoffman et al., 2011), which could be used without affecting the analytic properties of our approach.

2.2. Bayesian Neural Networks

The application of Bayesian methods to neural networks has a rich history in machine learning (MacKay, 1992; Hinton & van Camp, 1993; Buntine & Weigend, 1991; Neal, 1995; De Freitas, 2003). The goal of Bayesian neural networks is to uncover the full posterior distribution over the network weights in order to capture uncertainty, to act as a regularizer, and to provide a framework for model comparison. The full posterior is, however, intractable for most forms of neural networks, necessitating expensive approximate inference or Markov chain Monte Carlo simulation. More recently, full or approximate Bayesian inference has been considered for small pieces of the overall architecture. For example, in similar spirit to this work, Lázaro-Gredilla & Figueiras-Vidal (2010); Hinton & Salakhutdinov (2008) and Calandra et al. (2014a) considered inference over just the last layer of a neural network. Alternatively, variational approaches are developed in Kingma & Welling (2014); Rezende et al. (2014) and Mnih & Gregor (2014), where a neural network is used in a variational approximation to the posterior distribution over the latent variables of a directed generative neural network.

3. Adaptive Basis Regression with Deep Neural Networks

A key limitation of GP-based Bayesian optimization is that the computational cost of the technique scales cubically in the number of observations, limiting the applicability of the approach to objectives that require a relatively small number of observations to optimize. In this work, we aim to replace the GP traditionally used in Bayesian optimization with a model that scales in a less dramatic fashion, but retains most of the GP’s desirable properties such as flexibility and well-calibrated uncertainty. Bayesian neural networks are a natural consideration, not least because of the theoretical relationship between Gaussian processes and infinite Bayesian neural networks (Neal, 1995; Williams, 1996). However, deploying these at a large scale is very computationally expensive.

As such, we take a pragmatic approach and add a Bayesian linear regressor to the last hidden layer of a deep neural network, marginalizing only the output weights of the net while using a point estimate for the remaining parameters. This results in *adaptive basis regression*, a well-established

statistical technique which scales linearly in the number of observations, and cubically in the basis function dimensionality. This allows us to explicitly trade off evaluation time and model capacity. As such, we form the basis using the very flexible and powerful non-linear functions defined by the neural network.

First of all, without loss of generality and assuming compact support for each input dimension, we scale the input space to the unit hypercube. We denote by $\phi(\cdot) = [\phi_1(\cdot), \dots, \phi_D(\cdot)]^\top$ the vector of outputs from the last hidden layer of the network, trained on inputs and targets $\mathcal{D} := \{(\mathbf{x}_n, y_n)\}_{n=1}^N \subset \mathbb{R}^K \times \mathbb{R}$. We take these to be our set of basis functions. In addition, define Φ to be the design matrix arising from the data and this basis, where $\Phi_{nd} = \phi_d(\mathbf{x}_n)$ is the output design matrix, and \mathbf{y} the stacked target vector.

These basis functions are parameterized via the weights and biases of the deep neural network, and these parameters are trained via backpropagation and stochastic gradient descent with momentum. In this training phase, a linear output layer is also fit. This procedure can be viewed as a maximum *a posteriori* (MAP) estimate of all parameters in the network. Once this “basis function neural network” has been trained, we replace the MAP-parameterized output layer with a Bayesian linear regressor that captures uncertainty in the weights. See Section 3.1.2 for a more elaborate explanation of this choice.

The predictive mean $\mu(\mathbf{x}; \Theta)$ and variance $\sigma^2(\mathbf{x}; \Theta)$ of the model are then given by (see Bishop, 2006)

$$\mu(\mathbf{x}; \mathcal{D}, \Theta) = \mathbf{m}^\top \phi(\mathbf{x}) + \eta(\mathbf{x}), \quad (4)$$

$$\sigma^2(\mathbf{x}; \mathcal{D}, \Theta) = \phi(\mathbf{x})^\top \mathbf{K}^{-1} \phi(\mathbf{x}) + \frac{1}{\beta} \quad (5)$$

where

$$\mathbf{m} = \beta \mathbf{K}^{-1} \Phi^\top \tilde{\mathbf{y}} \in \mathbb{R}^D \quad (6)$$

$$\mathbf{K} = \beta \Phi^\top \Phi + \mathbf{I} \alpha \in \mathbb{R}^{D \times D}. \quad (7)$$

Here, $\eta(\mathbf{x})$ is a prior mean function which is described in Section 3.1.3, and $\tilde{\mathbf{y}} = \mathbf{y} - \eta(\mathbf{x})$. In addition, $\alpha, \beta \in \Theta$ are regression model hyperparameters. We integrate out α and β using slice sampling (Neal, 2000) according to the methodology of Snoek et al. (2012) over the marginal likelihood, which is given by

$$\log p(\mathbf{y} | \mathbf{X}, \alpha, \beta) = \frac{D}{2} \log \alpha + \frac{N}{2} \log \beta - \frac{N}{2} \log(2\pi) - \frac{\beta}{2} \|\tilde{\mathbf{y}} - \Phi \mathbf{m}\|^2 - \frac{\alpha}{2} \mathbf{m}^\top \mathbf{m} - \frac{1}{2} \log |\mathbf{K}|. \quad (8)$$

It is clear that the computational bottleneck of this procedure is the inversion of \mathbf{K} . However, note that the size of




Figure 1. The time per suggested experiment for our method compared to the state-of-the-art GP based approach from Snoek et al. (2014) on the six dimensional Hartmann function. We ran each algorithm on the same 32 core system with 80GB of RAM five times and plot the mean and standard deviation.

this matrix grows with the output dimensionality D , rather than the number of observations N as in the GP case. This allows us to scale to significantly more observations than with the GP as demonstrated in Figure 1.

3.1. Model details

3.1.1. NETWORK ARCHITECTURE

A natural concern with the use of deep networks is that they often require significant effort to tune and tailor to specific problems. One can consider adjusting the architecture and tuning the hyperparameters of the neural network as itself a difficult hyperparameter optimization problem. An additional challenge is that we aim to create an approach that generalizes across optimization problems. We found that design decisions such as the type of activation function used significantly altered the performance of the Bayesian optimization routine. For example, in Figure 2 we see that the commonly used rectified linear (ReLU) function can lead to very poor estimates of uncertainty, which causes the Bayesian optimization routine to explore excessively. Since the bounded tanh function results in smooth functions with realistic variance, we use this nonlinearity in this work; however, if the smoothness assumption needs to be relaxed, a combination of rectified linear functions with a tanh function only on the last layer can also be used in order to bound the basis.

In order to tune any remaining hyperparameters, such as the width of the hidden layers and the amount of regularization, we used GP-based Bayesian optimization. For each of one to four layers we ran Bayesian optimization using

the Spearmint (Snoek et al., 2014) package to minimize the average relative loss on a series of benchmark global optimization problems. We tuned a global learning rate, momentum, layer sizes, ℓ_2 normalization penalties for each set of weights and dropout rates (Hinton et al., 2012) for each layer. Interestingly, the optimal configuration featured no dropout and very modest ℓ_2 normalization. We suspect that dropout, despite having an approximate correction term, causes noise in the predicted mean resulting in a loss of precision. The optimizer instead preferred to restrict capacity via a small number of hidden units. Namely, the optimal architecture is a deep and narrow network with 3 hidden layers and approximately 50 hidden units per layer. We use the same architecture throughout our empirical evaluation, and this architecture is illustrated in Figure 2(d).

3.1.2. MARGINAL LIKELIHOOD VS MAP ESTIMATE

The standard empirical Bayesian approach to adaptive basis regression is to maximize the marginal likelihood with respect to the parameters of the basis (see Equation 8), thus taking the model uncertainty into account. However, in the context of our method, this requires evaluating the gradient of the marginal likelihood, which requires inverting a $D \times D$ matrix on each update of stochastic gradient descent. As this makes the optimization of the net significantly slower, we take a pragmatic approach and optimize the basis using a point estimate and apply the Bayesian linear regression layer *post-hoc*. We found that both approaches gave qualitatively and empirically similar results, and as such we in practice employ the more efficient one.

3.1.3. QUADRATIC PRIOR

One of the advantages of Bayesian optimization is that it provides natural avenues for incorporating prior information about the objective function and search space. For example, when choosing the boundaries of the search space, a typical assumption has been that the optimal solution lies somewhere in the interior of the input space. However, by the curse of dimensionality, most of the volume of the space lies very close to its boundaries. Therefore, we select a mean function $\eta(\mathbf{x})$ (see Equation 4) to reflect our subjective prior beliefs that the function is coarsely approximated by a convex quadratic function centered in the bounded search region, i.e.,

$$\eta(\mathbf{x}) = \lambda + (\mathbf{x} - \mathbf{c})^T \mathbf{\Lambda} (\mathbf{x} - \mathbf{c}) \quad (9)$$

where \mathbf{c} is the center of the quadratic, λ is an offset and $\mathbf{\Lambda}$ a diagonal scaling matrix. We place a Gaussian prior with mean 0.5 (the center of the unit hypercube) on \mathbf{c} , horseshoe (Carvalho et al., 2009) priors on the diagonal elements $\Lambda_{kk} \forall k \in \{1, \dots, K\}$ and integrate out b , λ and \mathbf{c} using slice sampling over the marginal likelihood.




Figure 2. A comparison of the predictive mean and uncertainty learned by the model when using 2(a) only tanh, 2(c) only rectified linear (ReLU) activation functions or 2(b) ReLU’s but a tanh on the last hidden layer. The shaded regions correspond to standard deviation envelopes around the mean. The choice of activation function significantly modifies the basis functions learned by the model. Although the ReLU, which is the standard for deep neural networks, is highly flexible, we found that its unbounded activation can lead to extremely large uncertainty estimates. Subfigure 2(d) illustrates the overall architecture of the DNGO model. Dashed lines correspond to weights that are marginalized.

The horseshoe is a so-called *one-group* prior for inducing sparsity and is a somewhat unusual choice for the weights of a regression model. Here we choose it because it 1) has support only on the positive reals, leading to convex functions, and 2) it has a large spike at zero with a heavy tail, resulting in strong shrinkage for small values while preserving large ones. This last effect is important for handling model misspecification as it allows the quadratic effect to disappear and become a simple offset if necessary.

3.2. Incorporating input space constraints

Many problems of interest have complex, possibly unknown bounds, or exhibit undefined behavior in some regions of the input space. These regions can be characterized as *constraints* on the search space. Recent work (Gelbart et al., 2014; Snoek, 2013; Gramacy & Lee, 2010) has developed approaches for modeling unknown constraints in GP-based Bayesian optimization by learning a constraint classifier and then discounting expected improvement by the probability of constraint violation.

More specifically, define $c_n \in \{0, 1\}$ to be a binary indicator of the validity of input \mathbf{x}_n . Also, denote the sets of valid and invalid inputs as $\mathcal{V} = \{(\mathbf{x}_n, y_n) \mid c_n = 1\}$ and $\mathcal{I} = \{(\mathbf{x}_n, y_n) \mid c_n = 0\}$, respectively. Note that $\mathcal{D} := \mathcal{V} \cup \mathcal{I}$. Lastly, let Ψ be the collection of constraint hyperparameters. The modified expected improvement function can be written as

$$a_{\text{CEI}}(\mathbf{x}; \mathcal{D}, \Theta, \Psi) = a_{\text{EI}}(\mathbf{x}; \mathcal{V}, \Theta) \mathbb{P}[c = 1 \mid \mathbf{x}, \mathcal{D}, \Psi] .$$

In this work, to model the constraint surface, we similarly replace the Gaussian process with the adaptive basis model,

integrating out the output layer weights:

$$\mathbb{P}[c = 1 \mid \mathbf{x}, \mathcal{D}, \Psi] = \int_{\mathbf{w}} \mathbb{P}[c = 1 \mid \mathbf{x}, \mathcal{D}, \mathbf{w}, \Psi] P(\mathbf{w}; \Psi) d\mathbf{w} . \quad (10)$$

In this case, we use a Laplace approximation to the posterior. For noisy constraints we perform Bayesian logistic regression, using a logistic likelihood function for $\mathbb{P}[c = 1 \mid \mathbf{x}, \mathcal{D}, \mathbf{w}, \Psi]$. For noiseless constraints, we replace the logistic function with a step function.

3.3. Parallel Bayesian Optimization

Obtaining a closed form expression for the joint acquisition function across multiple inputs is intractable in general (Ginsbourger & Riche, 2010). However, a successful Monte Carlo strategy for parallelizing Bayesian optimization was developed in Snoek et al. (2012). The idea is to marginalize over the possible outcomes of currently running experiments when making a decision about a new experiment to run. Following this strategy, we use the posterior predictive distribution given by Equations 4 and 5 to generate a set of fantasy outcomes for each running experiment which we then use to augment the existing dataset. By averaging over sets of fantasies, we can perform approximate marginalization when computing EI for a candidate point. We note that this same idea works with the constraint network, where instead of computing marginalized EI, we would compute the marginalized probability of violating a constraint.

To that end, given currently running jobs with inputs $\{\mathbf{x}_j\}_{j=1}^J$, the marginalized acquisition function

Experiment	# Evals	SMAC	TPE	Spearmint	DNGO
Branin (0.398)	200	0.655 ± 0.27	0.526 ± 0.13	0.398 ± 0.00	0.398 ± 0.00
Hartmann6 (-3.322)	200	-2.977 ± 0.11	-2.823 ± 0.18	-3.3166 ± 0.02	-3.319 ± 0.00
Logistic Regression	100	8.6 ± 0.9	8.2 ± 0.6	6.88 ± 0.0	6.89 ± 0.04
LDA (On grid)	50	1269.6 ± 2.9	1271.5 ± 3.5	1266.2 ± 0.1	1266.2 ± 0.0
SVM (On grid)	100	24.1 ± 0.1	24.2 ± 0.0	24.1 ± 0.1	24.1 ± 0.1

Table 1. Evaluation of DNGO on global optimization benchmark problems versus scalable (TPE, SMAC) and non-scalable (Spearmint) Bayesian optimization methods. All problems are minimization problems. For each problem, each method was run 10 times to produce error bars.

$a_{\text{MCEI}}(\cdot; \mathcal{D}, \Theta, \Psi)$ is given by

$$a_{\text{MCEI}}(\mathbf{x}; \mathcal{D}, \{\mathbf{x}_j\}_{j=1}^J, \Theta, \Psi) = \int a_{\text{CEI}}(\mathbf{x}; \mathcal{D} \cup \{(\mathbf{x}_j, y_j)\}_{j=1}^J, \Theta, \Psi) \times \mathbb{P}[\{c_j, y_j\}_{j=1}^J \mid \mathcal{D}, \{\mathbf{x}_j\}_{j=1}^J] dy_1 \dots dy_n dc_1 \dots dc_n.$$

When this strategy is applied to a GP, the cost of computing EI for a candidate point becomes cubic in the size of the augmented dataset. This restricts both the number of running experiments that can be tolerated, as well as the number of fantasy sets used for marginalization. With DNGO it is possible to scale both of these up to accommodate a much higher degree of parallelism.

Finally, following the approach of Snoek et al. (2012) we integrate out the hyperparameters of the model to obtain our final integrated acquisition function. For each iteration of the optimization routine we pick the next input, \mathbf{x}^* , to evaluate according to

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmax}} a_{\text{MCEI}}(\mathbf{x}; \mathcal{D}, \{\mathbf{x}_j\}_{j=1}^J), \quad (11)$$

where

$$a_{\text{MCEI}}(\mathbf{x}; \mathcal{D}, \{\mathbf{x}_j\}_{j=1}^J) = \int a_{\text{MCEI}}(\mathbf{x}; \mathcal{D}, \{\mathbf{x}_j\}_{j=1}^J, \Theta, \Psi) d\Theta d\Psi. \quad (12)$$

4. Experiments

4.1. HPOLib Benchmarks

In the literature, there exist several other methods for model-based optimization. Among these, the most popular variants in machine learning are the random forest-based SMAC procedure (Hutter et al., 2011) and the tree Parzen estimator (TPE) (Bergstra et al., 2011). These are faster to fit than a Gaussian process and scale more gracefully with large datasets, but this comes at the cost of a more heuristic treatment of uncertainty. By contrast, DNGO provides a balance between scalability and the Bayesian marginalization of model parameters and hyperparameters.

Method	Test BLEU
Human Expert LBL	24.3
Regularized LSTM	24.3
Soft-Attention LSTM	24.3
10 LSTM ensemble	24.4
Hard-Attention LSTM	25.0
Single LBL	25.1
2 LBL ensemble	25.9
3 LBL ensemble	26.7

Table 2. Image caption generation results using BLEU-4 on the Microsoft COCO 2014 test set. Regularized and ensembled LSTM results are reported in Zaremba et al. (2015). The baseline LBL tuned by a human expert and the Soft and Hard Attention models are reported in Xu et al. (2015). We see that ensembling our top models resulting from the optimization further improves results significantly. We noticed that there were distinct multiple local optima in the hyperparameter space, which may explain the dramatic improvement from ensembling a small subset of models.

To demonstrate the effectiveness of our approach, we compare DNGO to these scalable model-based optimization variants, as well as the input-warped Gaussian process method of Snoek et al. (2014) on the benchmark set of continuous problems from the HPOLib package (Eggenberger et al., 2013). As Table 1 shows, DNGO significantly outperforms SMAC and TPE, and is competitive with the Gaussian process approach. This shows that, despite vast improvements in scalability, DNGO retains the statistical efficiency of the Gaussian process method in terms of the number of evaluations required to find the minimum.

4.2. Image Caption Generation

In this experiment, we explore the effectiveness of DNGO on a practical and expensive problem where highly parallel evaluation is necessary to make progress in a reasonable amount of time. We consider the task of image caption generation using multi-modal neural language models. Specifically, we optimize the hyperparameters of the log-bilinear model (LBL) from Kiros et al. (2014) to maximize



(a) “A person riding a wave in the ocean.”



(b) “A bird sitting on top of a field.”



(c) “A horse is riding a horse.”

Figure 3. Sample test images and generated captions from the best LBL model on the COCO 2014 dataset. The first two captions sensibly describe the contents of their respective images, while the third is offensively inaccurate.

the BLEU score of a validation set from the recently released COCO dataset (Lin et al., 2014). In our experiments, each evaluation of this model took an average of 26.6 hours.

We optimize learning parameters such as learning rate, momentum and batch size; regularization parameters like dropout and weight decay for word and image representations; and architectural parameters such as the context size, whether to use the additive or multiplicative version, the size of the word embeddings and the multi-modal representation size¹. The final parameter is the number of factors, which is only relevant for the multiplicative model. This adds an interesting challenge, since it is only relevant for half of the hyperparameter space. This gives a total of 11 hyperparameters. Even though this number seems small, this problem offers a number of challenges which render its optimization quite difficult. For example, in order to not lose any generality, we choose broad box constraints for the hyperparameters; this, however, renders most of the volume of the model space infeasible. In addition, quite a few of the hyperparameters are categorical, which introduces severe non-stationarities in the objective surface.

Nevertheless, one of the advantages of a scalable method is the ability to highly parallelize hyperparameter optimization. In this way, high quality settings can be found after only a few sequential steps. To test DNGO in this scenario, we optimize the log-bilinear model with up to 800 parallel evaluations.

Running between 300 and 800 experiments in parallel (determined by cluster availability), we proposed and evaluated approximately 2500 experiments—the equivalent of over 2700 CPU days—in less than one week. Using the BLEU-4 metric, we optimized the validation set performance and the best LBL model found by DNGO outperforms recently proposed models using LSTM recurrent neural networks (Zaremba et al., 2015; Xu et al., 2015) on

¹Details are provided in the supplementary material.

Layer type	# Filters	Window	Stride
Convolution	96	3×3	
Convolution	96	3×3	
Max pooling		3×3	2
Convolution	192	3×3	
Convolution	192	3×3	
Convolution	192	3×3	
Max pooling		3×3	2
Convolution	192	3×3	
Convolution	192	1×1	
Convolution	10/100	1×1	
Global averaging		6×6	
Softmax			

Table 3. Our convolutional neural network architecture. This choice was chosen to be maximally generic. Each convolution layer is followed by a ReLU nonlinearity.

the test set. This is remarkable, as the LBL is a relatively simple approach. Ensembling this top model with the second and third best (under the validation metric) LBL models resulted in a test-set BLEU score² of 26.7, significantly outperforming the LSTM-based approaches. We noticed that there were distinct multiple local optima in the hyperparameter space, which may explain the dramatic improvement from ensembling a small number of models. We show qualitative examples of generated captions on test images in Figure 3. Further figures showing the BLEU score as a function of the iteration of Bayesian optimization are provided in the supplementary.

²We have verified that our BLEU score evaluation is consistent across reported results. We used a beam search decoding for our test predictions with the LBL model.

Method	CIFAR-10	CIFAR-100
Maxout	9.38%	38.57%
DropConnect	9.32%	N/A
Network in network	8.81%	35.68%
Deeply supervised	7.97%	34.57%
ALL-CNN	7.25%	33.71%
Tuned CNN	6.37%	27.4%

Table 4. We use our algorithm to optimize validation set error as a function of various hyperparameters of a convolutional neural network. We report the test errors of the models with the optimal hyperparameter configurations, as compared to current state-of-the-art results.

4.3. Deep Convolutional Neural Networks

Finally, we use DNGO on a pair of highly competitive deep learning visual object recognition benchmark problems. We tune the hyperparameters of a deep convolutional neural network on the CIFAR-10 and CIFAR-100 datasets. Our approach is to establish a single, generic architecture, and specialize it to various tasks via individualized hyperparameter tuning. As such, for both datasets, we employed the same generic architecture inspired by the configuration proposed in [Springenberg et al. \(2014\)](#), which was shown to attain strong classification results. This architecture is detailed in Table 3.

For this architecture, we tuned the momentum, learning rate, ℓ_2 weight decay coefficients, dropout rates, standard deviations of the random i.i.d. Gaussian weight initializations, and corruption bounds for various data augmentations: global perturbations of hue, saturation and value, random scalings, input pixel dropout and random horizontal reflections. We optimized these over a validation set of 10,000 examples drawn from the training set, running each network for 200 epochs. See Figure 4 for a visualization of the hyperparameter tuning procedure.

We performed the optimization on a cluster of Intel® Xeon Phi™ coprocessors, with 40 jobs running in parallel using a kernel library that has been highly optimized for efficient computation on the Intel® Xeon Phi™ coprocessor³. For the optimal hyperparameter configuration found, we ran a final experiment for 350 epochs on the entire training set, and report its result.

Our optimal models for CIFAR-10 and CIFAR-100 achieved test errors of 6.37% and 27.4% respectively. A comparison to published state-of-the-art results ([Goodfellow et al., 2013](#); [Wan et al., 2013](#); [Lin et al., 2013](#); [Lee et al., 2014](#); [Springenberg et al., 2014](#)) is given in Table 4. We see

³Available at <https://github.com/orippel/micmat>




Figure 4. Validation errors on CIFAR-100 corresponding to different hyperparameter configurations as evaluated over time. These are represented as a planar histogram, where the shade of each bin indicates the total count within it. The current best validation error discovered is traced in black. This projection demonstrates the exploration-versus-exploitation paradigm of Bayesian optimization, in which the algorithm trades off visiting unexplored parts of the space, and focusing on parts which show promise.

that the parallelized automated hyperparameter tuning procedure obtains models that are highly competitive with the state-of-the-art in just a few sequential steps.

A comprehensive overview of the setup, the architecture, the tuning and the optimum configuration can be found in the supplementary material.

5. Conclusion

In this paper, we introduced deep networks for global optimization, or DNGO, which enables efficient optimization of noisy, expensive black-box functions. While this model maintains desirable properties of the GP such as tractability and principled management of uncertainty, it greatly improves its scalability from cubic to linear as a function of the number of observations. We demonstrate that while this model allows efficient computation, its performance is nevertheless competitive with existing state-of-the-art approaches for Bayesian optimization. We demonstrate empirically that it is especially well suited to massively parallel hyperparameter optimization.

While adaptive basis regression with neural networks provides one approach to the enhancement of scalability, other models may also present promise. One promising line of work, for example by [Nickson et al. \(2014\)](#), is to introduce a similar methodology by instead employing the sparse Gaussian process as the underlying probabilistic model ([Snelson & Ghahramani, 2005](#); [Titsias, 2009](#); [Hensman et al., 2013](#)).

6. Acknowledgements

This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This work used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We would like to acknowledge the NERSC systems staff, in particular Helen He and Harvey Wasserman, for providing us with generous access to the Babbage Xeon Phi testbed.

The image caption generation computations in this paper were run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University. We would like to acknowledge the FASRC staff and in particular James Cuff for providing generous access to Odyssey.

Jasper Snoek is a fellow in the Harvard Center for Research on Computation and Society. Kevin Swersky is the recipient of an Ontario Graduate Scholarship (OGS). This work was partially funded by NSF IIS-1421780, the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canadian Institute for Advanced Research (CIFAR).

References

- Bardenet, R., Brendel, M., Kégl, B., and Sebag, M. Collaborative hyperparameter tuning. In *ICML*, 2013.
- Bergstra, J. and Bengio, Y. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305, 2012.
- Bergstra, J. S., Bardenet, R., Bengio, Y., and Kégl, B. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*. 2011.
- Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer-Verlag New York, Inc., 2006.
- Brochu, E., Brochu, T., and de Freitas, N. A Bayesian interactive optimization approach to procedural animation design. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2010.
- Bull, A. D. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, (3-4):2879–2904, 2011.
- Buntine, W. L. and Weigend, A. S. Bayesian back-propagation. *Complex systems*, 5(6):603–643, 1991.
- Calandra, R., Peters, J., Rasmussen, C. E., and Deisenroth, M. P. Manifold Gaussian processes for regression. *preprint arXiv:1402.5876*, 2014a.
- Calandra, R., Peters, J., Seyfarth, A., and Deisenroth, M. P. An experimental evaluation of Bayesian optimization on bipedal locomotion. In *International Conference on Robotics and Automation*, 2014b.
- Carvalho, C. M., Polson, N. G., and Scott, J. G. Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics*, 2009.
- De Freitas, J. F. *Bayesian methods for neural networks*. PhD thesis, Trinity College, University of Cambridge, 2003.
- de Freitas, N., Smola, A. J., and Zoghi, M. Exponential regret bounds for Gaussian process bandits with deterministic observations. In *ICML*, 2012.
- Djolonga, J., Krause, A., and Cevher, V. High dimensional Gaussian process bandits. In *Advances in Neural Information Processing Systems*, 2013.
- Eggensperger, K., Feurer, M., Hutter, F., Bergstra, J., Snoek, J., Hoos, H., and Leyton-Brown, K. Towards an empirical foundation for assessing Bayesian optimization of hyperparameters. In *NIPS Workshop on Bayesian Optimization in Theory and Practice*, 2013.
- Feurer, M., Springenberg, T., and Hutter, F. Initializing Bayesian hyperparameter optimization via meta-learning. In *AAAI Conference on Artificial Intelligence*, 2015.
- Garnett, R., Osborne, M. A., and Roberts, S. J. Bayesian optimization for sensor set selection. In *International Conference on Information Processing in Sensor Networks*, 2010.
- Gelbart, M. A., Snoek, J., and Adams, R. P. Bayesian optimization with unknown constraints. In *Uncertainty in Artificial Intelligence*, 2014.
- Ginsbourger, D. and Riche, R. L. Dealing with asynchronicity in parallel Gaussian process based global optimization. <http://hal.archives-ouvertes.fr/hal-00507632>, 2010.
- Goodfellow, I. J., Warde-Farley, D., Mirza, M., Courville, A. C., and Bengio, Y. Maxout networks. In *ICML*, 2013.
- Gramacy, R. B. and Lee, H. K. H. Optimization under unknown constraints, 2010. arXiv:1004.4027.
- Hensman, J., Fusi, N., and Lawrence, N. Gaussian processes for big data. In *Uncertainty in Artificial Intelligence*, 2013.
- Hinton, G. E. and van Camp, D. Keeping neural networks simple by minimizing the description length of the weights. In *ACM Conference on Computational Learning Theory*, 1993.
- Hinton, G. E. and Salakhutdinov, R. Using deep belief nets to learn covariance kernels for Gaussian processes. In *Advances in neural information processing systems*, pp. 1249–1256, 2008.
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- Hoffman, M., Brochu, E., and de Freitas, N. Portfolio allocation for Bayesian optimization. In *Uncertainty in Artificial Intelligence*, 2011.
- Hutter, F., Hoos, H. H., and Leyton-Brown, K. Sequential model-based optimization for general algorithm configuration. In *Learning and Intelligent Optimization 5*, 2011.
- Jones, D. R. A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization*, 21, 2001.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.

- Kiros, R., Salakhutdinov, R., and Zemel, R. S. Multimodal neural language models. In *ICML*, 2014.
- Krause, A. and Ong, C. S. Contextual Gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, 2011.
- Kushner, H. J. A new method for locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86, 1964.
- Lázaro-Gredilla, M. and Figueiras-Vidal, A. R. Marginalized neural network mixtures for large-scale regression. *Neural Networks, IEEE Transactions on*, 21(8):1345–1351, 2010.
- Lee, C.-Y., Xie, S., Gallagher, P., Zhang, Z., and Tu, Z. Deeply supervised nets. In *Deep Learning and Representation Learning Workshop, NIPS*, 2014.
- Lin, M., Chen, Q., and Yan, S. Network in network. *CoRR*, abs/1312.4400, 2013. URL <http://arxiv.org/abs/1312.4400>.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft COCO: Common objects in context. In *ECCV 2014*, pp. 740–755. Springer, 2014.
- Lizotte, D. *Practical Bayesian Optimization*. PhD thesis, University of Alberta, Edmonton, Alberta, 2008.
- MacKay, D. J. A practical Bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992.
- Mahendran, N., Wang, Z., Hamze, F., and de Freitas, N. Adaptive MCMC with Bayesian optimization. In *Artificial Intelligence and Statistics*, 2012.
- Mnih, A. and Gregor, K. Neural variational inference and learning in belief networks. In *ICML*, 2014.
- Mockus, J., Tiesis, V., and Zilinskas, A. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2, 1978.
- Neal, R. Slice sampling. *Annals of Statistics*, 31:705–767, 2000.
- Neal, R. M. *Bayesian learning for neural networks*. PhD thesis, University of Toronto, 1995.
- Nickson, T., Osborne, M. A., Reece, S., and Roberts, S. Automated machine learning using stochastic algorithm tuning. *NIPS Workshop on Bayesian Optimization*, 2014.
- Osborne, M. A., Garnett, R., and Roberts, S. J. Gaussian processes for global optimization. In *Learning and Intelligent Optimization*, 2009.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and variational inference in deep latent Gaussian models. In *ICML*, 2014.
- Snelson, E. and Ghahramani, Z. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pp. 1257–1264, 2005.
- Snoek, J. *Bayesian Optimization and Semiparametric Models with Applications to Assistive Technology*. PhD thesis, University of Toronto, Toronto, Canada, 2013.
- Snoek, J., Larochelle, H., and Adams, R. P. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, 2012.
- Snoek, J., Swersky, K., Zemel, R. S., and Adams, R. P. Input warping for Bayesian optimization of non-stationary functions. In *ICML*, 2014.
- Springenberg, J. T., Dosovitskiy, A., Brox, T., and Riedmiller, M. A. Striving for simplicity: The all convolutional net. *CoRR*, abs/1412.6806, 2014. URL <http://arxiv.org/abs/1412.6806>.
- Srinivas, N., Krause, A., Kakade, S., and Seeger, M. Gaussian process optimization in the bandit setting: no regret and experimental design. In *ICML*, 2010.
- Swersky, K., Snoek, J., and Adams, R. P. Multi-task Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2013.
- Titsias, M. K. Variational learning of inducing variables in sparse Gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pp. 567–574, 2009.
- Wan, L., Zeiler, M. D., Zhang, S., LeCun, Y., and Fergus, R. Regularization of neural networks using dropconnect. In *ICML*, 2013.
- Wang, Z., Zoghi, M., Hutter, F., Matheson, D., and de Freitas, N. Bayesian optimization in high dimensions via random embeddings. In *IJCAI*, 2013.
- Williams, C. K. I. Computing with infinite networks. In *Advances in Neural Information Processing Systems*, 1996.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., and Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. *arXiv preprint arXiv:1502.03044v2*, 2015.
- Zaremba, W., Sutskever, I., and Vinyals, O. Recurrent neural network regularization. *arXiv preprint arXiv:1207.0580*, 2015.