

Bayesian Monte Carlo for the Global Optimization of Expensive Functions¹

Perry Groot

Adriana Birlutiu

Tom Heskes

*Radboud University Nijmegen, Institute for Computing and Information Sciences,
Heyendaalseweg 135, 6525 AJ Nijmegen, the Netherlands*

Abstract

In the last decades enormous advances have been made possible for modelling complex (physical) systems by mathematical equations and computer algorithms. To deal with very long running times of such models a promising approach has been to replace them by stochastic approximations based on a few model evaluations. In this paper we focus on the often occurring case that the system modelled has two types of inputs $\mathbf{x} = (\mathbf{x}_c, \mathbf{x}_e)$ with \mathbf{x}_c representing control variables and \mathbf{x}_e representing environmental variables. Typically, \mathbf{x}_c needs to be optimised, whereas \mathbf{x}_e are uncontrollable but are assumed to adhere to some distribution. In this paper we use a Bayesian approach to address this problem: we specify a prior distribution on the underlying function using a Gaussian process (GP) and use Bayesian Monte Carlo (BMC) to obtain the objective function by integrating out environmental variables. We empirically evaluate several active learning criteria that were developed for the deterministic case (i.e., no environmental variables) and show that the ALC criterion appears significantly better than expected improvement and random selection.

1 Framework

Suppose we were to introduce a new cake mix into the consumer market that we like to be robust against inaccurate settings of oven temperature (T) and baking time (t). We would like to design experiments varying the control variables – the amount of flour (F), the amount of sugar (S), and the amount of egg powder (E) – and environmental variables (oven temperature and baking time) to see if we could create a cake mix that is better with respect to the environmental variables than the standard recipe so far produced by the product development laboratory.

We can represent this case study as having control variables $\mathbf{x}_c = \{F, S, E\}$, environmental variables $\mathbf{x}_e = \{T, t\}$, and as objective function $\ell(\mathbf{x}_c)$ the user rating of the cake mix. We focus on optimising the objective function $\ell(\mathbf{x}_c)$ with each output value the mean over the distribution of environmental variables. Formally $\mathbf{x}_c^* = \operatorname{argmax}_{\mathbf{x}_c} \ell(\mathbf{x}_c) = \operatorname{argmax}_{\mathbf{x}_c} \int_{\mathbf{x}_e} f(\mathbf{x}_c, \mathbf{x}_e) p(\mathbf{x}_e) d\mathbf{x}_e$. When f is expensive to evaluate, i.e., baking a cake, this means we can only obtain a small number of function evaluations leading to uncertainty about f and determining ℓ can be considered an inference problem. We address this problem by modelling f with a GP which places a prior distribution on f . We then use BMC, a Bayesian approach for evaluating integrals, to integrate out the environmental variables resulting in a stochastic approximation L of the objective function ℓ . Often L is hard to evaluate, but in some cases like the common case of a Gaussian kernel for modelling f and a Gaussian measure for $p(\mathbf{x}_e)$, the approximation L can be expressed in terms of analytical expressions. Technical details can be found in the full paper. The goal, however, is to find the value \mathbf{x}_c^* such that $\ell(\mathbf{x}_c^*)$ is maximised. The idea is to request more information about the true objective function ℓ (through f) in strategically chosen inputs, update our stochastic approximation L , and use the resulting model to make a prediction: $\tilde{\mathbf{x}}_c^* = \operatorname{argmax}_{\mathbf{x}_c} \bar{L}(\mathbf{x}_c | \mathcal{D}_n) = \operatorname{argmax}_{\mathbf{x}_c} \int_{\mathbf{x}_e} \bar{f}(\mathbf{x}_c, \mathbf{x}_e | \mathcal{D}_n) p(\mathbf{x}_e) d\mathbf{x}_e$ with \bar{L} the mean of L and $\mathcal{D}_n = \{((\mathbf{x}_c, \mathbf{x}_e)_i, f((\mathbf{x}_c, \mathbf{x}_e)_i))\}_{i=1}^n$ the n function evaluations collected so far.

There are two key aspects that distinguish our problem formulation with previous work. First, we do not optimise f , but the average that f takes over a distribution of the environmental variables. Second, our objective function L is a collection of stochastic variables which is only observed indirectly through

¹Published in the Proceedings of the 19th European Conference on Artificial Intelligence (ECAI), 2010.

$f(\mathbf{x}_c, \mathbf{x}_e)$ whereas previous work almost exclusively focusses on optimising a *deterministic* function h that is directly observed through observations $h(\mathbf{x})$.

We investigated several active learning criteria applicable to our problem formulation. The (generalised) Expected Improvement (EI^g) with non-negative parameter g (standard $g = 1$), which selects the point with the best expected improvement with respect to the current best value found so far; the ALM method selects the point with the highest variance; the ALC method selects the point with the highest model change (i.e., total variance reduction over a set of reference points); a Latin Hypercube Design (LHD) chooses a design of n points $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ beforehand such that for each dimension j all x_{ij} are distinct.

2 Experiments

We used DIRECT as global optimiser and the maximum value of the corresponding objective function ℓ was found to be 5.5330 and was obtained at $\mathbf{x}_c^* = (F^*, S^*, E^*) = (1.1852, -0.7407, 1.1084)$, implying an improved cake mix by using a higher amount of flour, a lower amount of sugar, and a higher amount of egg powder than the standard recipe set at $(0, 0, 0)$.

The goal of the various active learning criteria is to find the value \mathbf{x}_c^* that maximises ℓ as quickly as possible using properties of the stochastic approximation L . Therefore, let $\tilde{\mathbf{x}}_c^*$ be the value that maximises \bar{L} the current mean GP estimate of ℓ . We take as error measure ϵ the distance between the true maximum value and the true value at the predicted location $\tilde{\mathbf{x}}_c^*$, i.e., $\epsilon = |\max_{\mathbf{x}_c} \ell(\mathbf{x}_c) - \ell(\tilde{\mathbf{x}}_c^*)|$ with $\tilde{\mathbf{x}}_c^* = \operatorname{argmax}_{\mathbf{x}_c} \bar{L}(\mathbf{x}_c)$.

The results are shown in Figure 1 in which we plot the error measure and the standard deviation of the mean over 50 runs. Clearly, the ALM method performs worse than random sampling. The LHD approach performs better than ALM, but its performance is initially very similar to random sampling and after about 25 samples it is even outperformed by random sampling. Although LHDs are typically used as initialisation method in the literature these results suggest that an LHD is unnecessary and may lead to worse performance. Similar results for LHDs and deterministic functions have also been reported in the literature. The ALC criterion performs very well on the cake mix study. The downside, however, is that ALC is computationally more challenging and tries to optimise the global model fit, but not specifically the predicted maximum of the objective function. We also combined criteria to first obtain an $\hat{\mathbf{x}}_c$ which was then kept fixed in a second criterion to obtain a pair $(\hat{\mathbf{x}}_c, \hat{\mathbf{x}}_e)$ for further evaluation. We investigated the combination of ALC-ALC and EI-ALC. The ALC-ALC criterion is a bit worse than the ALC criterion, but performs quite well. It has the advantage that optimisation in a high dimensional space of both control and environmental variables can be split into two sequential optimisation steps in two lower dimensional spaces. The generalised expected improvement criterion has originally been developed for deterministic functions and has been the method of choice for surrogate modelling and optimisation. These results, however, show that the criterion cannot easily be augmented to be used for the optimisation of functions that are dependent on both control variables and environmental variables.

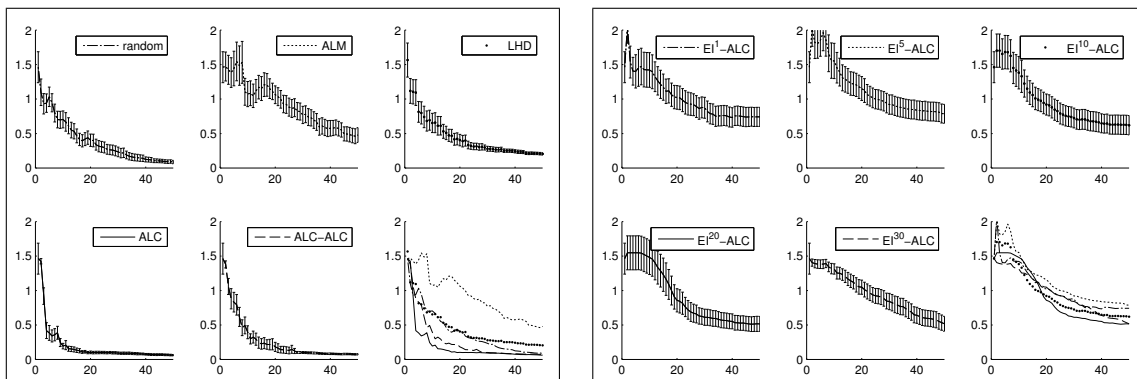


Figure 1: For each active learning criterion we computed a sequence of samples and observations to be added. At each step we computed the distance between the true maximum value and the value at the location where we predict the maximum value to be. We computed the mean performance and standard deviation of the mean over 50 runs. The bottom right subfigure in both panels superimposes the means.