Supplementary material for: Multi-attribute Bayesian optimization with interactive preference learning

Raul Astudillo
Cornell University

Peter I. Frazier
Cornell University, Uber

1 UNBIASED ESTIMATOR OF THE GRADIENT OF EI-UU

In this section we formally state and prove Proposition 1.

Proposition 1. *Suppose that $U(\cdot;\theta), \theta \in \Theta$ is differentiable for all $\theta \in \Theta$ and let $\mathbb{X}'$ be an open subset of $\mathbb{X}$ so that μ_n and K_n are differentiable on $\mathbb{X}'$ and there exists a measurable function $\eta : \mathbb{R}^k \to \mathbb{R}$ satisfying*

1. $\|\nabla U(\mu_n(x) + C_n(x)Z;\theta)\| < \eta(\theta, Z)$ *for all $x \in \mathbb{X}'$, $\theta \in \Theta$ and $Z \in \mathbb{R}^k$.*
2. $\mathbb{E}[\eta(\theta, Z)] < \infty$, *where Z is a m-variate standard normal random vector independent of θ, and the expectation is over both θ and Z.*

Further, suppose that for almost every $\theta \in \Theta$ and $Z \in \mathbb{R}^k$ the set $\{x \in \mathbb{X}' : U(\mu_n(x) + C_n(x)Z;\theta) = U_n^*(f;\theta)\}$ is countable. Then, EI-UU is differentiable on $\mathbb{X}'$ and its gradient, when it exists, is given by

$$\nabla \textit{EI-UU}(x) = \mathbb{E}\left[\gamma(x, \theta, Z)\right],$$

where the expectation is over θ and Z, and

$$\gamma(x, \theta, Z) = \begin{cases} \nabla U(\mu_n(x) + C_n(x)Z;\theta), \text{ if } U(\mu_n(x) + C_n(x)Z) > U_n^*(f;\theta), \\ 0, \text{ otherwise.} \end{cases}$$

Proof. From the given hypothesis it follows that, for any fixed $\theta \in \Theta$ and $Z \in \mathbb{R}^k$, the function $x \mapsto U(\mu_n(x) + C_n(x)Z;\theta)$ is differentiable on $\mathbb{X}'$. This in turn implies that the function $x \mapsto \{U(\mu_n(x) + C_n(x)Z;\theta - U_n^*(f;\theta)\}_+$ is continuous on $\mathbb{X}'$ and differentiable at every $x \in \mathbb{X}'$ such that $U(\mu_n(x) + C_n(x)Z;\theta \neq U_n^*(f;\theta)$, with gradient equal to $\gamma(x, \theta, Z)$. From our assumption that for almost every θ and Z the set $\{x \in \mathbb{X} : U(\mu_n(x) + C_n(x)Z;\theta) = U_n^*(f;\theta)\}$ is countable, it follows that for almost every θ and Z the function $x \mapsto \{U(\mu_n(x) + C_n(x)Z;\theta - U_n^*(f;\theta)\}_+$ is continuous on $\mathbb{X}'$ and differentiable on all $\mathbb{X}'$, except maybe on a countable subset. Using this, along with conditions 1 and 2, and Theorem 1 in L'Ecuyer (1990), the desired result follows. □

We note that, if one imposes the stronger condition $\mathbb{E}[\eta(\theta, Z)^2] < \infty$, then γ has finite second moment, and thus this unbiased estimator of ∇EI-UU(x) can be used within stochastic gradient ascent to find a stationary point of EI-UU (Bottou, 2010).

2 COMPUTATION OF EI-UU AND ITS GRADIENT WHEN U IS LINEAR

In this section we formally state and prove Propositions 2 and 3.

Proposition 2. *Suppose that $U(y;\theta) = \theta^\top y$ for all $\theta \in \Theta$ and $y \in \mathbb{R}^k$. Then,*

$$\text{EI-UU}(x) = \mathbb{E}_n \left[\Delta_n(x;\theta)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \sigma_n(x;\theta)\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)\right]$$

where the expectation is over θ, $\Delta_n(x;\theta) = \theta^\top \mu_n(x) - U_n^*(f;\theta)$, $\sigma_n(x;\theta) = \sqrt{\theta^\top K_n(x)\theta}$, and φ and Φ are the standard normal probability density function and cumulative distribution function, respectively.

Proof. Note that

$$\text{EI-UU}(x) = \mathbb{E}_n\left[\mathbb{E}_n\left[\left\{\theta^\top f(x) - U_n^*(f;\theta)\right\}_+ \mid \theta\right]\right].$$

Thus, it suffices to show that

$$\mathbb{E}_n\left[\left\{\theta^\top f(x) - U_n^*(f;\theta)\right\}_+ \mid \theta\right] = \Delta_n(x;\theta)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \sigma_n(x;\theta)\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right),$$

but this can be easily verified by noting that, conditioned on θ, the time-n posterior distribution of $\theta^\top f(x)$ is normal with mean $\theta^\top \mu_n(x)$ and variance $\theta^\top K_n(x)\theta$. □

Proposition 3. *Suppose that $U(y;\theta) = \theta^\top y$ for all $\theta \in \Theta \subset \mathbb{R}^k$ and $y \in \mathbb{R}^k$, μ_n and K_n are differentiable, and there exists a function $\eta : \Theta \to \mathbb{R}$ satisfying*

1. $\left\|\left(\theta^\top \nabla \mu_n(x)\right)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j \nabla K_n(x)_{i,j}\right\| \leq \eta(\theta)$ *for all $x \in \mathbb{X}$ and $\theta \in \Theta$.*
2. $\mathbb{E}[\eta(\theta)] < \infty$.

Then, EI-UU is differentiable and its gradient is given by

$$\nabla \text{EI-UU}(x) = \mathbb{E}_n\left[\left(\theta^\top \nabla \mu_n(x)\right)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j\nabla K_n(x)_{i,j}\right].$$

Proof. Recall that

$$\mathbb{E}_n\left[\left\{\theta^\top f(x) - U_n^*(f;\theta)\right\}_+ \mid \theta\right] = \Delta_n(x;\theta)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \sigma_n(x;\theta)\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right).$$

Moreover, standard calculations show that

$$\nabla\left[\Delta_n(x;\theta)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)\right] = \left(\theta^\top\nabla\mu_n(x)\right)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \Delta_n(x;\theta)\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)\nabla\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)},$$

and

$$\begin{aligned}\nabla\left[\sigma_n(x;\theta)\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)\right] &= \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j\nabla K_n(x)_{i,j} + \sigma_n(x;\theta)\left[-\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)\nabla\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right]\\ &= \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j\nabla K_n(x)_{i,j} - \Delta_n(x;\theta)\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)\nabla\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}.\end{aligned}$$

Thus, $\mathbb{E}_n\left[\left\{\theta^\top f(x) - U_n^*(f;\theta)\right\}_+ \mid \theta\right]$ is a differentiable function of x, and its gradient is given by

$$\nabla\mathbb{E}_n\left[\left\{\theta^\top f(x) - U_n^*(f;\theta)\right\}_+ \mid \theta\right] = \left(\theta^\top\nabla\mu_n(x)\right)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j\nabla K_n(x)_{i,j}.$$

From conditions 1 and 2, and theorem 16.8 in Clarke and Billingsley (1980), it follows that EI-UU is differentiable and its gradient is given by

$$\nabla\text{EI-UU}(x) = \mathbb{E}_n\left[\nabla\mathbb{E}_n\left[\left\{\theta^\top f(x) - U_n^*(f;\theta)\right\}_+ \mid \theta\right]\right]$$

i.e.,

$$\nabla\text{EI-UU}(x) = \mathbb{E}_n\left[\left(\theta^\top\nabla\mu_n(x)\right)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j\nabla K_n(x)_{i,j}\right].$$

□

We end by noting that if Θ is compact and μ_n and K_n are both continuously differentiable, then

$$(\theta,x)\to\left\|\left(\theta^\top\nabla\mu_n(x)\right)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j\nabla K_n(x)_{i,j}\right\|$$

is continuous and thus attains its maximum value on $\Theta\times\mathbb{X}$ (recall that $\mathbb{X}$ is compact as well). Thus, in this case conditions 1 and 2 are satisfied by the constant function

$$\eta \equiv \max_{(\theta,x)\in\Theta\times\mathbb{X}}\left\|\left(\theta^\top\nabla\mu_n(x)\right)\Phi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right) + \frac{\varphi\left(\frac{\Delta_n(x;\theta)}{\sigma_n(x;\theta)}\right)}{2\sigma_n(x;\theta)}\sum_{i,j=1}^m \theta_i\theta_j\nabla K_n(x)_{i,j}\right\|.$$

3 THOMPSON SAMPLING UNDER UTILITY UNCERTAINTY (TS-UU)

Thompson sampling for utility uncertainty (TS-UU) generalizes the well-known Thompson sampling method (Thompson, 1933) to our setting. TS-UU works as follows. It first samples θ from its posterior distribution. Then, it samples f from its Gaussian process posterior distribution. The point at which it evaluates f next is the one that maximizes $U(f(x);\theta)$ for the samples of f and θ. This contrasts with the point-estimate approach in that it samples θ from its posterior rather than simply setting it equal to a point estimate. For example, if we implemented this point-estimate approach using standard Thompson sampling, we would sample only f from its posterior and then optimize $U(f(x);\hat{\theta})$ where $\hat{\theta}$ is a point estimate, such as the maximum a posteriori estimate. TS-UU can induce substantially more exploration than this more classical approach.

TS-UU can be implemented by sampling $f(x)$ over a grid of points if x is low-dimensional. It can also be implemented for higher-dimensional x by optimizing f with a method for continuous nonlinear optimization (like CMA, Hansen (2016)), lazily sampling from the posterior on f each new point that CMA wants to evaluate, conditioning on previous real and sampled evaluations. We use the latter approach in our numerical experiments.

4 EXPLORATION AND EXPLOITATION TRADE-OFF

One of the key properties of the classical expected improvement acquisition function is that it is increasing with respect to both the posterior mean and variance. This means that it prefers to sample points that are either promising with respect to our current knowledge or are still highly uncertain, an appealing property for a sampling policy aiming to balance exploitation and exploration. The following result shows that, under certain conditions, the EI-UU sampling policy satisfies an analogous property.

Proposition 4. *Suppose that for every $\theta\in\Theta$ $U(\cdot;\theta)$ is convex and non-decreasing. Also suppose $x,x'\in\mathbb{X}$ are such that $\mu_n(x)\geq\mu_n(x')$ and $K_n(x)\gtrsim K_n(x')$, where the first inequality is coordinate-wise and $\gtrsim$ denotes the partial order defined by the cone of positive semi-definite matrices. Then,*

$$\textit{EI-UU}_n(x)\geq\textit{EI-UU}_n(x').$$

Proof. Since $K_n(x) \gtrsim K_n(x')$, we have that $f(x) \overset{d}{=} f(x') + (\mu_n(x) - \mu_n(x')) + W$, where W is a k-variate normal random vector with zero mean and covariance matrix $K_n(x) - K_n(x')$ independent of $f(x')$. Thus,

$$\begin{aligned}
\mathbb{E}_n\left[\{U(f(x);\theta) - U_n^*(f;\theta)\}_+ \mid \theta\right] &= \mathbb{E}_n\left[\{U(f(x') + (\mu_n(x) - \mu_n(x')) + W;\theta) - U_n^*(f;\theta)\}_+ \mid \theta\right] \\
&\geq \mathbb{E}_n\left[\{U(f(x') + W;\theta) - U_n^*(f;\theta)\}_+ \mid \theta\right] \\
&= \mathbb{E}_n\left[\mathbb{E}_n\left[\{U(f(x') + W;\theta) - U_n^*(f;\theta)\}_+ \mid \theta, f(x')\right]\right] \\
&\geq \mathbb{E}_n\left[\{U(f(x');\theta) - U_n^*(f;\theta)\}_+ \mid \theta\right],
\end{aligned}$$

where the first and second inequalities follow from the fact that the function $y \mapsto \{U(y;\theta) - U_n^*(f;\theta)\}_+$ is increasing and convex, respectively, along with Jensen's inequality. Finally, taking expectations with respect to θ yields the desired result. □

This result implies, for example, that for linear utility functions, the EI-UU sampling policy exhibits the behavior described above. We also note, however, that most utility functions used in practice are concave instead of convex.

5 SYNTHETIC TEST FUNCTIONS DEFINITIONS

5.1 DTLZ1a

A general form of this test function was first introduced in Deb et al. (2005). The version we use was defined in Knowles (2006). It is defined over $\mathbb{X} = [0,1]^6$, and has $k = 2$ attributes given by

$$\begin{aligned}
f_1(x) &= -0.5x_1\left((1 + g(x)\right) \\
f_2(x) &= -0.5(1 - x_1)\left((1 + g(x)\right),
\end{aligned}$$

where

$$g(x) = 100\left(5 + \sum_{i=2}^{6}\left[(x_i - 0.5)^2 - \cos\left(2\pi(x_i - 0.5)\right)\right]\right).$$

The Pareto optimal set of designs consists of those such that $x_i = 0.5$, $i = 2, \ldots, 6$, and x_1 may take any value in $[0,1]$. The Pareto front is a segment of the hyperplane $y_1 + y_2 = -0.5$.

5.2 DTLZ2

This function was first introduced in a general form in Deb et al. (2005). In our experiment, we use a concrete version of it with $k = 4$ attributes defined over $\mathbb{X} = [0,1]^5$. The attributes are

$$\begin{aligned}
f_1(x) &= -\left(1 + g(x)\right)\prod_{i=1}^{3}\cos\left(\frac{\pi}{2}x_i\right) \\
f_2(x) &= -(1 + g(x))\left(\prod_{i=1}^{2}\cos\left(\frac{\pi}{2}x_i\right)\right)\sin\left(\frac{\pi}{2}x_3\right), \\
f_3(x) &= -(1 + g(x))\cos\left(\frac{\pi}{2}x_1\right)\sin\left(\frac{\pi}{2}x_2\right), \\
f_4(x) &= -(1 + g(x))\sin\left(\frac{\pi}{2}x_1\right),
\end{aligned}$$

where

$$g(x) = \sum_{i=4}^{5}(x_i - 0.5).$$

.

5.3 VLMOP3

This test function first appeared in Van Veldhuizen and Lamont (1999). It is defined over $\mathbb{X} = [-3, 3]^2$ and has $k = 3$ attributes given by

$$
\begin{aligned}
f_1(x) &= -0.5(x_1^2 + x_2^2) - \sin(x_1^2 + x_2^2), \\
f_2(x) &= -\frac{(3x_1 - 2x_2 + 4)^2}{8} - \frac{(x_1 - x_2 + 1)^2}{27} - 15, \\
f_3(x) &= -\frac{1}{x_1^2 + x_2^2 + 1} + 1.1 \exp\left(-x_1^2 - x_2^2\right).
\end{aligned}
$$

References

Bottou, L. (2010). On-line Learning and Stochastic Approximations. *On-Line Learning in Neural Networks*, 17(9):9–42.

Clarke, L. E. and Billingsley, P. (1980). *Probability and Measure*, volume 64 of *Wiley Series in Probability and Statistics*. Wiley.

Deb, K., Thiele, L., Laumanns, M., and Zitzler, E. (2005). Scalable Test Problems for Evolutionary Multiobjective Optimization. In *Evolutionary Multiobjective Optimization*, pages 105–145. Springer.

Hansen, N. (2016). The CMA Evolution Strategy: A Tutorial. *arXiv preprint arXiv:1604.00772*.

Knowles, J. (2006). ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66.

L'Ecuyer, P. (1990). A unified view of the IPA, SF, and LR gradient estimation techniques. *Management Science*, 36(11):1364–1383.

Thompson, W. R. (1933). On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3/4):285.

Van Veldhuizen, D. A. and Lamont, G. B. (1999). Multiobjective evolutionary algorithm test suites. In *Proceedings of the ACM Symposium on Applied Computing*, pages 351–357.