Online Supplement to “Balanced Explorative and Exploitative Search with Estimation for Simulation Optimization”

Sigrún Andradóttir
H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA, sa@gatech.edu

Andrei A. Prudius
Quantitative Financial Development, Bloomberg LP, 731 Lexington Avenue, New York, NY 10022, USA, aprudius@gmail.com

---

The pseudo-code for the R-BEESE method applied to transient simulation optimization is given in Algorithm 4. The sampling distribution update procedure for A-BEESE is provided in Algorithm 5. The pseudo-code for the A-BEESE method applied to transient simulation optimization is given in Algorithm 6. The proofs of Theorems 1, 2, 3, and 4 appear below.

Proof of Theorem 1: (i) For $k \in \mathbb{N} \setminus \{0\}$ and $n \in \mathbb{N}$, define Ω_k^n to be the subset of Ω such that the R-BEES method samples a solution in the set A_k at iteration n using the global sampling distribution G. Fix $k \in \mathbb{N} \setminus \{0\}$ and let $\Omega_k = \{\omega \in \Omega : \Omega_k^n \text{ i.o.}\}$ (i.o. stands for infinitely often). Observe that $\mathbb{P}(\Omega_k^n) = pG(A_k) > 0$. Then $\sum_{n=0}^{\infty} \mathbb{P}(\Omega_k^n) = \sum_{n=0}^{\infty} pG(A_k) = \infty$. Note that $\{\Omega_k^n\}_{n=0}^{\infty}$ are independent events. Then the second Borel-Cantelli lemma yields that $\mathbb{P}(\Omega_k) = 1$. Let $\tilde{\Omega} = \bigcap_{k=1}^{\infty} \Omega_k$. Obviously $\mathbb{P}(\tilde{\Omega}) = 1$. Fix $\omega \in \tilde{\Omega}$ and $l \in \mathbb{N} \setminus \{0\}$. Then there exists an iteration number $N_l(\omega)$ such that some solution in the set A_l is sampled at iteration N_l. As l is arbitrary, we get that $f(\theta_n^*) \to f^*$ as $n \to \infty$ for this ω. This concludes the proof of (i).
(ii) The proof is the same as in (i) except that A_k needs to be substituted by B_k. ■

Proof of Theorem 2: Let $\Omega_1 \subset \Omega$ be such that on Ω_1, $\liminf_{n\to\infty} C_n(\theta)/n \geq \epsilon p(1-\alpha)$ and $\lim_{k\to\infty} \hat{f}_k(\theta) = f(\theta)$ for all $\theta \in \Theta$. Since Θ is finite and $G(\{\theta\}) \geq \epsilon$ for all $\theta \in \Theta$, Assumption 1 implies that $\mathbb{P}(\Omega_1) = 1$. Fix $\omega \in \Omega_1$ and let $\gamma = \max_{\theta\in\Theta} f(\theta) - \max_{\theta\in\Theta\setminus\Theta^*} f(\theta) > 0$. Then there exists an iteration number $N(\omega)$ such that for all $n \geq N(\omega)$ and $\theta \in \Theta$, we have that $C_n(\theta, \omega) \geq M_n$ and $|f_n(\theta, \omega) - f(\theta)| < \gamma/2$. Hence, $\theta_n^*(\omega) \in \Theta^*$ for all $n \geq N(\omega)$. ■

Proof of Theorem 3: (i) Suppose that all random elements in the A-BEES method (ones needed for sampling solutions) are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. First we show

Algorithm 4 (R-BEESE Algorithm for Transient Simulation Optimization)

$n \leftarrow 0$
Sample a solution θ from the global distribution G
Generate $X_\theta^1, \ldots, X_\theta^r$ independent of everything else
$C(\theta) \leftarrow r$, $\Sigma(\theta) \leftarrow \sum_{i=1}^r h_\theta(X_\theta^i)$, $f(\theta) \leftarrow \Sigma(\theta)/C(\theta)$
for $\theta' \in \Theta \setminus \{\theta\}$ **do**
  $C(\theta') \leftarrow 0$, $\Sigma(\theta') \leftarrow 0$
end for
$\theta_0 \leftarrow \theta$
while Stopping criterion is not satisfied **do**
  Draw a uniform $(0, 1)$ random variable U independent of everything else
  if $U \leq \alpha$ **then**
    $\theta \leftarrow \theta_n$
  else
    if $U \leq \alpha + (1 - \alpha)p$ **then**
      Sample a solution θ from the global distribution G independent of everything else
    else
      Sample a solution θ from a local distribution in $\mathcal{L}$
    end if
  end if
  Generate $X_\theta^{C(\theta)+1}, \ldots, X_\theta^{C(\theta)+r}$ independent of everything else
  $C(\theta) \leftarrow C(\theta) + r$, $\Sigma(\theta) \leftarrow \Sigma(\theta) + \sum_{i=1}^r h_\theta(X_\theta^{C(\theta)+i})$, $f(\theta) \leftarrow \Sigma(\theta)/C(\theta)$
  if $f(\theta) > f(\theta_n)$ **then**
    $\theta_{n+1} \leftarrow \theta$
  else
    $\theta_{n+1} \leftarrow \theta_n$
  end if
  $n \leftarrow n + 1$
end while
if $C(\theta_n) < M_n$ **and** $\{\theta \in \Theta : C(\theta) \geq M_n\} \neq \emptyset$ **then**
  $\theta_n \leftarrow \arg\max_{\{\theta \in \Theta : C(\theta) \geq M_n\}} f(\theta)$
end if
Present $\theta_n^* = \theta_n$ as the estimate of the optimal solution

Algorithm 5 (Sampling Distribution Update Procedure for A-BEESE)

if LocalSearch=true **then**
  if $\Delta \leq \delta$ **then**
    LocalSearch $\leftarrow$ false
    GScounter $\leftarrow 0$, $k \leftarrow k_g$
    $v_l^* \leftarrow v^*$
  end if
else
  GScounter $\leftarrow$ GScounter $+1$
  if GScounter $= g$ **then**
    LocalSearch $\leftarrow$ true
    $k \leftarrow k_l$
  else
    if $\Delta \leq \delta$ **then**
      if $v^* - v_l^* \geq \delta$ **then**
        LocalSearch $\leftarrow$ true
        $k \leftarrow k_l$
      end if
    else
      if $D \leq d$ **then**
        LocalSearch $\leftarrow$ true
        $k \leftarrow k_l$
      end if
    end if
  end if
end if

that the A-BEES method uses the global sampling distribution G i.o. for every $\omega \in \Omega$. Contrary to this, suppose that there exists an $\omega \in \Omega$ such that the global sampling distribution is used a finite number of times. Then there exists an iteration number $N_0(\omega)$ such that local sampling distributions are used from this iteration on. Let $\bar{f}(\omega)$ be the objective function value at the best point found prior to iteration $N_0(\omega)$. Since Algorithm 2 switches from local to global search whenever $\Delta \leq \delta$, the maximum number of successive reviews for which the local search will be continued after iteration $N_0(\omega)$ is $\lfloor (f^* - \bar{f}(\omega))/\delta \rfloor + 1 < \infty$, where $\lfloor \cdot \rfloor$ is the floor function. This provides a contradiction. The remainder of the proof is identical to the proof of part (i) of Theorem 1 except that Ω_k^n is defined to be the subset of Ω such that the A-BEES method samples a solution in the set A_k when the global sampling distribution G is used for the n^{th} time for all $k \in \mathbb{N} \setminus \{0\}$ and $n \in \mathbb{N}$, so that $\mathbb{P}(\Omega_k^n) = G(A_k) > 0$. This concludes the proof of (i).

Algorithm 6 (A-BEESE Algorithm)

counter $\leftarrow 0$, $n \leftarrow 0$
LocalSearch $\leftarrow$ false
GScounter $\leftarrow 0$, $k \leftarrow k_g$
Sample a solution θ from the global distribution G
Generate $X_\theta^1, \ldots, X_\theta^r$ independent of everything else
$C(\theta) \leftarrow r$, $\Sigma(\theta) \leftarrow \sum_{i=1}^r h_\theta(X_\theta^i)$, $f(\theta) \leftarrow \Sigma(\theta)/C(\theta)$
for $\theta' \in \Theta \setminus \{\theta\}$ **do**
  $C(\theta') \leftarrow 0$, $\Sigma(\theta') \leftarrow 0$
end for
Let $v^*, v_l^* \leftarrow f(\theta)$ and $\theta_0 \leftarrow \theta$
while Stopping criterion is not satisfied **do**
  Draw a uniform $(0, 1)$ random variable U independent of everything else
  if $U \leq \alpha$ **then**
    $\theta \leftarrow \theta_n$
  else
    if LocalSearch=true **then**
      Sample a solution θ from a local distribution in $\mathcal{L}$
    else
      Sample a solution θ from the global distribution G independent of everything else
    end if
  end if
  Generate $X_\theta^{C(\theta)+1}, \ldots, X_\theta^{C(\theta)+r}$ independent of everything else
  Update $C(\theta)$, $\Sigma(\theta)$, and $f(\theta)$
  counter $\leftarrow$ counter+1
  Compute θ_{n+1} and update v^* (if needed)
  if counter=k **then**
    Update Δ and D
    Update search nature (use Algorithm 5)
    counter $\leftarrow 0$
  end if
  $n \leftarrow n + 1$
end while
if $C(\theta_n) < M_n$ **and** $\{\theta \in \Theta : C(\theta) \geq M_n\} \neq \emptyset$ **then**
  $\theta_n \leftarrow \arg\max_{\{\theta \in \Theta : C(\theta) \geq M_n\}} f(\theta)$
end if
Present $\theta_n^* = \theta_n$ as the estimate of the optimal solution

(ii) Let $\tilde{\Omega}_1 \subset \Omega$ be such that global sampling is performed a finite number of times. Fix $\omega \in \tilde{\Omega}_1$. Since the only reason for not switching back to global search is that the objective function value improvement between successive reviews is at least δ, we conclude that $f(\theta_n^*)$ goes to infinity for this ω. It remains to show that $f(\theta_n^*)$ diverges to infinity for almost every

$\omega \in \tilde{\Omega}_2 = \Omega \setminus \tilde{\Omega}_1$ (the set of ω's such that the global sampling distribution is used i.o.).

Fix $l \in \mathbb{N} \setminus \{0\}$ and $\omega \in \tilde{\Omega}_2 \bigcap \tilde{\Omega}'$, where $\tilde{\Omega}'$ is the set $\tilde{\Omega}$ of part (i) of the proof of Theorem 1 with the exceptions that A_k is substituted by B_k and Ω_k^n is defined as in part (i) of this proof. Then there exists an iteration number $N_l(\omega)$ such that some solution in the set B_l is sampled at iteration $N_l(\omega)$. As l is arbitrary, we get that $f(\theta_n^*) \to f^*$ as $n \to \infty$ for this ω.

Observe that $\mathbb{P}(\tilde{\Omega}') = 1$ (the proof of this is similar to the proof that $\mathbb{P}(\tilde{\Omega}) = 1$ in part (i) of Theorem 1) which implies that $\mathbb{P}(\tilde{\Omega}_2 \bigcap \tilde{\Omega}') = \mathbb{P}(\tilde{\Omega}_2)$, and the proof is complete. ■

Lemma 1. *Suppose that Assumption 1 holds and that Θ is finite. Then the A-BEESE algorithm uses the global sampling distribution G infinitely often with probability one.*

Proof: Suppose that all random elements in A-BEESE (needed for sampling and simulating solutions) are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\Omega_1 \subset \Omega$ be such that Assumption 1 holds for all $\theta \in \Theta$. Because Θ is finite, we have that $\mathbb{P}(\Omega_1) = 1$. It suffices to show that the global distribution G is used i.o. for every $\omega \in \Omega_1$. Suppose that there exists an $\omega \in \Omega_1$ for which G is used only a finite number of times. Let $\Theta(\omega) \subset \Theta$ be the set of alternatives that are sampled infinitely often. Since Θ is finite, there exists an iteration number $N_1(\omega)$ such that for all $n \geq N_1(\omega)$, local search is performed in iteration n, the algorithm samples points in the set $\Theta(\omega)$ in iteration n, and we have $|f_n(\theta) - f(\theta)| < \delta/4$ for all $\theta \in \Theta(\omega)$. Note that for any iterations $n_1, n_2 \geq N_1(\omega)$ and $\theta \in \Theta$, we have that $|f_{n_1}(\theta) - f_{n_2}(\theta)| < \delta/2$. Now consider two successive search reviews at iterations n_1 and n_2 ($n_1 < n_2$) that occur after iteration $N_1(\omega)$. We have that

$$\Delta = f_{n_2}(\theta_{n_2}) - f_{n_2}(\theta_{n_1}) \leq |f_{n_2}(\theta_{n_2}) - f_{n_1}(\theta_{n_2})| + [f_{n_1}(\theta_{n_2}) - f_{n_1}(\theta_{n_1})] + |f_{n_1}(\theta_{n_1}) - f_{n_2}(\theta_{n_1})| < \delta,$$

where we have used the definition of θ_{n_1}. Since $\Delta < \delta$, the algorithm switches from local to global search. This provides a contradiction and hence the proof is complete. ■

Proof of Theorem 4: Suppose that all random elements in A-BEESE (needed for sampling and simulating solutions) are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Observe that the random elements in the A-BEESE method are of four types, namely, the ones needed for estimating the objective function value at each solution (defined on Ω_s), the ones needed for sampling solutions using local sampling distributions in $\mathcal{L}$ (defined on Ω_l), the ones needed for sampling solutions using the global sampling distribution G (defined on Ω_g), and the ones needed to decide whether to sample the point with the highest estimated objective function value or use the current sampling distribution (defined on Ω_a). Hence, without loss of generality, we assume that $\Omega = \Omega_s \times \Omega_l \times \Omega_g \times \Omega_a$. Also, without loss of generality, we assume that $\Omega_g = \prod_{i=0}^{\infty} \Omega_g^i$, where Ω_g^i is a sample space on which random elements for sampling from the global distribution G for the i^{th} time are defined.

Let $\tilde{\Omega}_s \subset \Omega_s$ be such that Assumption 1 holds for each $\theta \in \Theta$. Because Θ is finite, we have that $\mathbb{P}(\tilde{\Omega}_s) = 1$. For each $\theta \in \Theta$, $l \in \mathbb{N}\backslash\{0\}$, and $\omega_g \in \Omega_g$, let $H_l(\theta, \omega_g)$ be the number of times a solution θ has been sampled by the time the global distribution G has been used for the l^{th} time (we also let $H_l(\theta, \omega) = H_l(\theta, \omega_g)$, where $\omega = (\omega_s, \omega_l, \omega_g, \omega_a) \in \Omega$). Let $\tilde{\Omega}_g \subset \Omega_g$ be such that $\liminf_{l\to\infty} H_l(\theta, \omega_g)/l \geq \epsilon$ for all $\theta \in \Theta$. Since $G(\{\theta\}) > \epsilon$ for all $\theta \in \Theta$ and Θ is finite, the Strong Law of Large Numbers implies that $\mathbb{P}(\tilde{\Omega}_g) = 1$. Let $\tilde{\Omega}_a \subset \Omega_a$ be such that the long-run average fraction of time the point with the highest estimated objective function value is sampled equals α. By the Strong Law of Large Numbers we have that $\mathbb{P}(\tilde{\Omega}_a) = 1$.

Fix $\omega \in \tilde{\Omega}_s \times \Omega_l \times \tilde{\Omega}_g \times \tilde{\Omega}_a$ and let $\gamma = \max_{\theta\in\Theta} f(\theta) - \max_{\theta\in\Theta\backslash\Theta^*} f(\theta) > 0$. Also define $\xi = \min(\gamma/2, \delta/4)$. By Lemma 1, the finiteness of Θ, and the choice of ω, it follows that there exists an iteration number $N_2(\omega)$ such that for all $\theta \in \Theta$ and $n \geq N_2(\omega)$, we have that $|f_n(\theta, \omega) - f(\theta)| < \xi$ (note that $\tilde{\Omega}_s = \Omega_1$ in the proof of Lemma 1). Recall that $C_n(\theta)$ denotes the number of times a solution θ has been simulated by iteration n. Then it suffices to show that there exists $N(\omega)$ such that $C_n(\theta, \omega) \geq M_n$ for all $\theta \in \Theta$ and $n \geq N(\omega)$.

Let $l_n(\omega)$ be the number of times the global distribution G is used by iteration n. Since $f_n(\theta)$ is within $\delta/4$ of $f(\theta)$ for all $\theta \in \Theta$ and $n \geq N_2(\omega)$, the improvement in the estimated objective function values between iterations n_1 and n_2 will be less than δ for all $n_1, n_2 \geq N_2(\omega)$. This implies that after some iteration $N_3(\omega)$, the A-BEESE method will perform $k_g * g$ iterations of global search followed by k_l iterations of local search, and then the cycle will repeat. Taking into account the sampling of θ_n with probability α, this implies that

$$\lim_{n\to\infty} \frac{l_n(\omega)}{n} = (1-\alpha)\frac{gk_g}{k_l + gk_g} > 0.$$

Then, for all $\theta \in \Theta$,

$$\begin{aligned}
\liminf_{n\to\infty} \frac{C_n(\theta,\omega)}{n} &= \liminf_{n\to\infty} \frac{C_n(\theta,\omega)}{l_n(\omega)} \times \lim_{n\to\infty} \frac{l_n(\omega)}{n} \\
&\geq \liminf_{n\to\infty} \frac{H_{l_n(\omega)}(\theta,\omega)}{l_n(\omega)} \times (1-\alpha)\frac{gk_g}{k_l+gk_g} \\
&\geq \epsilon(1-\alpha)\frac{gk_g}{k_l+gk_g} > 0.
\end{aligned}$$

Because $M_n = o(n)$, this implies that there exists an iteration number $N(\omega)$ so that $C_n(\theta, \omega) \geq M_n$ for every $\theta \in \Theta$ and all $n \geq N(\omega)$, and hence the proof is complete. ■