Balanced Explorative and Exploitative Search with Estimation for Simulation Optimization

雷可鹏

2021.11.29

Introduction

- Exploration: searching globally for promising solutions within the entire feasible region.
- Exploitation : locally search for improved solutions in promising subregions.
- Estimation : obtaining more precise objective function estimates at desirable alternatives and an improved estimator of the optimal solution.

Introduction

- Form: $\max_{\theta \in \Theta} f(\theta)$
- Θ is discrete, global sampling from Θ is possible, and $f(\theta)$ at any $\theta \in \Theta$ cannot be evaluated exactly and needs to be estimated via a "blackbox" simulation procedure.

Difficulties of optimization via simulation

- The noise in the estimated objective function values

Can be almost completely eliminated by performing a lot of simulation runs , but simulations are usually computationally expensive.

- The fact that simulation optimization problems often have little known structure (and solving even a deterministic optimization problem with little known structure is difficult).

We want to design specialized algorithms to solve the above problem that will search the feasible space thoroughly and yet be able to identify optimal or near-optimal solutions in the presence of noise.

Structural assumption

- NFL(no free lunch) theorems for deterministic optimization: that the average performance of each algorithm over all possible discrete optimization problems is identical. This suggests that a deterministic optimization problem will only be solved efficiently if it possesses some known structure and the optimization algorithm exploits that structure.

- Solutions located close to each other have similar performance.

Properties that optimization algorithms should possess in order to be efficient

Maintain balance between exploration and exploitation

Only know little about the structure of the objective function.

→ Start by exploring the entire feasible region.

→ Exploit good subregions by searching locally for better solutions.

→ The effectiveness of the search algorithm depend heavily on the ability of the method to identify when it should switch focus from global search (exploration) to local search (exploitation).

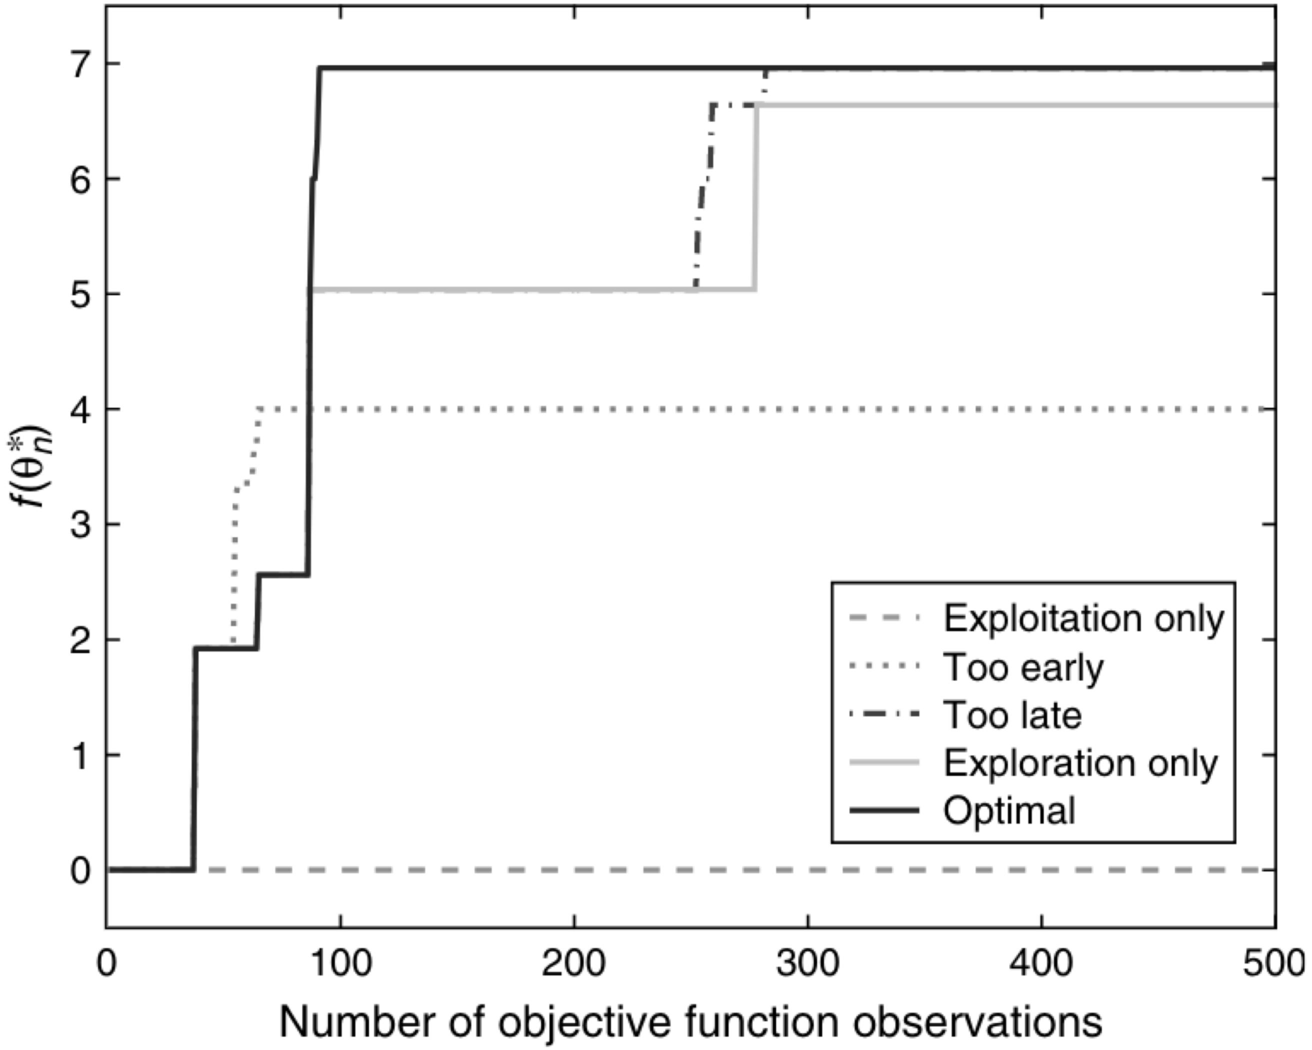

Figure 1 **Identification of a Proper Switch Point from Exploration to Exploitation**

Differences between deterministic and simulation optimization

- Presence of stochastic errors , which leads to two complications :

  1. More difficult to effectively guide the search for improved solutions.

  2. Select the best solution identified by the search.

- Simulation optimization problems are more likely to possess little known structure.

Where future simulations are to be conducted

Some solutions might appear to be good when in fact they are bad and vice versa.

→ Not be misled by such information for long.

→ Consider where additional simulations should be conducted to benefit the search the most as well as to be careful in deciding how much faith to put in the available function estimates, especially when choosing the estimate of the optimal solution. This issue will be further referred to as estimation.

Obtaining more precise function estimates

- Identifying an optimal solution among very good solutions rather than on locating good alternatives.
- How ? Achieved by allocating simulation effort to points with good estimated objective function values, under our assumption that solutions located close to one another have similar performance , local search of desirable regions will also yield improved objective function estimates at good points.

Local search help with exploitation and estimation.

- the good empirical performance of these methods is at least to some extent due to the fact that they happen to do estimation well (this observation has not been made by the original authors).

- Simulated annealing (SA) Algorithm

  Nested partitions(NP)method

Estimation of the optimal solution.

- The current solution
- The most visited solution
- The solution with the best estimated objective function value
- The solution with the best estimated objective function value, provided it has been simulated “sufficiently often”.

Contents

- R-BEESE : the randomized balanced explorative and exploitative search with estimation method

- A-BEESE : the adaptive explorative and exploitative search with estimation method

R-BEES for deterministic optimization

- The global sampling distribution G
- The family of local sampling distributions $\mathcal{L}$
- At any iteration, with probability $0<p\leqslant 1$, the global distribution is used , and with probability $1-p$,a local distribution in $\mathcal{L}$ is used.

Algorithm 1 (R-BEES Algorithm)

$n \leftarrow 0$
Sample a solution θ from the global distribution G
Evaluate the objective function at θ
$\theta_0 \leftarrow \theta$
while Stopping criterion is not satisfied **do**
  Draw a uniform $(0, 1)$ random variable U independent of everything else
  if $U \leq p$ **then**
    Sample a solution θ from the global distribution G independent of everything else
  else
    Sample a solution θ from a local distribution in $\mathscr{L}$
  end if
  Evaluate the objective function f at θ (if needed)
  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
Present $\theta_n^* = \theta_n$ as the estimate of the optimal solution

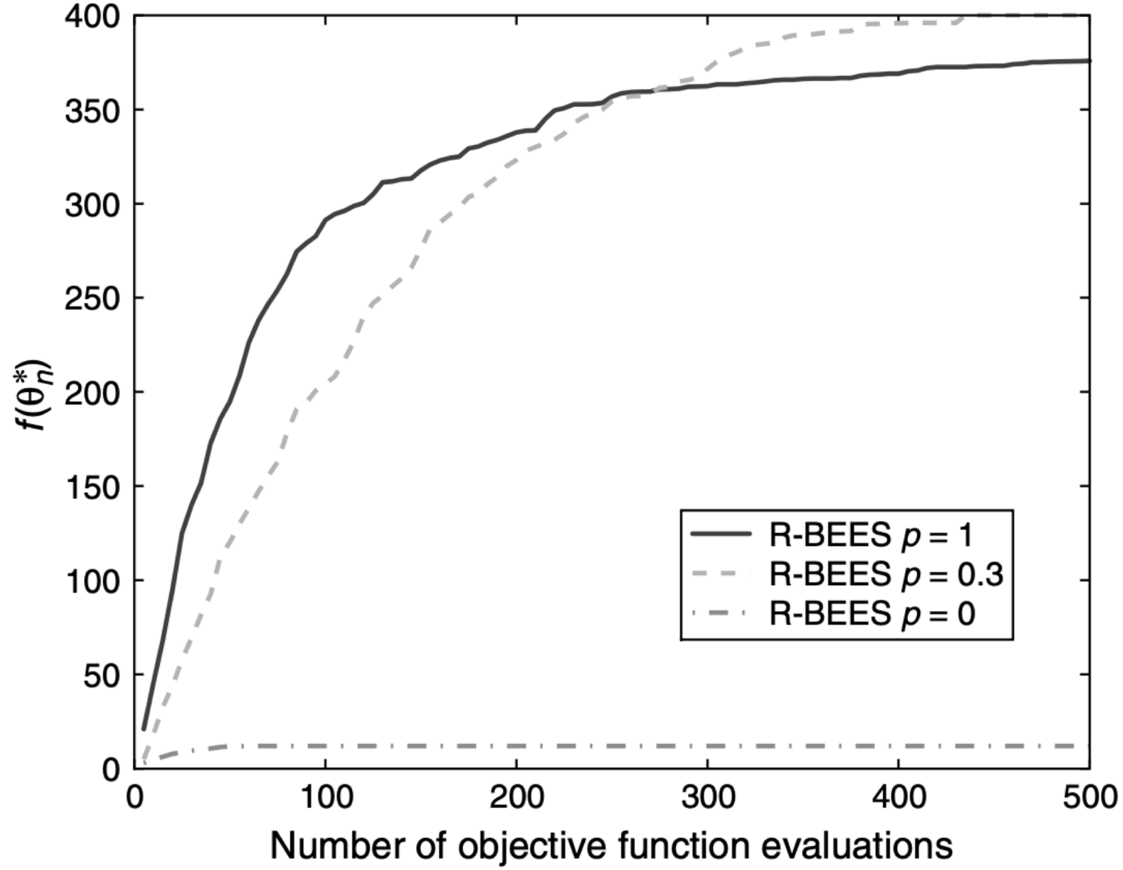

Figure 2 **Performance of the R-BEES Method on the Unimodal Problem with $\sigma^2 = 0$**

THEOREM 1. (i) *Suppose that* $f^* < \infty$. *Assume that the global sampling distribution* G *on* Θ *is such that for every* $k \in \mathbb{N}\backslash\{0\}$, $G(A_k) > 0$, *where* $A_k = \{\theta \in \Theta\colon f(\theta) \geq f^* - 1/k\}$. *Then, with probability one,* $f(\theta_n^*)$ *converges to* f^* *as* $n \to \infty$.

(ii) *Suppose that* $f^* = \infty$. *Assume that the global sampling distribution* G *on* Θ *is such that for every* k *integer* $G(B_k) > 0$, *where* $B_k = \{\theta \in \Theta\colon f(\theta) \geq k\}$. *Then, with probability one,* $f(\theta_n^*)$ *diverges to* $+\infty$ *as* $n \to \infty$.

$\mathcal{L}$:vehicle for exploiting the special structure

- One reasonable choice: uniform distribution over a ball $B_l(\theta_n)$ of radius l around θ_n. As the search progresses, information about the differences between objective function values of points located within balls of radius l becomes available , and we could alter the value of l accordingly.
- Generally, $\mathrm{N}(\theta) \subset \Theta$ is the neighborhood of each $\theta \in \Theta$, then it would be reasonable for the local distribution used in iteration n to be uniform on $\mathrm{N}(\theta_n)$.

R-BEESE : adding estimation component for simulation optimization

With probability $0 \leqslant \alpha < 1$, the point θ_n with the highest estimated objective function value is sampled. γ simulation replications are conducted at each sampled point.

Algorithm 1 (R-BEES Algorithm)

$n \leftarrow 0$

Sample a solution θ from the global distribution G

Evaluate the objective function at θ

$\theta_0 \leftarrow \theta$

while Stopping criterion is not satisfied **do**

  Draw a uniform $(0, 1)$ random variable U independent of everything else

  if $U \leq p$ **then**

    Sample a solution θ from the global distribution G independent of everything else

  else

    Sample a solution θ from a local distribution in $\mathscr{L}$

  end if

  Evaluate the objective function f at θ (if needed)

  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

Present $\theta_n^* = \theta_n$ as the estimate of the optimal solution

Let $\theta_n^* \in \Theta$ be the point with the highest estimated objective function value among solutions that have been simulated at least Mn times. if this set of solutions is empty, then $\theta_n^* = \theta_n$

R-BEESE with α =0, Mn =1, and γ =1 reduces to the R-BEES method

Definitions

- For each $\theta \in \Theta$, define $f_n(\theta)$ to be the estimate of $f(\theta)$ available at the end of iteration n(let $f_n(\theta) = -\infty$ if $C_n(\theta)$=0 ,where $C_n(\theta)$ is the number of times θ has been simulated by the end of iteration n) and $\widehat{f_k}(\theta)$ to be the estimate of $f(\theta)$ after θ has been sampled k times.

ASSUMPTION 1. *For each* $\theta \in \Theta$, $\mathbb{P}\{\lim_{k\to\infty} \hat{f}_k(\theta) = f(\theta)\} = 1$.

Assumption 1 can be easily satisfied. In the case of transient simulation, let X_θ^i be the ith observation of X_θ collected by R-BEESE. Then, Assumption 1 holds with $\hat{f}_k(\theta) = \sum_{i=1}^{kr} h_\theta(X_\theta^i)/kr$ provided that $X_\theta^1, X_\theta^2, \ldots$ are independent random elements with the distribution of X_θ and $\mathbb{E}[|h_\theta(X_\theta)|] < \infty$ (this follows from the strong law of large numbers).

THEOREM 2. *Suppose that Assumption* 1 *holds,* Θ *is finite, and* $M_n = o(n)$. *Also, assume that* $G(\{\theta\}) \geq \epsilon > 0$ *for all* $\theta \in \Theta$. *Then, the R-BEESE method converges almost surely to the set of optimal solutions* $\Theta^* = \{\theta \in \Theta\colon f(\theta) \geq f(\theta') \text{ for all } \theta' \in \Theta\}$

Note that Theorem 2 does not guarantee the effectiveness of R-BEESE. This is not surprising given the generality of the assumptions (e.g., no assumptions are made about local sampling, even though the chosen approach will likely impact performance in a major way). Consequently, Theorem 2 provides sufficient conditions under which R-BEESE is valid, leaving much flexibility for algorithm development aimed at achieving effective performance (for some efforts in that direction; see §5). The assumptions on $\{M_n\}$ and G in Theorem 2, for example, are satisfied when G is uniform on Θ and $M_n = \sqrt{n}$ for all n.

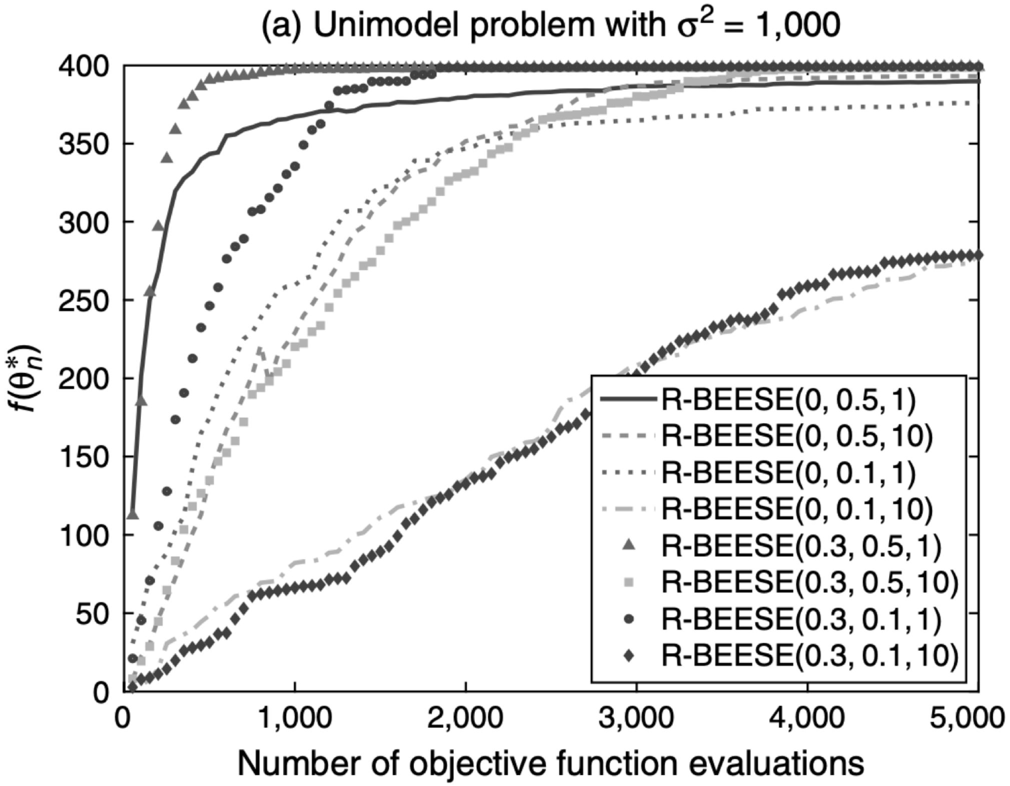

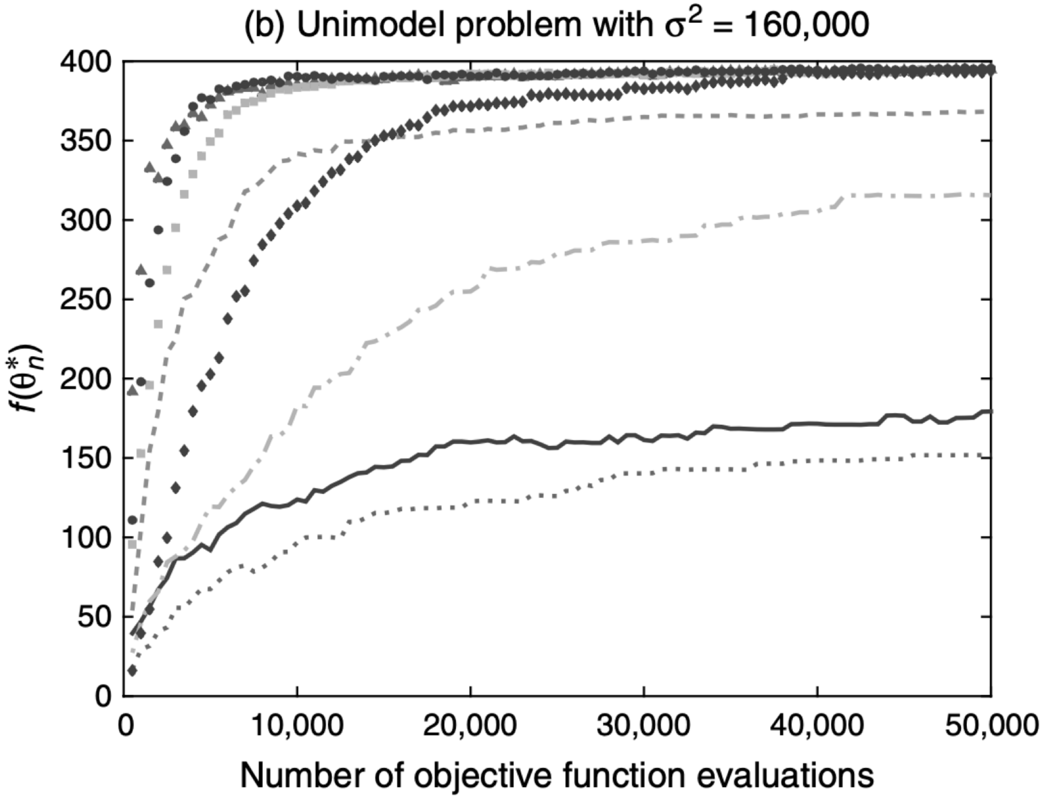

(α, p, r) : α is the most important!

A-BEES

- R-BEES samples randomly either from local or global distributions , A-BEES adaptively alternates between sampling from local or global distributions, with the goal of using the “appropriate” type of distribution at each stage of the search.

A-BEES

- After sampling k points since the last review (decision about the nature of the search),a decision is made about whether the next k sampled points will be selected using local or global distributions. Let ν^* be the function value of the best solution θ_n found so far and v_l^* be the function value at the best point found the last time a local search was performed. Let Δ be the improvement in the function value between the current and preceding reviews and D be the distance between the points where the corresponding function values were achieved.

Algorithm 2 (Sampling Distribution Update Procedure for A-BEES)

```
if LocalSearch = true then
    if Δ ≤ δ then
        LocalSearch ← false
        v_l* ← v*
    end if
else
    if Δ ≤ δ then
        if v* − v_l* ≥ δ then
            LocalSearch ← true
        end if
    else
        if D ≤ d then
            LocalSearch ← true
        end if
    end if
end if
```

local → global :

1. the improvement Δ in the objective function value between successive reviews is small. Usually this means that the local search has identified a near-local optimal solution, and hence, there is little merit in continuing searching locally.

global → local:

1. the improvement Δ is small but substantial improvement in the objective function value has been achieved since the last switch from a local to a global search. This means that A-BEES has identified a promising region , and the global search is not yielding substantial progress.

2. Δ is large but the distance D is small . This makes sense because the improvement has been local in nature, and hence, a local search may be preferable.

Algorithm 3 (A-BEES Algorithm)

counter $\leftarrow 0$, $n \leftarrow 0$
LocalSearch $\leftarrow$ false
Sample a solution θ from the global distribution G
Evaluate the objective function at θ
Let $v^*, v_l^* \leftarrow f(\theta)$ and $\theta_0 \leftarrow \theta$
while Stopping criterion is not satisfied **do**
    if LocalSearch $=$ true **then**
        Sample a solution θ from a local distribution in $\mathscr{L}$
    else
        Sample a solution θ from the global distribution G independent of everything else
    end if
    Evaluate the objective function f at θ (if needed)
    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 2)
        counter $\leftarrow 0$
    end if
    $n \leftarrow n+1$
end while
Present $\theta_n^* = \theta_n$ as the estimate of the optimal solution

Our convergence result for the A-BEES method is given below.

THEOREM 3. *Under the conditions of Theorem* 1, *we have* $\lim_{n\to\infty} f(\theta_n^*) = f^*$ *almost surely.*

A-BEESE

- Suppose that two successive reviews occur in iterations n_1 and n_2 , where $n_1 < n_2$. Then , D is the distance from θ_{n_1} to θ_{n_2} ,$\Delta = f_{n_2}\left(\theta_{n_2}\right) - f_{n_2}(\theta_{n_1})$, $v^* = f_{n_2}\left(\theta_{n_2}\right)$, and $v_l^* = f_{n_2}(\theta_l)$, where l is the last iteration number in which a local search was performed.

Algorithm 3 (A-BEES Algorithm)

counter $\leftarrow 0$, $n \leftarrow 0$
LocalSearch $\leftarrow$ false
Sample a solution θ from the global distribution G
Evaluate the objective function at θ
Let $v^*, v_l^* \leftarrow f(\theta)$ and $\theta_0 \leftarrow \theta$
while Stopping criterion is not satisfied **do**
  if LocalSearch $=$ true **then**
    Sample a solution θ from a local distribution in $\mathscr{L}$
  else
    Sample a solution θ from the global distribution G independent of everything else
  end if
  Evaluate the objective function f at θ (if needed)
  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 2)
    counter $\leftarrow 0$
  end if
  $n \leftarrow n+1$
end while
Present $\theta_n^* = \theta_n$ as the estimate of the optimal solution

With probability $0 \leqslant \alpha < 1$, the point θ_n with the highest estimated objective function value is sampled.

The optimal solution θ_n^* is estimated as in the R-BEESE method.

Algorithm 2 (Sampling Distribution Update Procedure for A-BEES)

```
if LocalSearch = true then
    if Δ ≤ δ then
        LocalSearch ← false
        v_l* ← v*
    end if
else
    if Δ ≤ δ then
        if v* − v_l* ≥ δ then
            LocalSearch ← true
        end if
    else
        if D ≤ d then
            LocalSearch ← true
        end if
    end if
end if
```

Conduct a local (global) search for k_l (k_g) iterations before attempting to switch to a global (local) search (by invoking Algorithm 2). Typically, the parameters k_l and k_g satisfy $k_l \geqslant k_g$.

local $\longrightarrow$ global :

1. the improvement Δ in the objective function value between successive reviews is small. Usually this means that the local search has identified a near-local optimal solution, and hence, there is little merit in continuing searching locally.

global $\longrightarrow$ local:

1. the improvement Δ is small but substantial improvement in the objective function value has been achieved since the last switch from a local to a global search. This means that A-BEES has identified a promising region , and the global search is not yielding substantial progress.

2. Δ is large but the distance D is small . This makes sense because the improvement has been local in nature, and hence, a local search may be preferable.

3. if a global search has been conducted for g consecutive reviews.

Theorem 4. *Under the conditions of Theorem 2 (including the assumption that* Θ *is finite), the A-BEESE method converges almost surely to the set of optimal solutions* Θ^*.

Numerical examples

Unimodal problem:

$$f(\theta_1, \theta_2) = \max\{0, -(\theta_1 - 30)^2 - (\theta_2 - 30)^2 + 400\}$$

and $\Theta = \{(i, j) \in \mathbb{N}^2\colon\ 0 \le i, j \le 199\}$.

Two-hills problem:

$$f(\theta) = \max\{f_1(\theta), f_2(\theta), 0\},$$

where $f_1(\theta) = -(0.4\theta_1 - 5)^2 - 2(0.4\theta_2 - 17.2)^2 + 7$ and $f_2(\theta) = -(0.4\theta_1 - 12)^2 - (0.4\theta_2 - 4)^2 + 4$. The feasible space $\Theta = \{\theta = (\theta_1, \theta_2) \in \mathbb{N}^2\colon\ 0 \le \theta_1, \theta_2 \le 49\}$

Buffer allocation problem:

$$\Theta = \Big\{(\theta_1, \ldots, \theta_5) \in \mathbb{N}^5\colon\ \theta_1 + \theta_2 + \theta_3 \le 20;\ \theta_4 + \theta_5 = 20;\ 1 \le \theta_k \le 20 \text{ for } k = 1, \ldots, 5\Big\}$$

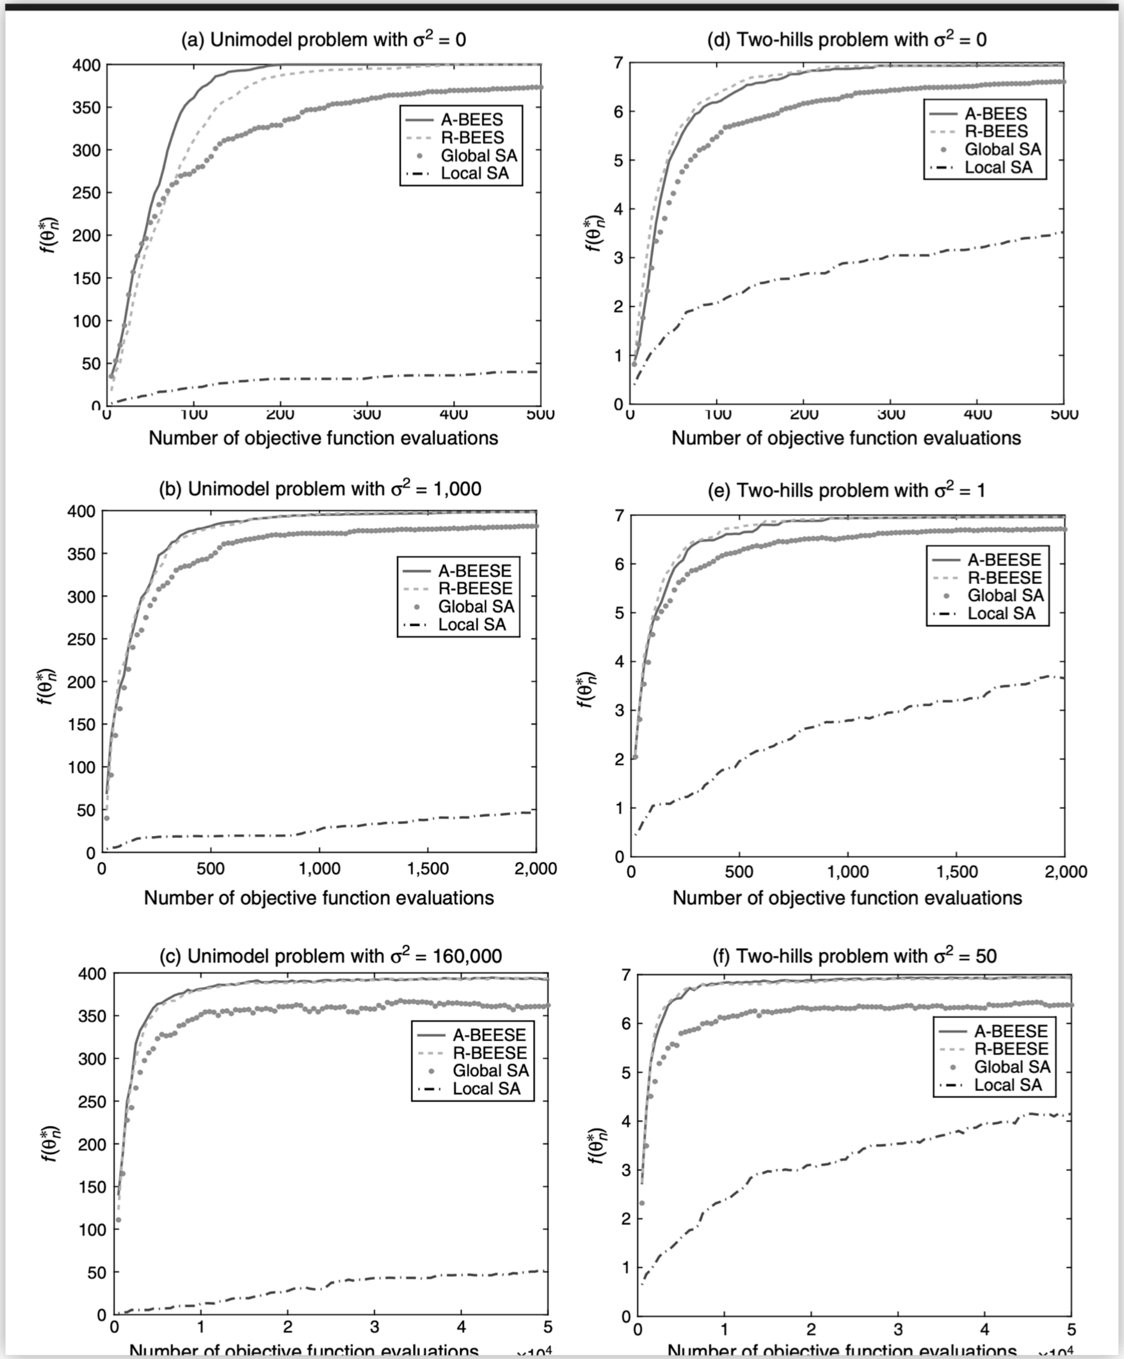

A-BEES is considerably better than R-BEES when σ^2= 0 and that A-BEESE has similar performance to R-BEESE when $\sigma^2 \in \{1000,16000\}$.

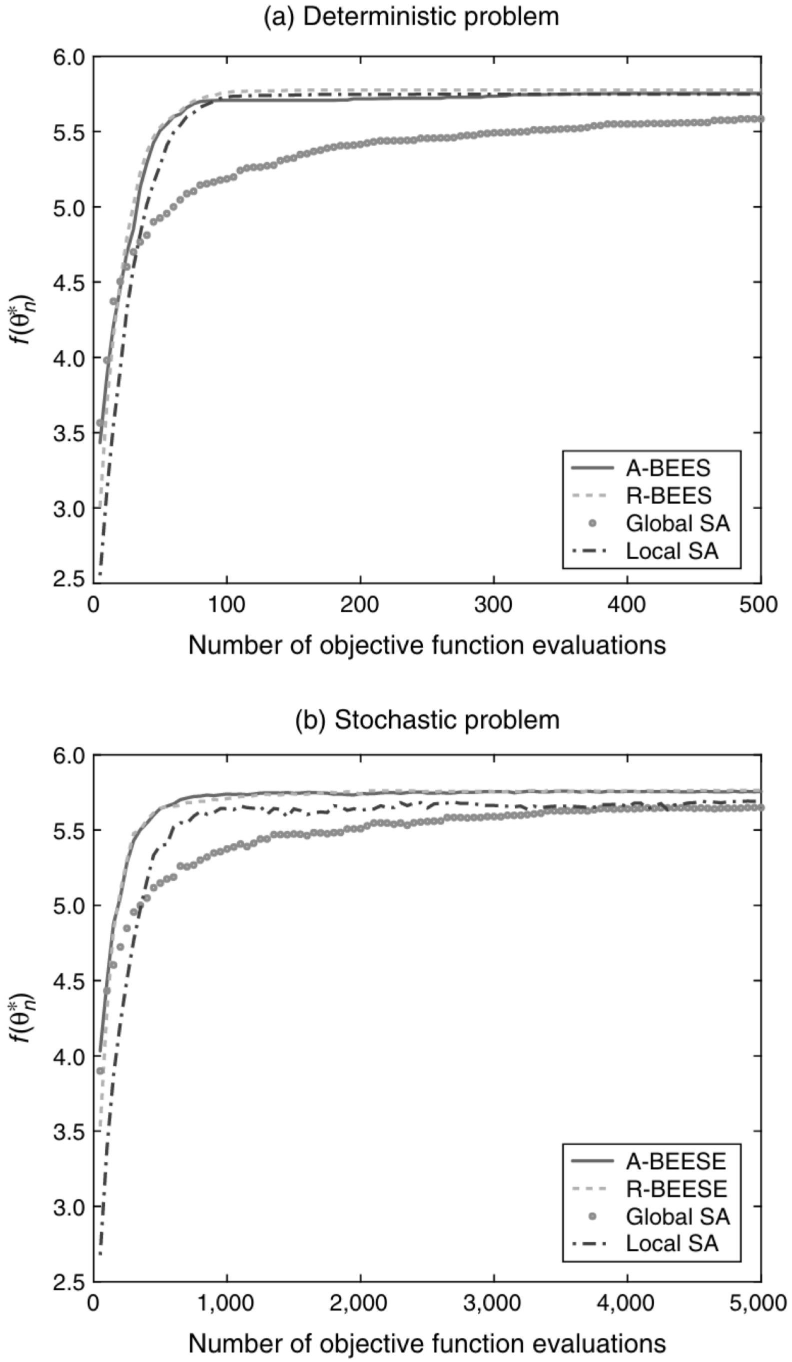

A-BEES(E) and R-BEES(E) perform similarly, and they outperform both SA algorithms.

- The convergence of each method slows down as the noise increases. However, the relative performance of the methods does not depend heavily on the noise level . Moreover , the difference in the empirical performance of the R-BEESE and A-BEESE methods becomes smaller as σ^2 grows.
- In practice , one may not have the luxury of identifying good parameter settings, and hence, the robustness of algorithms to parameter values is important. Our experience is that R-BEESE is quite robust to parameter value.

Thanks!