ICML 2022

Generalizing Gaussian Smoothing for Random Search

Katelyn Gao and Ozan Sener

intel®

Derivative-free Optimization

- In many real-world applications, analytical gradient of the loss function is expensive to compute
  - Ex: search and rescue robot on complex terrain
- But evaluations of the loss are cheap
- DFO: optimize objective $F(\theta)$ only using zeroth order (noisy) evaluations ($\mathrm{f}(\theta, \xi)$)

Gaussian Smoothing (GS)

- GS: estimate gradient via forward difference, using evaluations at randomly perturbed parameters

$$\nabla_\theta F^{FD}(\theta) = \frac{1}{c}(F(\theta + c\epsilon) - F(\theta))\epsilon, \qquad \epsilon \sim \mathcal{N}(0, I), c > 0$$

- Under regularity conditions, converges to stationary point if $c \to 0$
- Average over L perturbations

Generalizing GS

- Main idea: we can sample perturbations from arbitrary distributions to optimize a criterion
- Proposal: select distribution that reduces gradient estimate MSE of forward difference estimator w/ noisy evaluations:

$$\nabla_\theta \hat{F}^{FD}(\theta) = \frac{1}{cL} \sum_{l=1}^{L} \sum_{i=1}^{N} \left(f(\theta + c\epsilon_l, \xi_i) - f(\theta, \xi_i) \right) \epsilon_l$$

- Algorithms have same computational complexity as GS and do not depend on characteristics of objective
  - BeS: ϵ_l standardized Bernoulli with prob. 0.5
  - GS-shrinkage and BeS-shrinkage: decrease variances of GS and BeS by scaling

Theoretical motivation

- Theorem (informal): Assume gradient estimates have at most MSE M and bias $\mathrm{B}\nabla_\theta F(\theta)$. After T steps of SGD,

$$\frac{1}{\mathrm{T}}\sum_t \|\nabla_\theta F(\theta^t)\|_2^2 \leq \frac{M + O_F(1)}{(1-2B)\sqrt{T}}$$

  where $O_F(1)$ depends on characteristics of $F(\theta)$
- To improve convergence, reduce M
- BeS has smaller MSE than GS
  - GS-shrinkage decreases the MSE for Gaussian perturbations; similarly for BeS-shrinkage

Experiments: RL (MuJoCo)

Linear policy, L=20, N=1

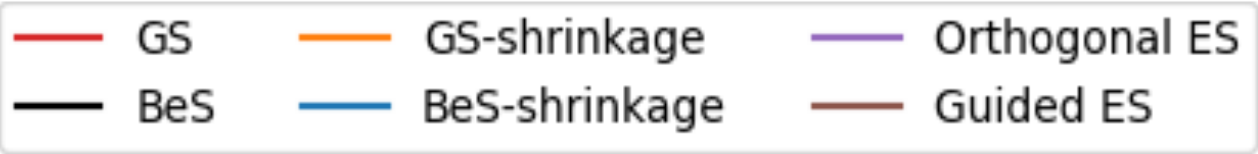

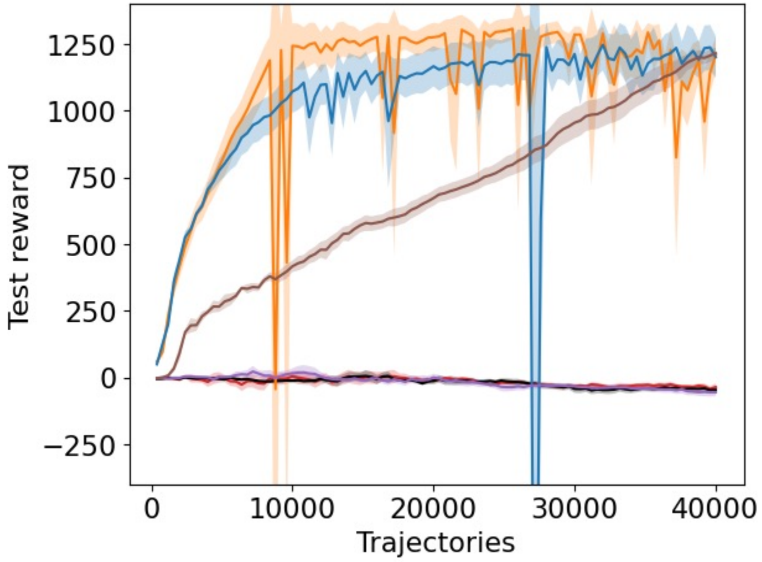

Ant

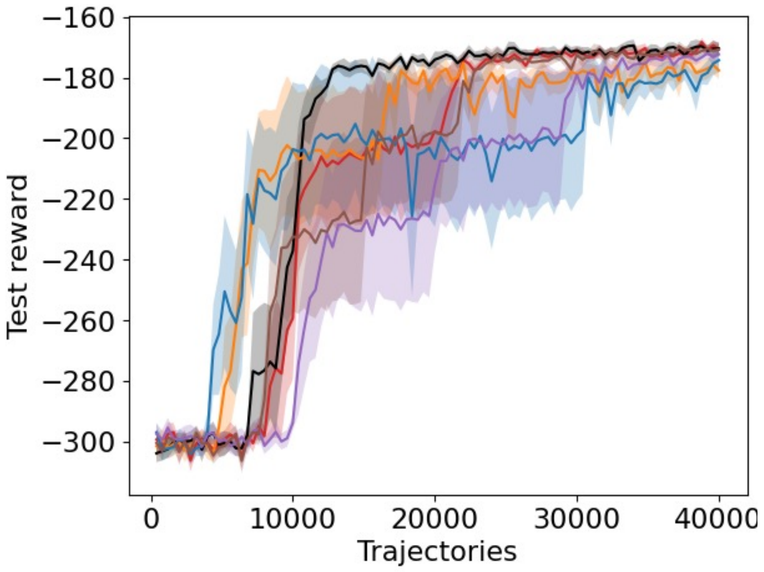

Half Cheetah Random Velocity

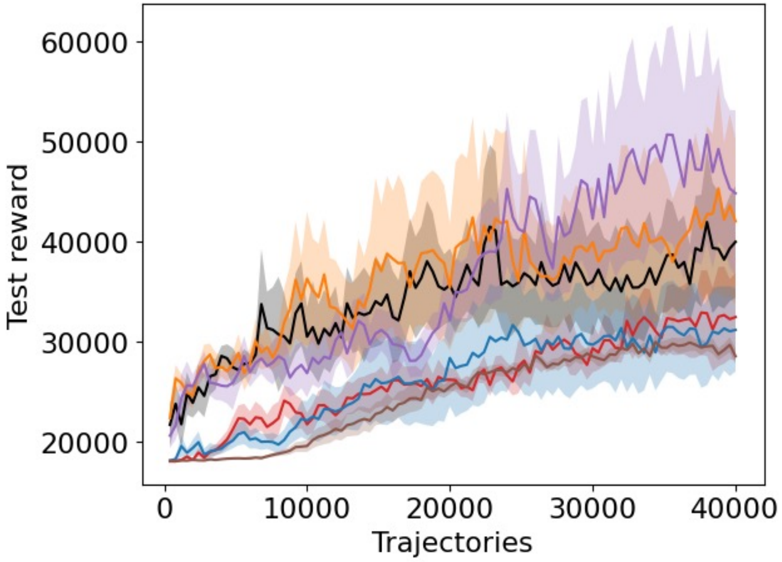

ML1-Reach

Experiments: DFO (Nevergrad)

d=100, L=10, N=1

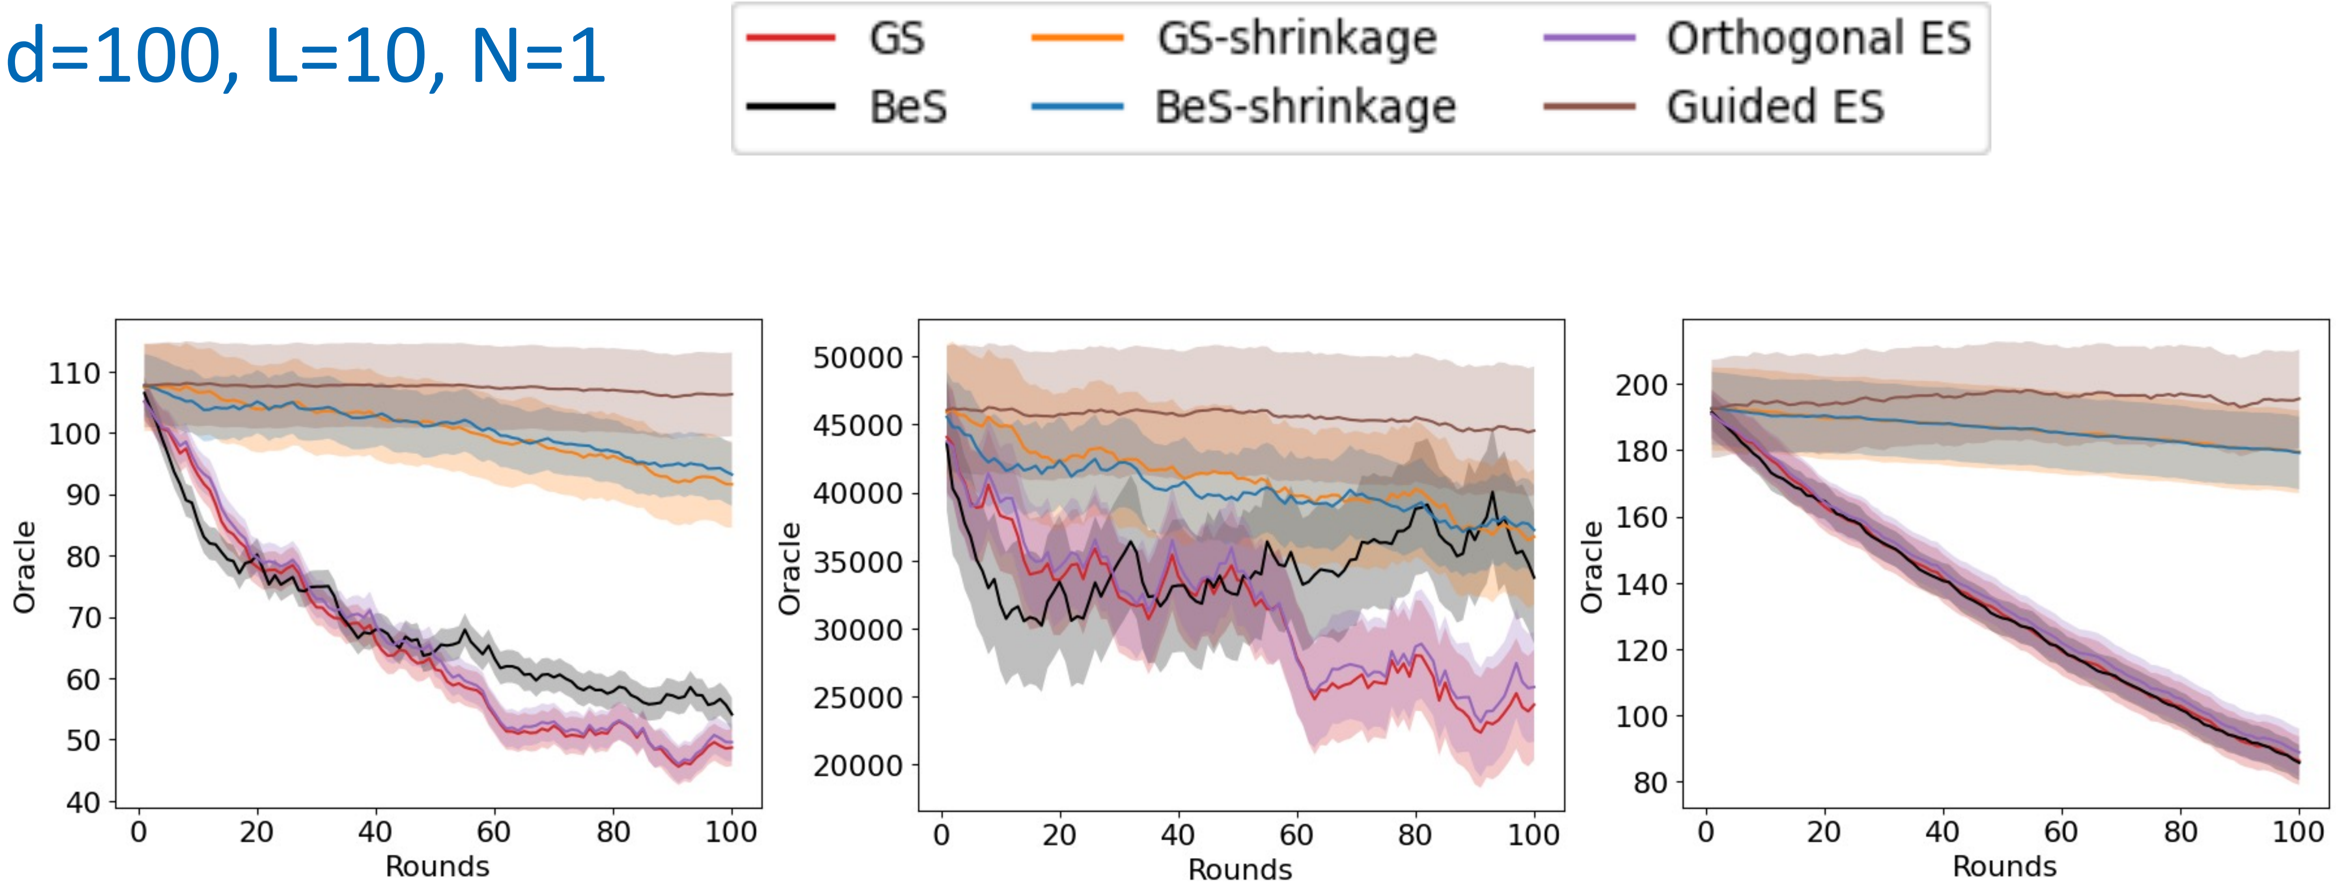

sphere rosenbrock hm

Thank you for your attention!