Quantile Stein Variational Gradient Descent for Batch Bayesian Optimization

Chengyue Gong[1] Jian Peng[2] Qiang Liu[1]

[1] The University of Texas at Austin
[2] University of Illinois at Urbana-Champaign

Bayesian Optimization

- **Goal: black-box optimization**

$$\max_x f(x), \qquad f(\cdot)\text{: expensive, black-box function.}$$

- **Bayesian Optimization:**

Iteratively acquire new points based on an acquisition function:

$$x^{new} \leftarrow \arg\max_x \alpha(x \mid \mathcal{D}),$$

$$\mathcal{D}^{new} \leftarrow \mathcal{D} \cup \{x^{new},\ f(x^{new})\},$$

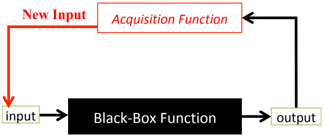

Acquisition function:

$$\alpha(x \mid \mathcal{D}) := \mathbb{E}_f[f(x) \mid \mathcal{D}] + \eta\sqrt{\mathrm{var}_f[f(x) \mid \mathcal{D}]}. \qquad (UCB)$$

- **Batch Bayesian Optimization**:

  Find multiple query points $\{x_i\}_{i=1}^m$ in parallel at every iteration.

- Much more challenging; two desiderata:
  - **Diversity**: Everyone should be unique.
  - **Qualification**: Everyone should be good.

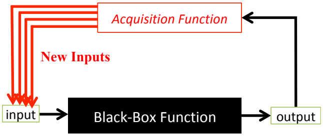

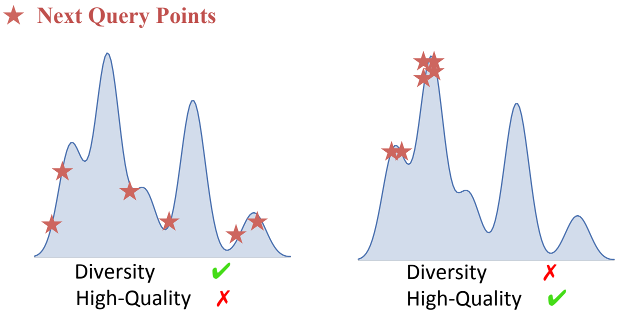

- Optimizing the distribution ρ of query points $\{x_i\}$ by

$$\max_{\rho} \left\{ F[\rho] := \mathbb{E}_{\rho}^{\omega}[\alpha(x)] \ + \ \eta H[\rho] \right\}.$$

- $H[\rho]$ is the entropy. It encourages the **diversity**.
- $\mathbb{E}_{\rho}^{\omega}[\cdot]$ is a quantile distorted expectation. It enforces **qualification**,

$$\mathbb{E}_{\rho}^{\omega}[\alpha(x)] = \int_0^1 Q_{f,\rho}^{\beta}\omega(\beta)d\beta,$$

  - $Q_{f,\rho}^{\beta}$ is the β-th quantile of $\alpha(x)$, when $x \sim \rho$.
  - $\omega \colon [0,1] \to \mathbb{R}_+$ is a distortion function:

  - **Risk neutral**: $\omega(\beta) = 1$.
  - **Risk aversion**: $\omega(\beta)$ is monotonic decreasing.
  - **Risk seeking**: $\omega(\beta)$ is monotonic increasing.

We want risk aversion: Take $\omega(\beta) = \beta^{-\lambda}$, where $\lambda \geq 0$.

Quantile Stein Variational Gradient Descent [Liu, Wang 16]

Idea: Find particle distributions

$$\rho := \sum_{i=1}^{n} \delta_{x_i}/n$$

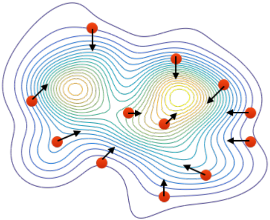

to approximately solve the optimization.
The particles $\{x_i\}_{i=1}^n$ are iteratively moved to maximize the objective by gradient-like updates

$$x_i' \leftarrow x_i + \epsilon \phi^*(x_i), \qquad \phi^* = \arg\max_{\phi \in \mathcal{H}} \left\{ \frac{d}{d\epsilon} F[\rho'] \Big|_{\epsilon=0} \;\; s.t. \;\; ||\phi||_{\mathcal{H}} \leq 1 \right\},$$

ϵ: step-size; ϕ^*: chosen to maximize the objective function as fast as possible. $\mathcal{H}$: a reproducing kernel Hilbert space (RKHS) with positive definite kernel $k(x, x')$.

Quantile Stein Variational Gradient Descent [Liu, Wang 16]

Optimization:

$$\max_{\rho} \left\{ F[\rho] := \mathbb{E}_{\rho}^{\omega}[\alpha(x)] \; + \; \eta H[\rho] \right\}.$$

Algorithm:

$$x_i \leftarrow x_i + \frac{\epsilon}{n} \sum_{i=1}^{n} [\underbrace{\xi(x_j)}_{\text{quantile}} \; \underbrace{\nabla_x \alpha(x_j) k(x_j, x_i)}_{\text{gradient}} + \underbrace{\eta \nabla_{x_j} k(x_j, x_i)}_{\text{repulsive force}}], \quad \forall i = 1, \ldots, n.$$

Here, each particle is assigned a weight to account the distortion function:

$$\xi(x_j) = \omega \left(\frac{rank(x_j)}{n} \right), \qquad rank(x_j) = \sum_{\ell=1}^{n} \mathbb{I}[\alpha(x_\ell) \le \alpha(x_j)].$$

Empirical Results

Standard Benchmarks

	LP-UCB (Gonzalez et. al., 2016)	DPP (Kathuria et. al., 2016)	MACE (Lyu et. al., 2018)	QSBO-UCB Ours
Branin	3.28e-4	9.63e-4	**2.85e-5**	5.14e-5
Eggholder	51.34	82.81	74.14	**46.86**
Dropwave	0.14	0.13	0.09	**0.07**
CrossInTray	6.83e-3	7.64e-3	3.78e-4	**1.35e-4**
gSobol5	1.85	2.34	1.14	**0.32**
gSobol10	1.04e2	1.07e3	48.92	**31.19**
gSobol15	2.34e3	5.28e3	6.39e2	**3.61e2**
Ackley5	3.71	3.74	2.36	**2.23**
Ackley10	3.87	4.23	3.01	**2.41**
Alpine2	75.92	73.39	**63.29**	73.01

Table: Negative Rewards

Empirical Results

Automatic Chemical Design (Gomez-Bombarelli et. al., 2018; Griffiths, 2017)

	LP-UCB	DPP	MACE	QSBO-UCB
QED	0.91±0.05	0.91±0.06	0.92±0.03	**0.93±0.03**
SAS	2.18± 0.06	2.29±0.08	2.16±0.04	**2.08±0.05**
LogP	0.50±0.11	0.47±0.07	0.41±0.06	**0.33±0.08**

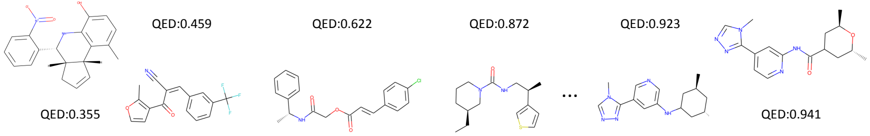

Figure: Illustration of the search process of our QSBO-UCB.

Conclusions

1. A new algorithm (QSVGD) for risk-sensitive objective
2. Risk-aversion samples for batch Bayesian optimization
3. Good empirical results

Thank You

Poster #239, Today 06:30 PM –09:00 PM @ Pacific Ballroom