

# Batched Energy-Entropy acquisition for Bayesian optimization


Felix Teufel, Carsten Stahlhut, Jesper Ferkinghoff-Borg



Bayesian optimization (BO) enables round-based optimization of black-box problems. In many application domains, it is often most efficient to conduct experiments that acquire points in parallel. However, commonly used acquisition functions are often high-dimensional and intractable in batch mode, leading to the use of sampling-based alternatives.

We propose a statistical physics inspired acquisition function that can natively handle batches. Batched Energy-Entropy acquisition for BO (BEEBO) enables tight control of the explore-exploit trade-off of the optimization process.

- Parallel gradient-based optimization of points
- No sampling and Monte Carlo integrals
- Tight control of the explore-exploit trade-off
- Risk-averse BO under heteroskedastic noise



Acquiring 100 points on a surrogate of the Ackley function (background). BEEBO enables controllable acquisition.

## The BEEBO acquisition function

$$a_{\text{BEEBO}}(\mathbf{x}) = -E(\mathbf{x}) + T \cdot I(\mathbf{x})$$

Energy-Exploit                      Entropy-Explore

### Entropy - Explore

- Assume we have a posterior probability distribution over the surrogate function  $f$  evaluated at a batch of points  $\mathbf{x}$ ,  $f(\mathbf{x}) \sim P(\mathbf{f} | D, \mathbf{x})$ . The lack of knowledge of  $f$  at  $\mathbf{x}$  is quantified by the differential entropy  $H$ :

$$H(\mathbf{f} | D, \mathbf{x}) = - \int P(\mathbf{f} | D, \mathbf{x}) \ln(P(\mathbf{f} | D, \mathbf{x})) d\mathbf{f}$$

- We can contrast  $H$  with the entropy *after* we obtain measurements at

$$H_{\text{aug}}(\mathbf{f} | D, \mathbf{x}) = \int P(y | D, \mathbf{x}) H(\mathbf{f} | D_{\text{aug}}(y)) dy$$

- Using these two terms, we can compute the *information gain*, the expected reduction in entropy.

$$I(\mathbf{x}) = H(\mathbf{f} | D, \mathbf{x}) - H_{\text{aug}}(\mathbf{f} | D, \mathbf{x})$$

- In BEEBO, we use  $I(\mathbf{x})$  as the explore component of the acquisition function. When using a GP, the Gaussian posterior covariance  $C(\mathbf{x})$ , and therefore the entropy, only depends on the positions of  $\mathbf{x}$ , *not* on the actual observed values  $y$ .

$$C(\mathbf{x}) = K(\mathbf{x}, \mathbf{x}) - K(\mathbf{x}, \mathbf{x}_D) \cdot (K(\mathbf{x}_D, \mathbf{x}_D) + \sigma^2(\mathbf{x}_D))^{-1} \cdot K(\mathbf{x}_D, \mathbf{x})$$

$$H(\mathbf{f} | D, \mathbf{x}) = \frac{Q}{2} \ln(2\pi e) + \frac{1}{2} \ln \det(C(\mathbf{x}))$$

- We can therefore compute  $I(\mathbf{x})$  in closed form by adding  $\mathbf{x}$  to the GP's training data without observing  $y$ .

### Energy - Exploit

- To quantify the optimality of points, we need to summarize the distribution  $f(\mathbf{x})$  using the expectation of a scalar function,  $\tilde{E}: \mathbb{R}^Q \rightarrow \mathbb{R}$ .
- We use a softmax weighted sum, as this allows us to smoothly interpolate between the expected mean and the expected maximum of the distribution using the softmax inverse temperature  $\beta$  (mean at 0, maximum at infinity).
- We multiply the expectation with the batch size  $Q$  so that it scales linearly, as does  $I(\mathbf{x})$ .

$$E(\mathbf{x}) = -\mathbb{E}[\tilde{E}(\mathbf{x})] \cdot Q = -\mathbb{E} \left[ \sum_{q=1}^Q \text{softmax}(\beta \mathbf{f}_q) f_q \right] \cdot Q$$

- The softmax-weighted sum of a multivariate normal is not available in closed form. We apply Taylor expansion to derive a fully differentiable closed form approximation.

## Availability

```
!pip install beebo
from beebo import BatchedEnergyEntropyBO
from torch.optim.optimizer import Optimizer

batch_size = 96
acq_fn = BatchedEnergyEntropyBO(model, temperature = 1.0)
points, value = optimize_acqf(acq_fn, bounds, batch_size, 10, 100)
```

## BEEBO with GPs

```
Input: model GP, initial batch points x, temperature T
repeat
  Calculate mu(x), C(x) using GP
  E ← -softmax_expectation(mu(x), C(x), beta)
  GP_aug ← augment(GP, x)
  Calculate I_aug(x) using GP_aug
  I ← 1/2 ln det(C(x)) - 1/2 ln det(C_aug(x))
  a ← -E + T * I
  x ← x + gamma * a
until converged
Output: optimized batch points x
```

**Optimizing BEEBO with GPs**  
All terms in BEEBO are fully differentiable. To maximize the acquisition function, we perform joint gradient descent on all batch points simultaneously.

Low-rank updates can be used to add points to (augment) the posterior covariance.

## Benchmark experiments

- 10 rounds, batch size 100
- Final batch at full exploit (T=0)
- 26 test problems
- 10 replicates
- Compare to  $q$ -UCB at equal explore-exploit rates
- Additional comparisons to default  $q$ -EI, Thompson sampling (TS), Kriging Believer (KB), GIBBON and TuRBO

### BO on test problems

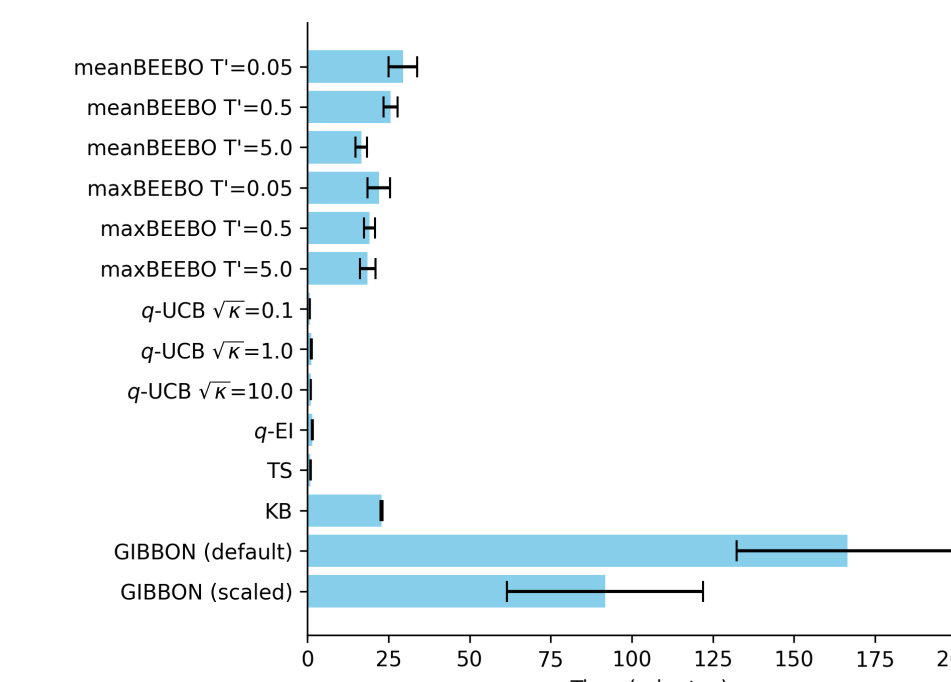
Problem	#	meanBEEBO	q-UCB	TS	KB	GIBBON	TuRBO
Ackley	2	0.901 ± 0.005	0.985 ± 0.011	0.975 ± 0.015	0.982 ± 0.023	0.980 ± 0.035	0.988 ± 0.013
Ley	2	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	2	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	4	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	8	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	10	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	10	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	10	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	10	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	10	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	20	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	20	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	20	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	20	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	20	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	50	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	50	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	50	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	50	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	50	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	100	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	100	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	100	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	100	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	100	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Mean		0.905	0.928	0.928	0.928	0.928	0.928
Median		0.905	0.928	0.928	0.928	0.928	0.928

**Optimization benchmark.** Best observed value after 10 rounds of batched BO. BEEBO outperforms  $q$ -UCB at equal trade-offs.

Problem	#	meanBEEBO	q-UCB	TS	KB	GIBBON	TuRBO
Ackley	2	0.901 ± 0.005	0.985 ± 0.011	0.975 ± 0.015	0.982 ± 0.023	0.980 ± 0.035	0.988 ± 0.013
Ley	2	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	2	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	4	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	8	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	10	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	10	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	10	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	10	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	10	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	20	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	20	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	20	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	20	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	20	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	50	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	50	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	50	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	50	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	50	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Ackley	100	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Ley	100	0.981 ± 0.024	0.989 ± 0.016	0.985 ± 0.016	0.983 ± 0.007	0.983 ± 0.011	0.993 ± 0.006
Rastrigin	100	0.999 ± 0.002	0.996 ± 0.001	0.999 ± 0.001	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Shubert	100	0.810 ± 0.010	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001	0.988 ± 0.001
Coste	100	0.945 ± 0.115	0.914 ± 0.041	0.922 ± 0.063	0.923 ± 0.063	0.923 ± 0.063	0.923 ± 0.063
Mean		0.905	0.928	0.928	0.928	0.928	0.928
Median		0.905	0.928	0.928	0.928	0.928	0.928

**Controllability benchmark.** Batch instantaneous regret at round 10 (acquired at T=0, full exploit). Hyperparameter-free baselines are run at default.

### Runtime



BEEBO's optimization runtime is competitive with iterative approaches such as KB and greatly improves upon GIBBON. Reparametrization trick methods (q-UCB, q-EI) are orders of magnitude faster at a cost of MC integration accuracy. All methods were run in BoTorch.

## Summary

- Competitive performance to existing sampling-based or greedy heuristic batched BO methods
- Trade-off hyperparameter has predictable behaviour regardless of batch size
- Information gain enables risk-averse BO under heteroskedastic noise (sensitive to good surrogate for  $\sigma$ ).

## Outlook

- More memory-efficient predictive covariances (GP cubic scaling is constraining larger-scale BO on GPU)
- Generalization to multi-objective BO (vector-valued energy term)
- Optimal scheduling of the temperature
- Information gain approximations for non-GP surrogate models

