Efficient Black-Box Combinatorial Optimization

Hamid Dadkhahi Karthikeyan Shanmugam
Jesus Rios Payel Das

IBM Research NY

RealML @ ICML2020

Overview

Black-box function optimization over purely categorical variables

The black-box functions of interest:

- Intrinsically expensive to evaluate
- Noisy
- No trivial means to find the minimum

Problem Statement

Problem: Given the categorical domain $\mathcal{X} = [k]^n$, with n variables each of cardinality k, the objective is to find

$$x^* = \arg\min_{x \in \mathcal{X}} f(x)$$

where $f : \mathcal{X} \mapsto \mathbb{R}$ is a real-valued combinatorial function.

- Exhaustive search infeasible in practice
- Find x^* (or an approximation of it) in as few function evaluations as possible

Learning Framework

Learning framework at each time step t:

- *Surrogate model* provides an estimate for the black-box function via observations $\{(x_i, f(x_i)) : i \in [t]\}$ seen so far.
- *Acquisition function* selects a new candidate point x_t.
- The black-box function returns the evaluation $f(x_t)$.

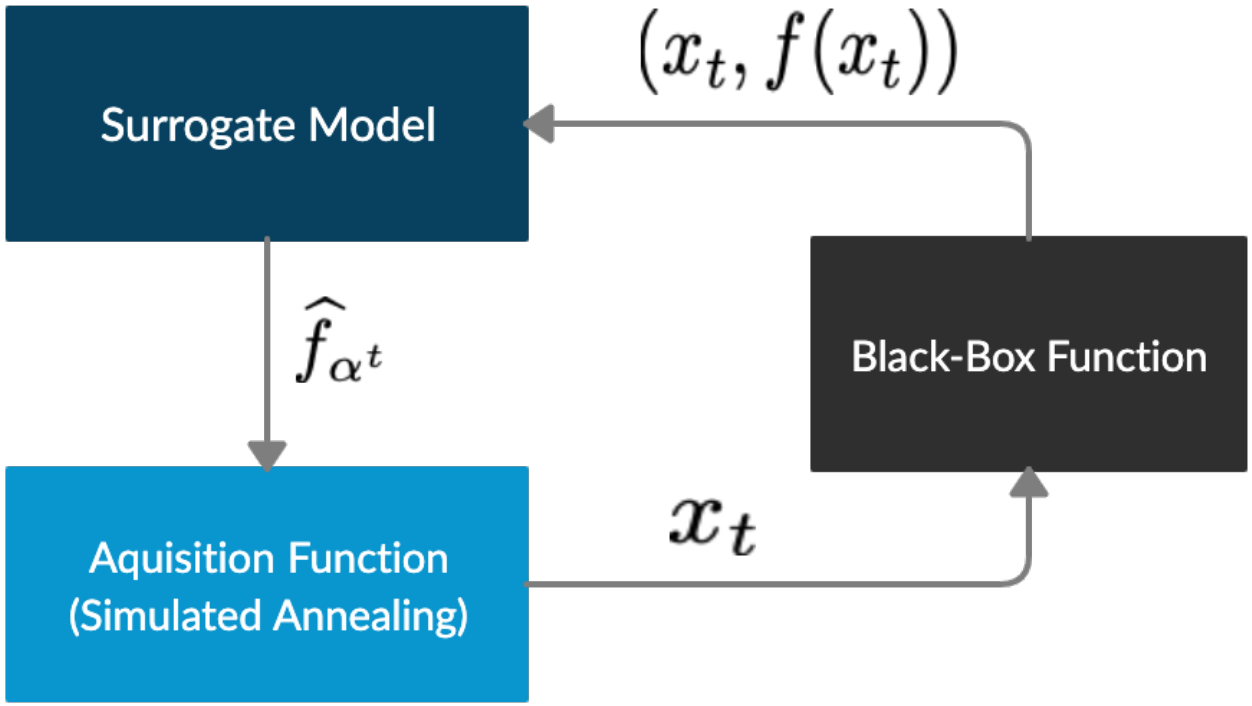

Surrogate Model

Boolean Case: Multilinear Polynomial Representation (Fourier expansion)

$$f(x) = \sum_{\mathcal{I} \subseteq [n]} \alpha_{\mathcal{I}} \psi_{\mathcal{I}}(x)$$

- $\alpha_{\mathcal{I}}$: Fourier coefficient of f on $\mathcal{I}$
- $\psi_{\mathcal{I}}(x) = \Pi_{i \in \mathcal{I}} x_i$: monomials of order $|\mathcal{I}|$

Categorical Case: Fourier representation on finite Abelian groups

$$f(x) = \sum_{\mathcal{I} \in [k]^n} \alpha_{\mathcal{I}} \psi_{\mathcal{I}}(x)$$

- $\alpha_{\mathcal{I}}$: Fourier coefficients
- $\psi_{\mathcal{I}}(x) = \exp(2\pi j \langle x, \mathcal{I} \rangle / k)$: characters ($k$-th roots of unity)

The ECO Algorithm

Surrogate Model Update Rule:

- ▷ Exponential weight update rule from the Hedge algorithm
- ▷ We maintain a pool of monomials (Boolean case) or characters (categorical case) where each term plays the role of an expert
- ▷ Find the optimal coefficient α_i for *expert* ψ_i.

Acquisition Function: A version of simulated Annealing

Results: RNA Sequence Optimization Problem

- RNA sequence as a string $A = a_1 \dots a_n$ of n letters (nucleotides) over the alphabet $\Sigma = \{A, U, G, C\}$
- Given a sequence length n, find a sequence with Minimum Free Energy (MFE)
- Experiments: RNA sequences of length $n = 30$

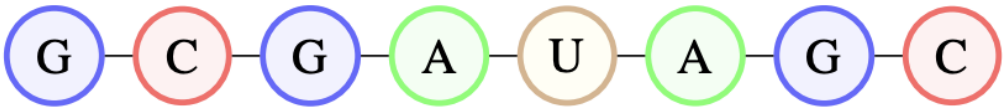

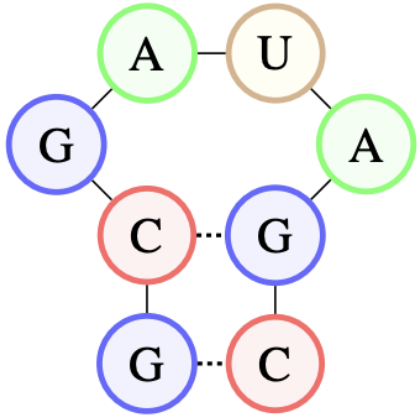

Results: RNA Sequence Optimization Problem

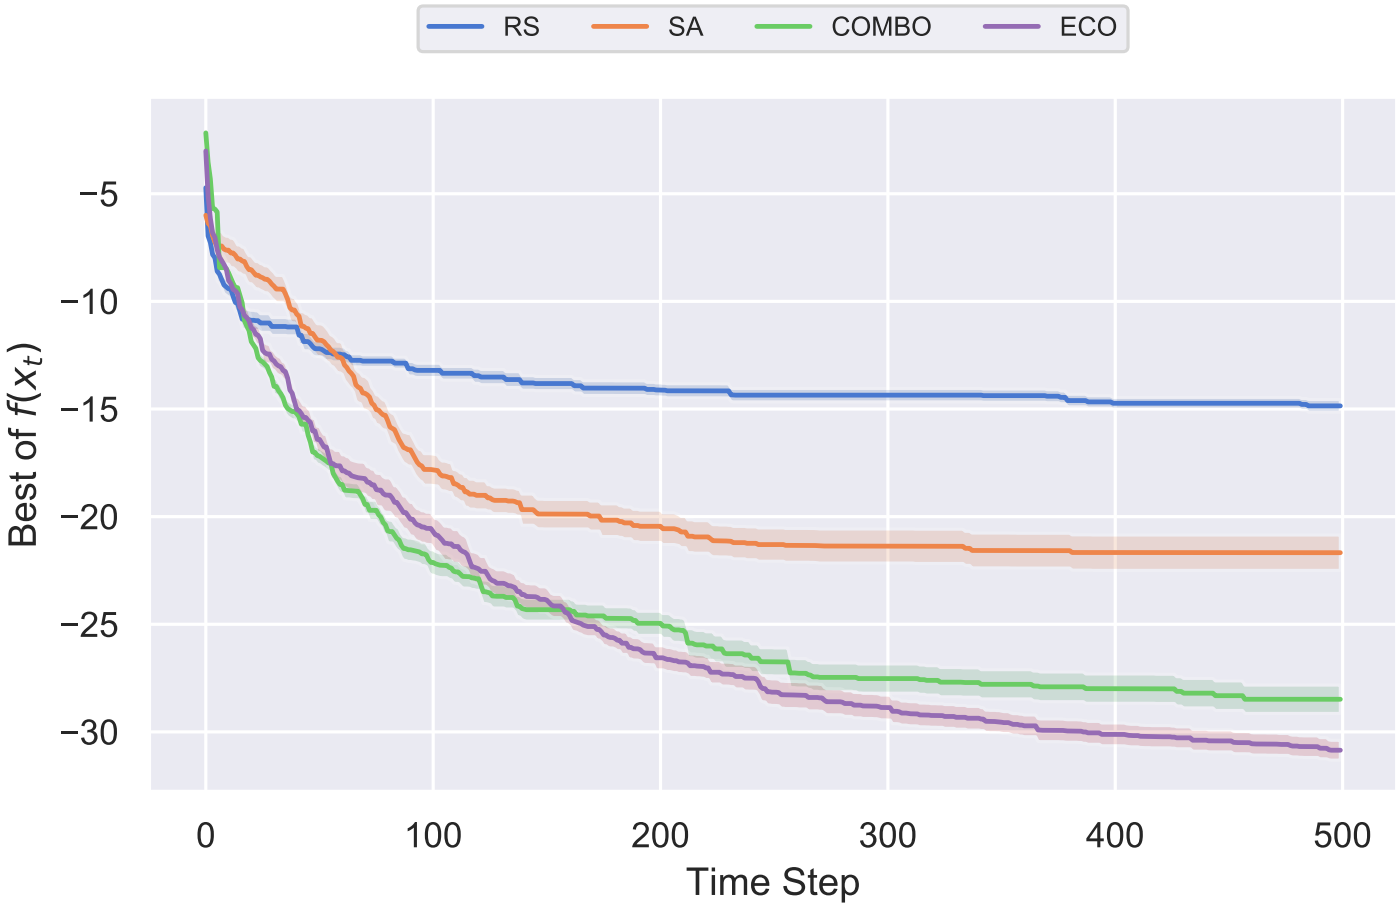

Results: Computation Times

Average computation time per step (in Seconds)

Dataset	n	k	COMBO	ECO
Sequence Optimization	30	4	253.8	5.7

Thank you!

Questions?

hdadkhahi@ibm.com
karthikeyan.shanmugam2@ibm.com
jriosal@us.ibm.com
daspa@us.ibm.com