Generative Boltzmann Optimization via Conditional Diffusion Models for Black-Box Optimization

Georgia Tech College of Engineering
H. Milton Stewart School of Industrial and Systems Engineering

Kaitao Liao · Advisor: Enlu Zhou

MAIN TAKEAWAY Low-cost pseudo-labels guide the diffusion model toward the global optimum with only a few real function evaluations.

BACKGROUND

Black-box optimization searches for the best solution when the relationship between inputs and outcomes is unknown. We can test different inputs and observe their objective values, but we cannot see how those values are calculated.

The goal is to find the best solution with as few costly evaluations as possible.

CANDIDATES → BLACK BOX → SCORES

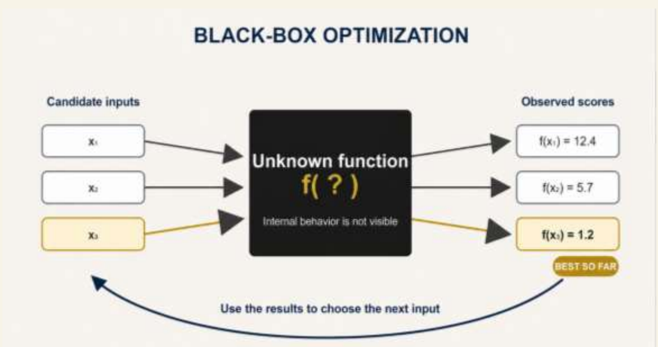

Keep the highest score

EXPERIMENTAL SETUP

Fair comparison across algorithms

Same 50 random seeds · Matched initial points · Equal real-evaluation budgets

GBO-CD configuration

30D
40 iterations
3 layers · 128 neurons
200 epochs

100D
60 iterations
5 layers · 256 neurons
200 epochs

Shared GBO-CD settings
5,000 pseudo-samples per iteration · 50 new evaluations per iteration
Pseudo-samples require no additional objective evaluations.

HOW TO READ THE PLOTS

X-axis: Optimization iteration
Y-axis: Mean objective value
Colors: Blue = GBO-CD, Orange = CE, Green = PSA
Solid line: Mean across 50 independent seeds
Shaded band: 95% confidence interval
Dashed line: Known global optimum

Higher objective values indicate better solutions.

SCAN FOR DETAILS

SUPPLEMENTARY MATERIALS

Benchmark definitions
Baseline explanations
Experimental setup
Theory and limitations

GBO-CD WORKFLOW

1. **Evaluate initial points** — Start with 50 random points and calculate their true objective values.
2. **Create a pseudo-sample pool** — Sample 5,000 points uniformly across the search space without evaluating their true objective values.
3. **Assign pseudo-labels** — Use the nearest evaluated point to estimate each pseudo-sample's objective value.
4. **Temperature-based Boltzmann weights** — Convert pseudo-labels into temperature-controlled weights that favor better estimated points.
5. **Conditional diffusion model training** — Train the model with pseudo-samples and their Boltzmann weights to learn where higher-value points are located.
6. **Generate new candidate points** — Create 50 new points in regions with higher estimated objective values.
7. **Update the evaluated dataset** — Calculate their true objective values and add them to the evaluated dataset for the next round.

Each round adds new evaluated results, helping the model search better regions as the temperature decreases.

RESULTS ACROSS OBJECTIVE FUNCTIONS

● GBO-CD ● CE ● PSA – – – Global optimum

Ackley · 30D

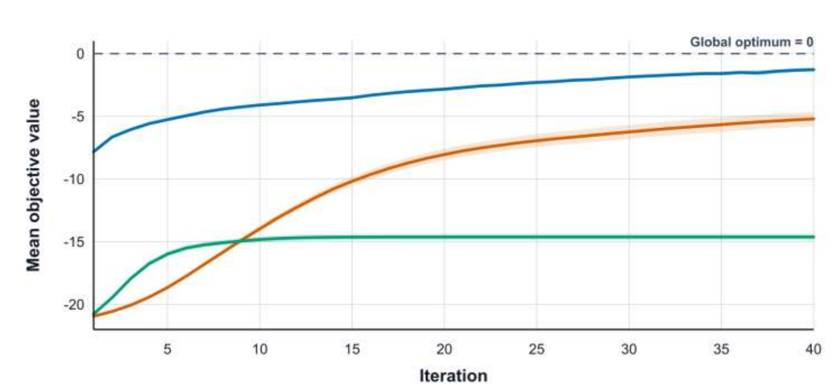

Alpine · 30D

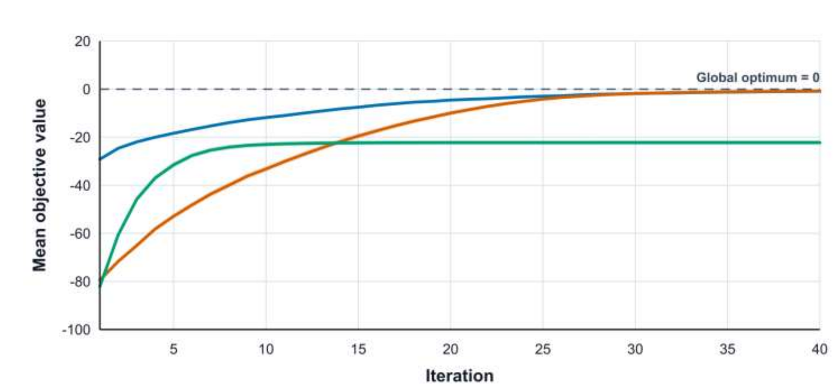

Lévy · 30D

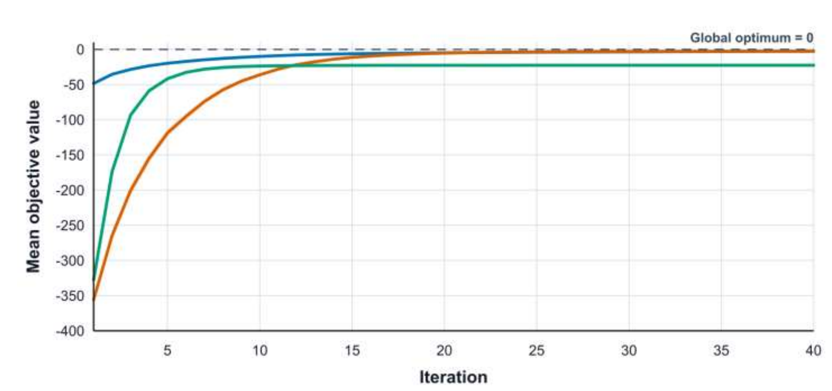

Pintér · 30D

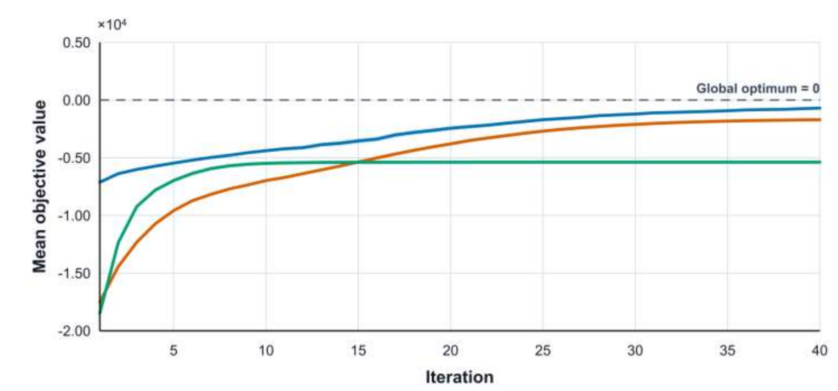

Ackley · 100D

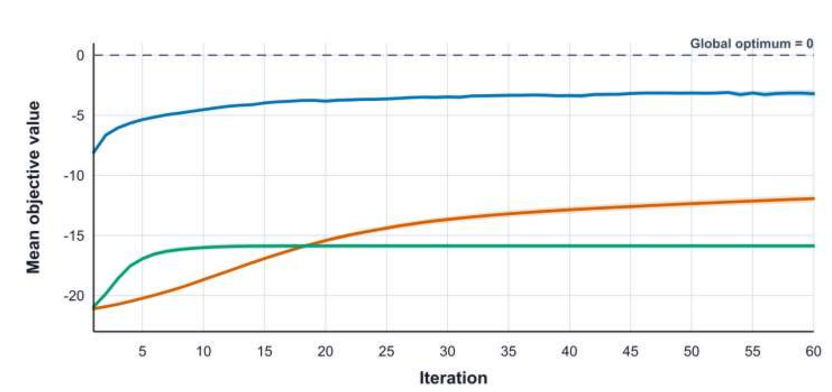

Alpine · 100D

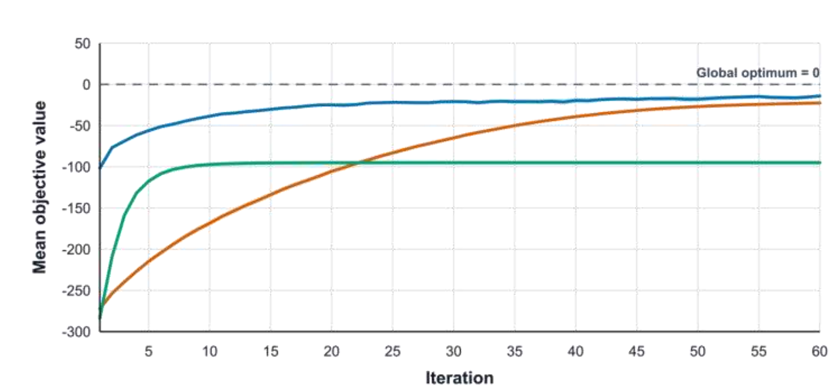

Lévy · 100D

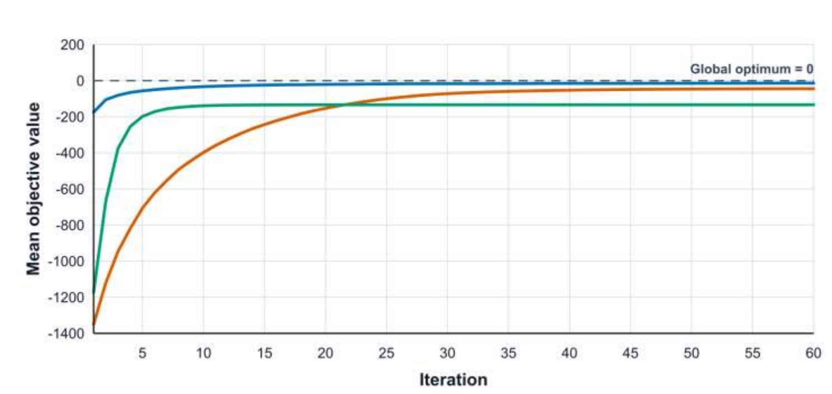

Pintér · 100D

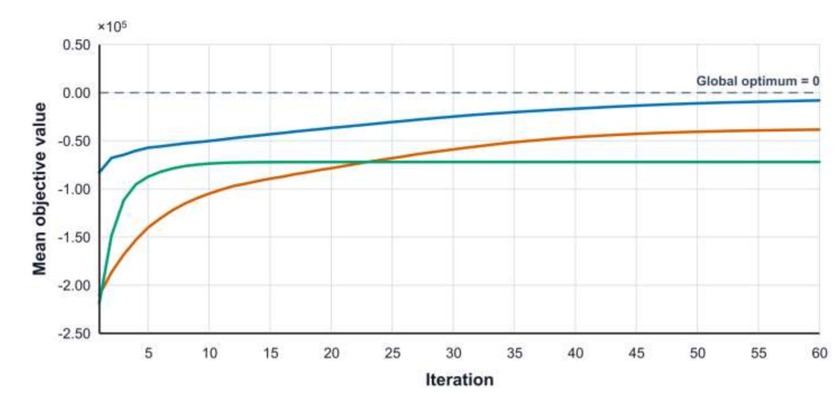

PERFORMANCE

CE = Cross-Entropy Method baseline · PSA = Population-Based Simulated Annealing baseline
Mean = average objective across 50 seeds at that method's best iteration
95% CI = uncertainty interval for that mean across the 50 seeds

Experiment	Dim.	GBO-CD	CE baseline	PSA baseline
Ackley	30	Mean: -1.28 95% CI: [-1.42, -1.13]	Mean: -5.21 95% CI: [-5.79, -4.64]	Mean: -14.62 95% CI: [-14.89, -14.35]
Alpine	30	Mean: -0.88 95% CI: [-0.97, -0.78]	Mean: -0.82 95% CI: [-0.99, -0.65]	Mean: -22.23 95% CI: [-23.25, -21.21]
Lévy	30	Mean: -2.69 95% CI: [-2.80, -2.57]	Mean: -2.62 95% CI: [-3.11, -2.14]	Mean: -22.58 95% CI: [-24.91, -20.25]
Pintér	30	Mean: -694 95% CI: [-779, -609]	Mean: -1,702 95% CI: [-1,842, -1,561]	Mean: -5,381 95% CI: [-5,550, -5,212]
Ackley	100	Mean: -3.10 95% CI: [-3.27, -2.94]	Mean: -11.93 95% CI: [-12.26, -11.61]	Mean: -15.87 95% CI: [-16.10, -15.64]
Alpine	100	Mean: -14.07 95% CI: [-15.93, -12.22]	Mean: -22.58 95% CI: [-23.66, -21.50]	Mean: -94.99 95% CI: [-97.59, -92.39]
Lévy	100	Mean: -13.72 95% CI: [-14.84, -12.60]	Mean: -44.60 95% CI: [-47.88, -41.33]	Mean: -133.58 95% CI: [-140.24, -126.92]
Pintér	100	Mean: -8,067 95% CI: [-8,389, -7,744]	Mean: -38,354 95% CI: [-39,392, -37,317]	Mean: -71,970 95% CI: [-73,482, -70,459]

KEY FINDINGS

01 **GBO-CD achieved the best results on 6 of the 8 benchmark functions, including all four 100D benchmarks.**

02 **GBO-CD demonstrated a clearer performance advantage over CE and PSA on the 100D benchmarks.**

03 **GBO-CD generally reached better solutions earlier and continued improving across iterations.**

CONCLUSION

GBO-CD demonstrates strong potential for high-dimensional black-box optimization under limited real-evaluation budgets. Its pseudo-labels help the conditional diffusion model locate higher-value regions without additional objective evaluations. Future directions include reducing model training time and testing GBO-CD on a wider range of high-dimensional benchmark functions.