Bayesian Optimization for Categorical and Mixed Variables Using a Multinomial Logit Surrogate

Muhammad Amir Saeed **Antonio Candelieri**

Department of Economics, Quantitative Methods and Business Strategies
University of Milano-Bicocca, Milan, Italy

Abstract

Bayesian optimization is effective for expensive black-box problems, but standard Gaussian-process surrogates are less suitable for categorical and mixed-variable search spaces. We propose MNL-BO, a preference-based Bayesian optimization method that uses a multinomial logit surrogate trained from pairwise comparisons. The model provides interpretable utility estimates for categorical alternatives and supports continuous, discrete, and categorical variables in a unified framework. Experiments on categorical benchmarks, the Traveling Salesman Problem, and mixed-variable pressure vessel design show competitive performance against random search, local search, metaheuristics, and SMAC-inspired tree-based Bayesian optimization baselines.

Problem Statement and Contributions

Many optimization problems involve categorical and mixed-variable decision spaces where traditional Gaussian-process surrogates require specialized kernels or encodings. This work introduces MNL-BO, a preference-driven optimization framework based on a multinomial logit surrogate that directly models structured decision alternatives through pairwise comparisons.

- Preference-based surrogate learning using MNL models
- Utility and uncertainty estimation for acquisition-driven search
- Unified handling of categorical, discrete, and continuous variables
- Empirical validation on benchmark, combinatorial, and engineering problems

Proposed Method: MNL-BO

MNL-BO learns a utility-based ranking of candidate configurations from preference comparisons. The MNL choice probability is

$$\pi_j(x) = \frac{\exp(x^\top \boldsymbol{\beta}_j)}{\sum_{k=1}^{m} \exp(x^\top \boldsymbol{\beta}_k)}.$$

Instead of modeling raw objective values, the surrogate represents candidate quality through a latent utility function:

$$P(x) \propto \exp(u(x)).$$

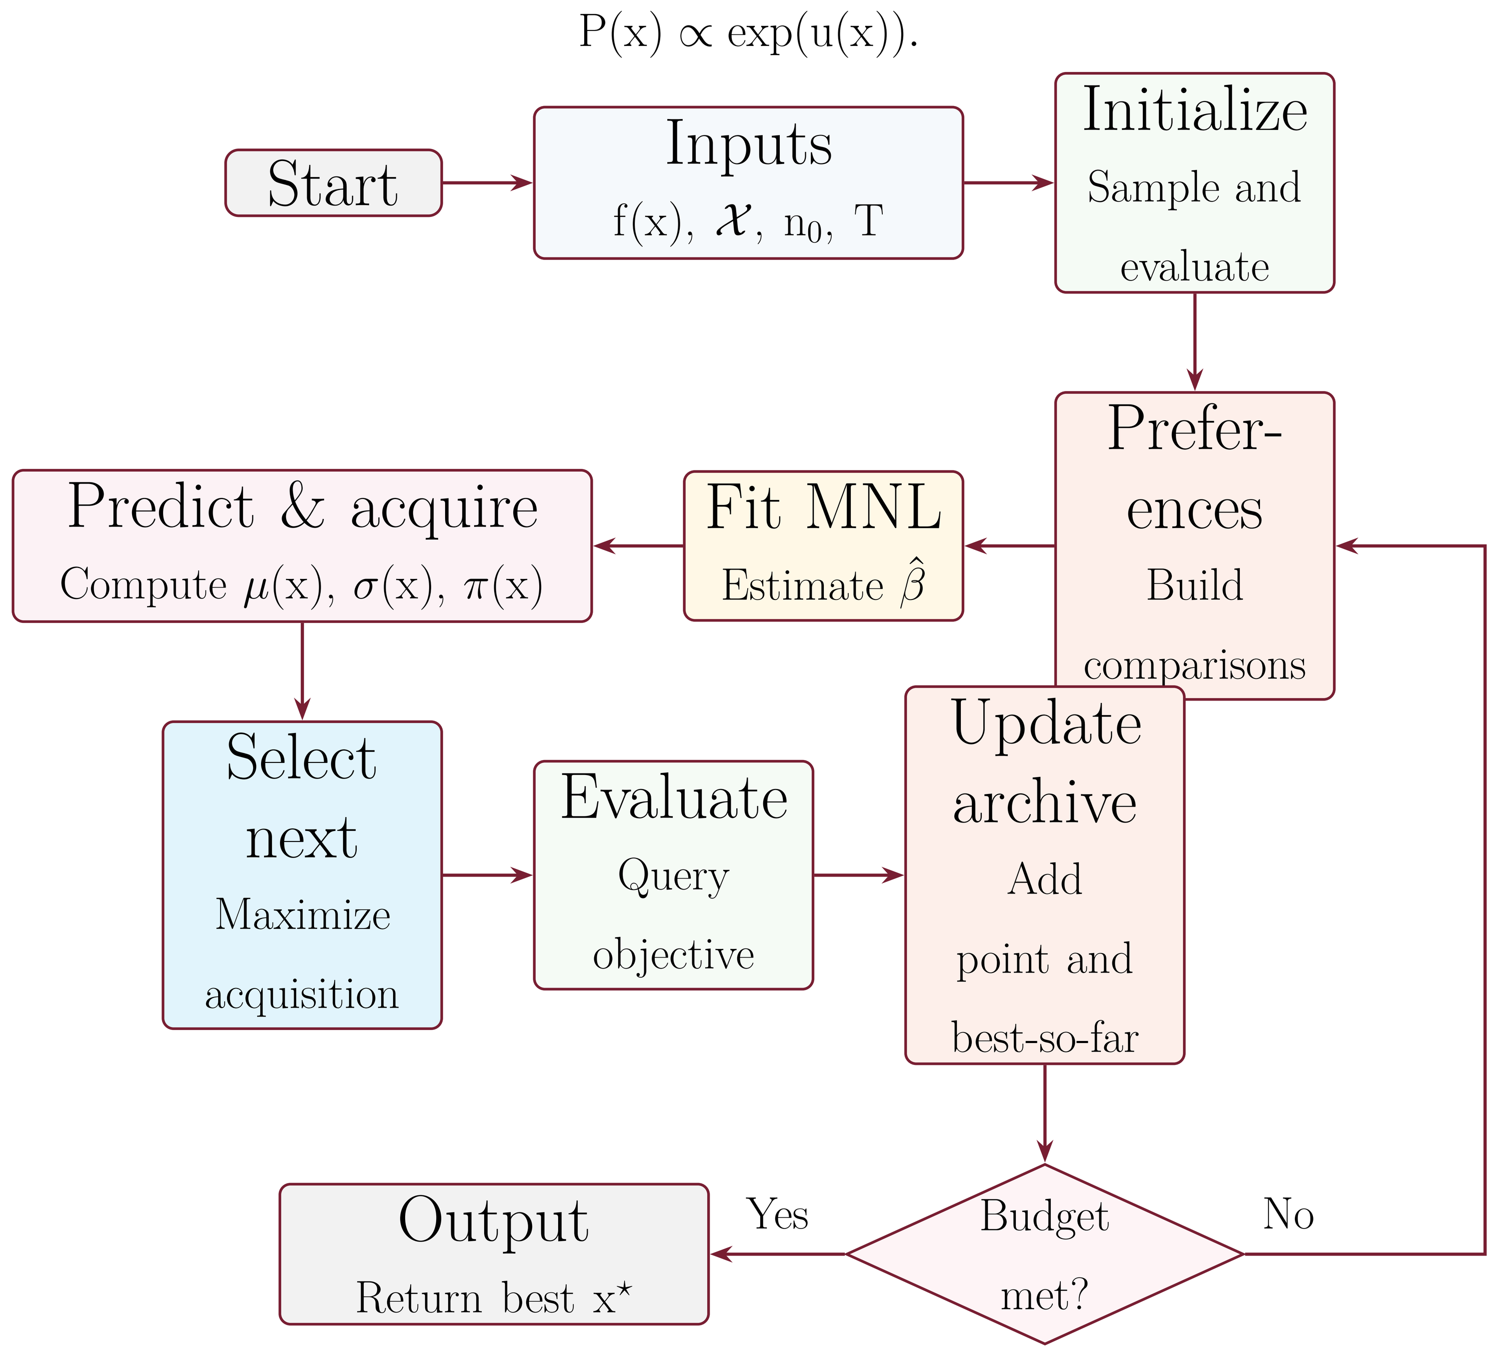

Figure 1. Flowchart of the proposed MNL-BO procedure.

Preference-Based Surrogate Modeling

The surrogate learns preferences between candidate configurations instead of directly modeling raw objective values.

For two candidate solutions x_a and x_b:

$$\Pr(x_a \succ x_b) = \frac{\exp(u(x_a))}{\exp(u(x_a)) + \exp(u(x_b))}$$

Advantages

- Natural handling of categorical variables
- Interpretable utility estimates
- Works for combinatorial and mixed-variable optimization

Experimental Setup

The proposed framework is evaluated on three benchmark problems.

Problem	Variable Type
Categorical Benchmark	Pure categorical
Traveling Salesman Problem	Permutation
Pressure Vessel Design	Mixed-variable

Table 1. Optimization benchmarks used in the study.

Traveling Salesman Problem Case Study

The TSP is used to evaluate MNL-BO on a combinatorial optimization problem where each solution is a permutation of $n = 15$ U.S. cities. The objective is to find the shortest closed tour visiting every city exactly once.

$$L(x) = \sum_{i=2}^{n} C_{x_{i-1},x_i} + C_{x_n,x_1}$$

Candidate tours are generated using 2-opt, swap, insertion, and random permutation moves. The MNL surrogate then selects the next tour using the acquisition function.

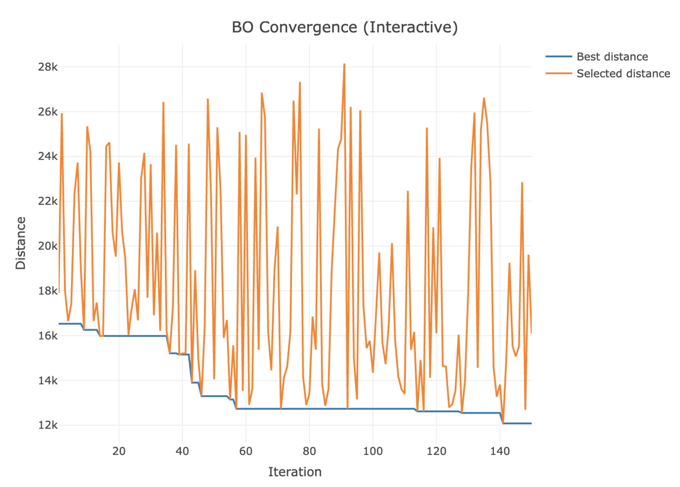

Figure 2. Representative TSP convergence: selected tour distance and best-so-far distance.

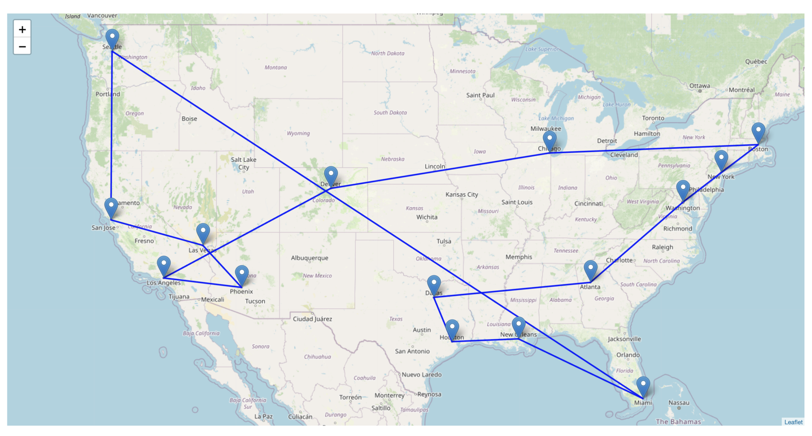

Figure 3. Best route found by MNL-BO for the 15-city TSP instance.

Method	Mean	SD	Best
MNL-BO (EI)	12028.15	1114.61	9790
MNL-BO (UCB)	11922.00	987.52	10073
Random	16121.30	1015.47	13445
2-opt	11533.20	1203.30	9181
RF-SMAC (EI)	10816.25	1061.57	9232
RF-SMAC (UCB)	11271.05	499.68	10230

Table 2. TSP multi-run comparison over $R = 20$ runs. Lower is better.

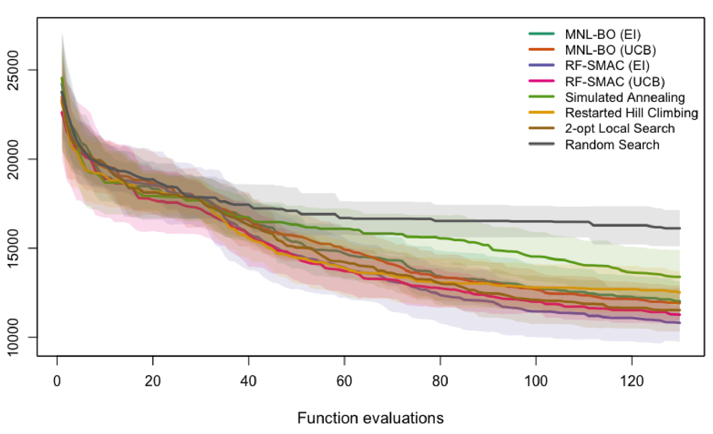

Figure 4. TSP multi-run convergence comparison.

Pressure Vessel Design with Material Selection

The pressure vessel benchmark evaluates MNL-BO on a constrained engineering design problem involving continuous, discrete, and categorical variables.

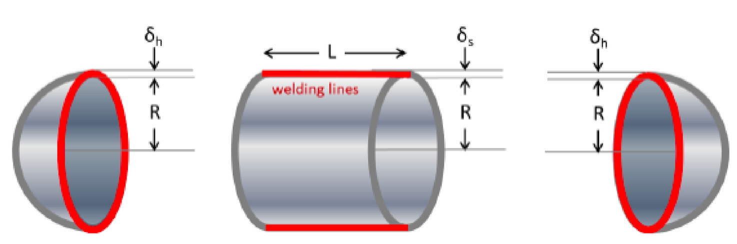

Figure 5. Pressure vessel geometry and dimensions.

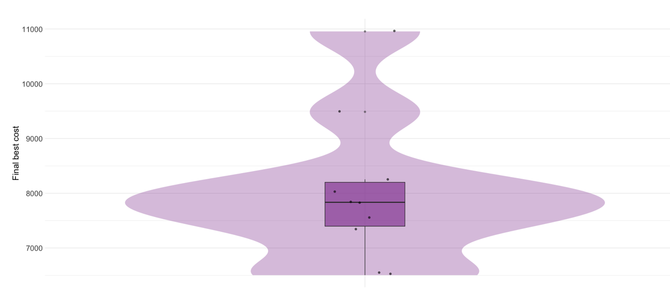

Figure 6. Distribution of final best costs across 10 independent runs.

Conclusion

MNL-BO provides a preference-based Bayesian optimization framework for categorical and mixed-variable search spaces using a multinomial logit surrogate. The method demonstrates competitive performance on combinatorial and engineering optimization benchmarks while maintaining interpretability and flexibility. The proposed framework naturally handles categorical, discrete, and continuous variables within a unified optimization setting without requiring specialized kernels or handcrafted encodings. Experimental results on benchmark, combinatorial, and constrained engineering problems show stable convergence behavior and effective exploration of complex search spaces.

References

[1] F. Archetti and A. Candelieri, Bayesian Optimization and Data Science, Springer, Vol. 849, 2019.

[2] R. Garnett, Bayesian Optimization, Cambridge University Press, 2023.

[3] Y. Croissant, "Estimation of Multinomial Logit Models in R: The mlogit Package," Journal of Statistical Software, Vol. 50, pp. 1–24, 2013.