algorithms

MDPI

Article

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

Muhammad Amir Saeed * and Antonio Candelieri *

Department of Economics, Quantitative Methods and Business Strategies, University of Milano-Bicocca, 20126 Milan, Italy

* Correspondence: m.saeed@campus.unimib.it (M.A.S.); antonio.candelieri@unimib.it (A.C.)

Abstract

Bayesian optimization (BO) is a widely used framework for optimizing expensive black-box functions. Most BO methods rely on Gaussian process (GP) surrogates, which perform well in continuous domains but encounter difficulties when decision variables include categorical or mixed discrete–continuous components. In particular, GP-based approaches typically require ad hoc numerical encodings of categorical variables that may fail to capture the structure of discrete decision spaces. In this work, we propose MNL-BO (Multinomial Logit Bayesian Optimization), a preference-based Bayesian optimization framework that replaces the GP surrogate with a multinomial logit (MNL) model trained from pairwise preference comparisons. The resulting surrogate provides a natural and interpretable representation of categorical alternatives while allowing continuous, discrete, and categorical variables to be handled within a unified optimization framework. The predictive utility estimates and uncertainty indicators generated by the MNL model are employed to formulate acquisition functions that reconcile exploration with exploitation. The proposed methodology is evaluated on three progressively complex optimization challenges: a purely categorical benchmark, a combinatorial Traveling Salesman problem, and a constrained mixed-variable engineering design problem concerning material selection in pressure vessel optimization. Multi-run tests provide consistent advantages over random search and exhibit stable convergence behavior across diverse random initializations. In addition to heuristic baselines such as local search and classical metaheuristics, we also compare against tree-based Bayesian optimization baselines inspired by the Sequential Model-based Algorithm Configuration (SMAC) framework. The results indicate that the proposed MNL-BO method achieves competitive performance under comparable evaluation budgets while providing an interpretable probabilistic surrogate for categorical decision spaces. These findings suggest that preference-based surrogate modeling provides a practical and flexible alternative for Bayesian optimization in categorical and mixed-variable optimization problems.

Keywords: Bayesian optimization; categorical variables; multinomial logit model; black-box optimization; discrete optimization

Check for updates

Academic Editor: Frank Werner

Received: 9 March 2026

Revised: 24 April 2026

Accepted: 29 April 2026

Published: 4 May 2026

1. Introduction

Bayesian optimization (BO) is a widely used strategy for the global optimization of expensive black-box functions [1,2]. The method has garnered considerable attention in machine learning and engineering domains where objective evaluations are computationally expensive, including hyperparameter optimization (HPO), automated machine learning (AutoML), and simulation-based design optimization [3–5].

In these settings, evaluating the objective function may require expensive procedures such as training machine learning models using cross-validation or running computational simulations. Bayesian optimization solves this problem by making a probabilistic surrogate model of the objective function and then repeatedly picking evaluation points that balance exploration and exploitation using an acquisition function [1,2].

Gaussian Process (GP) regression has traditionally been the predominant surrogate model employed in Bayesian Optimization (BO) due to its ability to quantify predictive uncertainty and provide adaptable non-parametric modeling capabilities [1]. Then, acquisition functions such as Expected Improvement (EI) and Upper Confidence Bound (UCB) are optimized to select the next evaluation point [6,7]. Contemporary Bayesian optimization frameworks, including Predictive Entropy Search, use information-theoretic acquisition functions that aim to minimize uncertainty about the optimal location.

Fortunately, many practical optimization problems involve mixed-variable search spaces, in which decision variables may be continuous, integer, or categorical. Examples occur in algorithm configuration, material design, engineering optimization, and AutoML pipelines [8,9]. Consequently, managing such heterogeneous domains has emerged as a significant research focus in Bayesian optimization.

Recent studies have shown that Gaussian Process-based Bayesian Optimization can be modified to handle mixed and categorical variables through the use of specialized kernels, latent-variable encodings, and structured search strategies [10–12]. For instance, Garrido-Merchán and Hernández-Lobato [10] suggested kernel transformations that impose integer and categorical constraints in Gaussian processes, whereas Roustant et al. [12] presented group kernels for the modeling of categorical inputs. Some other methods use latent-variable embeddings [11] or additive structures [13].

Beyond kernel engineering, several frameworks have been developed to support mixed-variable optimization. Examples include surrogate-based engineering optimization methods [8], comparative studies of mixed-variable BO algorithms [9], and modern toolkits such as SMT 2.0 for surrogate modeling with hierarchical and mixed-variable Gaussian processes [14]. General optimization frameworks such as BoTorch [15] and NOMAD version 4 [16] also provide practical implementations for mixed-variable and derivative-free optimization.

Even with these improvements, the search for other surrogate models that can naturally represent discrete or categorical choices is still a lively area of research. Surrogate models based on discrete choice theory offer an attractive framework for tackling optimization problems where choices represent structured design configurations.

This research examines the utilization of a multinomial logit (MNL) model [17,18] as a probabilistic surrogate within the Bayesian Optimization (BO) framework. Unlike Gaussian-process-based approaches, which rely on specialized kernels, latent embeddings, or continuous relaxations to handle categorical variables, the proposed framework adopts a parametric discrete-choice model that directly operates on categorical and mixed-variable inputs without requiring ad hoc encoding schemes.

We introduce **MNL-BO (Multinomial Logit Bayesian Optimization)**, a preference-based optimization framework that constructs the surrogate from pairwise comparisons and probabilistic utility estimates. In contrast to existing preference-based BO methods based on Gaussian processes or Bradley–Terry-type models, the proposed approach leverages a multinomial logit structure to obtain interpretable utility coefficients and a tractable uncertainty proxy derived from asymptotic covariance. This enables efficient modeling of heterogeneous variable types while maintaining interpretability and computational simplicity.

https://doi.org/10.3390/a19050361

We note that the performance of the proposed method depends on problem structure, and while it is effective in several settings, it is not uniformly superior across all problem classes; this is discussed in detail in the experimental section.

1.1. Related Work

There have been several proposed solutions for how to make Bayesian optimization work in domains with discrete and mixed variables.

A prevalent approach is to modify Gaussian process models to accommodate discrete variables using specialized kernels or encodings. Garrido-Merchán and Hernández-Lobato [10] presented kernel transformations that facilitate the direct integration of integer and categorical variables into Gaussian process models. Zhang et al. [11] introduced a latent variable representation for qualitative inputs, and Roustant et al. [12] created group kernels to identify correlations among categorical levels.

Another area of research is exploring hybrid strategies that combine BO with other optimization methods. For instance, CoCaBO [19] combines multi-armed bandit algorithms with Gaussian process optimization to solve problems with both categorical and continuous inputs. Sequential model-based optimization frameworks, such as SMAC [20], use random forest models as surrogates for predictors. They are often used for algorithm configuration and hyperparameter optimization.

Density-estimation methods are also options for modeling that don't use GP. The Tree-structured Parzen Estimator (TPE) algorithm [5,21] employs kernel density estimation to delineate promising and non-promising areas within the search space and is extensively utilized in hyperparameter optimization libraries.

Another new method is GRYFFIN, a Bayesian optimization framework designed specifically for categorical variables in scientific discovery applications. GRYFFIN uses kernel density estimation on categorical spaces and can use expert descriptors to help the optimization process.

For mixed-variable problems, derivative-free optimization methods have been created in addition to Bayesian optimization methods. Pattern search techniques [22] and the MADS-based NOMAD algorithm [16] are two examples. These techniques provide robust theoretical convergence guarantees and have been successfully applied to engineering design challenges.

Numerous Bayesian optimization (BO) methodologies have been proposed for mixed-variable domains; however, most rely on Gaussian process surrogates with specialized kernels or encodings. This study, on the other hand, looks at a discrete-choice modeling point of view by using a multinomial logit surrogate within the BO framework.

There is a lot of writing about Bayesian optimization. Still, most of the time, surrogate models are designed for continuous search spaces and require specialized kernels or ad hoc encodings to handle categorical variables. This drives the creation of surrogate models that inherently depict categorical decision frameworks while maintaining the exploration-exploitation dynamics of Bayesian optimization.

In contrast to GP-based approaches for categorical optimization, such as CoCaBO and kernel-based methods, the proposed MNL surrogate avoids the need for specialized kernels, latent embeddings, or continuous relaxations of discrete variables. Instead, it directly models preferences over discrete configurations using a parametric utility-based formulation.

Compared to density-estimation-based methods such as TPE and GRYFFIN, which rely on sampling or kernel density estimation to approximate the search distribution, the proposed approach provides an explicit parametric model of choice probabilities with interpretable coefficients. Furthermore, unlike tree-based surrogates such as SMAC, which model the objective value directly, the MNL surrogate operates on preference comparisons, offering a different modeling perspective particularly suited to discrete and structured decision spaces.

Overall, the proposed framework should be viewed as a complementary surrogate-based optimization approach, particularly advantageous in settings where categorical structure and interpretability are important. While direct empirical comparisons with CoCaBO, GRYFFIN, and GP-kernel-based methods are not included in the current study, as such comparisons would require careful experimental design to ensure fair evaluation across methods with different candidate-generation and hyperparameter requirements, we provide a structured discussion of the expected relative strengths and limitations of MNL-BO with respect to these approaches. A comprehensive empirical benchmarking study is identified as an important direction for future work. In this broader context, we also note connections to surrogate-based modeling beyond standard regression settings, such as the surrogate approach in ordinal regression [23] and recent extensions to mixed-type data analysis [24], which further highlight the relevance of structured surrogate modeling in complex data domains.

1.2. Contributions

This paper presents MNL-BO (Multinomial Logit Bayesian Optimization), a preference-driven Bayesian optimization framework tailored for optimization challenges that encompass categorical, discrete, and mixed decision variables. The principal contributions of this study are stated as follows:

- We introduce MNL-BO, a Bayesian optimization framework that employs a multinomial logit (MNL) model as a surrogate representation of the objective function. Unlike Gaussian-process-based approaches that require specialized kernels or latent embeddings, the proposed surrogate directly models categorical alternatives through a parametric discrete-choice formulation derived from random utility theory.
- We develop a method for deriving predictive utility estimates and uncertainty measures from the MNL surrogate, enabling the construction of acquisition strategies that balance exploration and exploitation within the BO loop.
- We formulate a pairwise-comparison-based learning procedure for training the surrogate model. This preference-based formulation allows continuous, categorical, and discrete variables to be handled within a unified optimization framework without requiring specialized kernels or encoding strategies.
- We empirically evaluate the proposed method on three representative optimization problems: a categorical benchmark problem, a Traveling Salesman Problem instance, and a constrained engineering design problem involving pressure vessel material selection.
- We empirically compare the proposed approach against multiple baselines, including random search, local-search and metaheuristic methods, and tree-based Bayesian optimization surrogates such as RF-SMAC. The results demonstrate that MNL-BO provides competitive performance in several settings, while also highlighting scenarios where performance depends on problem structure.

2. Methodology

Bayesian Optimization (BO) is a sequential strategy for optimizing expensive black-box functions using a probabilistic surrogate model and an acquisition rule that balances exploration and exploitation [1,2,25]. At each iteration, the surrogate is updated using the set of previously evaluated configurations, and the next query point is selected by maximizing an acquisition function.

Gaussian Process (GP) regression is the most widely used surrogate model in classical BO because it provides both a predictive mean and a predictive uncertainty estimate in continuous domains [1]. However, GP-based BO becomes less natural in mixed-variable and categorical search spaces, where notions such as smoothness, local continuity, and distance-based correlation are difficult to define in a principled way. As a result, GP-based methods often require specialized kernels, latent embeddings, or ad hoc encodings when categorical variables are present [10–12].

To address these limitations, we propose MNL-BO (Multinomial Logit Bayesian Optimization), a preference-based Bayesian optimization framework that replaces the GP surrogate with a multinomial logit (MNL) model. The proposed framework is designed for single-objective optimization problems over categorical, discrete, or mixed discrete–continuous search spaces, where preference-based utility modeling provides a natural surrogate representation. We retain the term Bayesian Optimization to refer to the sequential, model-based optimization loop following established usage in the literature [26]; the surrogate model itself is fitted via maximum-likelihood estimation and is therefore frequentist in nature.

2.1. MNL Surrogate and Preference-Based Representation

The multinomial logit (MNL) model is a classical probabilistic model for discrete choice and categorical decision behavior [17,27,28]. It is rooted in random utility theory, where each alternative is associated with a latent utility and the observed choice corresponds to the utility-maximizing alternative up to a stochastic perturbation [27,28]. This perspective is especially attractive for optimization over discrete design spaces, since candidate configurations can be interpreted as competing alternatives with latent quality scores.

Let $\mathbf{x}$ denote the feature representation of a candidate configuration. In the standard MNL formulation, the probability of selecting alternative $j \in \{1, \ldots, m\}$ is

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

where $\boldsymbol{\beta}_j$ is the parameter vector is associated with category j. For notational convenience, we define $\pi_{ij} = \pi_j(\mathbf{x}_i)$, which represents the probability that alternative j is selected in observation i, given the feature vector $\mathbf{x}_i$. This makes explicit the dependence of the choice probabilities on individual observations. The parameters are estimated by maximizing the log-likelihood,

$$\log L(\boldsymbol{\beta}) = \sum_{i=1}^{N} \sum_{j=1}^{m} y_{ij} \log \pi_{ij}, \tag{2}$$

where $y_{ij} = 1$ if alternative j is selected in observation i, and $y_{ij} = 0$ otherwise.

In the proposed framework, the MNL model is not used to directly regress the raw objective values. Instead, it is used to represent a latent utility function $u(\mathbf{x})$ inferred from observed comparisons between candidate configurations. This leads to a preference-based surrogate, where the relative quality of configurations is encoded through a probability model of the form

$$P(\mathbf{x}) \propto \exp(u(\mathbf{x})). \tag{3}$$

Accordingly, MNL-BO treats optimization as a sequential process of learning a utility-based ranking over the feasible set and then querying new configurations using acquisition functions defined on that learned utility landscape.

2.2. Relation to Pairwise Comparison Models

A useful interpretation of the proposed surrogate arises in the pairwise-comparison setting. The multinomial logit model simplifies to the Bradley-Terry form for paired comparisons when the choice set consists of only two alternatives [29–31]. If we compare two configurations, $\mathbf{x}_a$ and $\mathbf{x}_b$, then

$$\Pr(\mathbf{x}_a \succ \mathbf{x}_b) = \frac{\exp(u(\mathbf{x}_a))}{\exp(u(\mathbf{x}_a)) + \exp(u(\mathbf{x}_b))}. \tag{4}$$

The proposed framework is based on an important theoretical idea: the surrogate used in MNL-BO can be thought of as a generalization of standard pairwise ranking models. It also makes sense to use pairwise preference data inside the BO loop because the resulting likelihood fits well with well-known models for ranking and comparative judgments.

This perspective additionally connects the proposed method to preference learning more broadly. In particular, Gaussian-process preference learning also models latent utilities from comparative observations but does so through a nonparametric GP prior rather than a parametric choice model [32,33]. MNL-BO retains the preference-learning viewpoint while replacing the GP surrogate with a discrete-choice surrogate that is naturally aligned with categorical decision spaces.

2.3. Surrogate Prediction and Uncertainty Quantification

A BO surrogate must provide two key quantities: a point estimate of performance and a measure of uncertainty. In MNL-BO, the surrogate prediction is represented by the estimated latent utility

$$\mu(\mathbf{x}) = \phi(\mathbf{x})^\top \hat{\beta}, \tag{5}$$

where $\phi(\mathbf{x})$ is the feature representation of configuration $\mathbf{x}$ the parameter vector $\hat{\beta}$ is obtained by maximizing the log-likelihood in Equation (2) with respect to $\boldsymbol{\beta}$. Since no closed-form solution exists, the optimization is carried out using standard numerical procedures for multinomial logistic regression.

Unlike GP regression, the standard MNL model does not directly produce a posterior predictive variance. To obtain an uncertainty proxy suitable for BO, we approximate parameter uncertainty using the asymptotic covariance matrix of the maximum-likelihood estimator,

$$\widehat{\Sigma}_\beta = \mathrm{Cov}(\hat{\beta}), \tag{6}$$

where $\mathrm{Cov}(\hat{\beta})$ denotes the estimated covariance matrix of the maximum-likelihood estimator, typically obtained as the inverse of the observed Fisher information matrix, and propagate it to the latent utility through a first-order delta-method approximation:

$$\sigma(\mathbf{x}) = \sqrt{\phi(\mathbf{x})^\top \widehat{\Sigma}_\beta \phi(\mathbf{x})}. \tag{7}$$

The values of $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ are similar to the predictive mean and predictive uncertainty in GP-based BO, but they still fit with the parametric structure of the MNL surrogate.

Robustness in Early Iterations

The delta-method approximation employed for uncertainty quantification may exhibit reduced reliability in small-sample contexts, especially in the initial iterations of the optimization process when preference observations are scarce. In these instances, the covariance estimate may be unreliable or poorly conditioned due to the inadequate identifiability of the MNL model.

The implementation employs basic fallback methods to guarantee numerical stability. In the event that the MNL fit is unsuccessful due to separation, non-identifiability, or near-singularity of the Hessian, the method defaults to an exploratory phase, characterized by uniform random candidate selection for that iteration. Furthermore, the prediction variance is subjected to regularization as

$$\sigma(\mathbf{x}) \leftarrow \sqrt{\max\{\sigma^2(\mathbf{x}), \epsilon\}}, \tag{8}$$

for a small constant $\epsilon > 0$, which prevents degenerate acquisition scores and premature over-exploitation.

In fact, this instability is primarily limited to the initial phase of optimization and swiftly decreases as additional preference data is gathered. Empirically, we discover that these occurrences are rare and do not substantially impact the overall sample efficiency of the approach.

2.4. Theoretical Perspective on the MNL Surrogate

The multinomial logit model employed as a surrogate in the proposed MNL-BO framework is based on the random utility model prevalent in discrete choice theory [17,28]. According to standard assumptions, each candidate configuration has a hidden utility value, and the observed pairwise preferences are based on comparisons between these hidden utilities.

When the MNL model is correctly specified and there are sufficient observations, maximum-likelihood estimation yields consistent estimates of the underlying utility parameters. As a result, the surrogate can be seen as a way to estimate the hidden objective structure that controls how candidates prefer one configuration over another.

2.4.1. Theoretical Implications of Heuristic Acquisition Functions

A central theoretical distinction between GP-based BO and the proposed MNL-BO framework concerns the availability of formal performance guarantees. In GP-based BO, regret bounds are derived from the properties of the Gaussian process posterior. Specifically, Srinivas et al. [34] established that GP-UCB achieves sublinear cumulative regret under mild assumptions on the kernel, with the regret bound depending on the maximum information gain of the chosen kernel. This result relies critically on the availability of a proper posterior predictive distribution, from which both the predictive mean and variance are derived as exact Bayesian quantities. Similarly, theoretical analyses of expected improvement under GP surrogates [35] exploit the fact that the posterior is a well-defined probability measure over functions, enabling convergence guarantees under appropriate regularity conditions.

In contrast, the MNL surrogate does not produce a formal posterior predictive distribution. The uncertainty estimate $\sigma(\mathbf{x})$ in Equation (7) is derived from the asymptotic covariance of the maximum-likelihood estimator via the delta method, which provides a first-order approximation rather than an exact Bayesian quantity. As a consequence, the mathematical machinery underlying GP-based regret analysis does not directly transfer to the MNL-BO setting, and no analogous regret bounds are available for the proposed acquisition functions.

Nevertheless, certain weaker theoretical properties can still be established. Under standard regularity conditions for multinomial logistic regression, the MLE $\hat{\beta}$ is consistent and asymptotically normal as the number of preference observations grows [28]. This implies that the surrogate utility estimates converge to the true latent utility parameters as more data is collected, providing an informal convergence guarantee for the surrogate itself. Furthermore, in finite discrete domains, the feasible set is exhaustible in principle: if the acquisition function ensures that every candidate is selected with nonzero probability, the algorithm will eventually evaluate all configurations, guaranteeing identification of the global optimum in finite iterations. While this property is weaker than sublinear regret, it provides a theoretical basis for the soundness of the approach in bounded discrete spaces.

The primary justification for the proposed framework therefore rests on this consistency property together with the empirical evidence presented in Section 3.

2.4.2. Modeling Assumptions and Limitations

The multinomial logit surrogate presupposes a linear utility framework and adheres to the independence of irrelevant alternatives (IIA) assumption, potentially constraining its capacity to encapsulate intricate nonlinear relationships among choice variables. In combinatorial challenges like the traveling salesman problem or engineering design problems, the actual objective function may have significant relationships that are inadequately captured by a linear model.

In the proposed framework, this choice reflects a trade-off between model expressiveness and computational tractability. The parametric structure of the MNL model enables efficient estimation, interpretability, and stable integration within the BO loop. Moreover, structured feature representations can capture part of the combinatorial structure, allowing the surrogate to provide informative rankings of candidate configurations.

From an optimization perspective, the surrogate does not need to perfectly model the objective function globally; rather, it must provide sufficiently accurate relative comparisons to guide the search. Empirical results suggest that this level of approximation is sufficient to support effective optimization. Extending the framework to more expressive models, such as mixed logit or nonlinear preference models, is a promising direction for future work.

2.5. MNL-BO Framework

The proposed MNL-BO framework uses the standard sequential logic of Bayesian optimization, but it changes each step to fit a preference-based discrete-choice surrogate.

The algorithm starts by looking at a first set of possible configurations. From the resulting archive of observed objective values, pairwise or choice-based preference data are constructed. Next, a multinomial logit surrogate is fitted to these preference observations. This gives us an estimated latent utility model for the candidate configurations.

Then, an acquisition function that is defined on the latent utility space is tested on the set of configurations that are possible but have not yet been tested. The acquisition rule picks the best candidate, tests it with the real objective function, and adds it to the archive. After that, the surrogate model and acquisition function are updated, and the process starts over until the evaluation budget runs out.

This method keeps the main idea of BO, which is making decisions one at a time when we don't know what will happen, but it replaces the GP surrogate with a preference-based utility model that works better for categorical and mixed-variable domains.

2.6. Adaptation of Acquisition Functions

Replacing the GP surrogate with an MNL surrogate requires revisiting the interpretation of classical BO acquisition functions. In GP-based BO, acquisition criteria such as Expected Improvement (EI) and Upper Confidence Bound (UCB) are derived from a predictive posterior distribution over objective values. In MNL-BO, the surrogate instead provides latent utility estimates and an approximate uncertainty measure derived from parameter covariance. Consequently, the acquisition functions proposed here should be interpreted as BO-inspired heuristic criteria rather than exact posterior quantities.

2.6.1. Utility Scale and Normalization

The latent utility $\mu(\mathbf{x})$ is identifiable only up to an additive constant, and its scale may vary across iterations. For numerical stability, utilities are re-centered internally during probability computation, for example by replacing $\mu(\mathbf{x})$ with $\mu(\mathbf{x}) - \max_{\mathbf{x}'} \mu(\mathbf{x}')$ inside softmax calculations. More generally, standardization over the candidate set can be applied without changing the ranking induced by the surrogate.

2.6.2. Probability-Weighted Expected Improvement

For maximization problems, we define the probability-weighted expected improvement heuristic

$$\mathrm{EI}_{\mathrm{MNL}}(\mathbf{x}) = \max\big(\mu(\mathbf{x}) - f(\mathbf{x}^{+}), 0\big)\, \pi(\mathbf{x}), \tag{9}$$

where $f(\mathbf{x}^{+})$ is the best objective value observed so far and $\pi(\mathbf{x})$ is the probability assigned to configuration $\mathbf{x}$ by the MNL surrogate. For minimization problems, the improvement term is defined analogously with reversed sign. This criterion prioritizes alternatives that simultaneously exhibit high utility and high surrogate probability of being preferred.

2.6.3. Confidence-Bound Criteria

To incorporate uncertainty more explicitly, we define MNL-based confidence-bound acquisition rules:

$$\mathrm{UCB}_{\mathrm{MNL}}(\mathbf{x}) = \mu(\mathbf{x}) + \kappa\, \sigma(\mathbf{x})\, \pi(\mathbf{x}), \tag{10}$$

$$\mathrm{LCB}_{\mathrm{MNL}}(\mathbf{x}) = \mu(\mathbf{x}) - \kappa\, \sigma(\mathbf{x})\, \pi(\mathbf{x}), \tag{11}$$

where $\kappa > 0$ controls the exploration-exploitation trade-off. The uncertainty term encourages exploration of under-informed configurations, while the probability weighting discourages excessive exploration of implausible alternatives.

2.6.4. Choice of κ

For reproducibility, all experiments in this work use a fixed exploration parameter $\kappa = 2.0$, held constant across iterations, following the convention of GP-UCB [34]. This value reflects a moderate exploration bias and was found to yield stable performance across the benchmark problems considered. Higher values of κ promote broader exploration and may be preferable in highly multimodal or poorly characterized search spaces, while lower values concentrate acquisition on exploitation of currently promising regions. Adaptive or theoretically calibrated schedules for κ are possible and represent a natural direction for future work.

2.6.5. Interpretation and Justification

The acquisition rules in (9)–(11) are motivated by the same exploration-exploitation principle that underlies classical BO, but they are not exact analogues of GP-derived EI or UCB. Instead, they combine three ingredients that are naturally available from the MNL surrogate: estimated utility, approximate uncertainty, and choice probability.

The probability-weighted EI criterion in (9) combines a deterministic improvement term with a probabilistic weighting that reflects how strongly the surrogate favors a given configuration. Similarly, the confidence-bound rules in (10) and (11) extend optimism- or pessimism-based search to the discrete-choice setting by adjusting utility estimates according to local uncertainty.

These criteria are best understood as principled heuristics tailored to the preference-learning surrogate, rather than posterior-optimal decision rules. This distinction is important for readers to understand: the proposed framework keeps the operational structure of BO while purposefully relaxing the stronger probabilistic assumptions that are possible in GP-based formulations. As a result, the proposed acquisition functions should be interpreted as surrogate-based ranking criteria that balance exploitation (through the estimated utility) and exploration (through the uncertainty proxy). While this framework does not provide the same theoretical guarantees as fully Bayesian approaches, it is well-suited to finite discrete domains, where such approximations can still lead to effective optimization performance in practice.

It is important to note that the acquisition functions defined in (9)–(11) are heuristic adaptations of classical Bayesian optimization criteria. Unlike Gaussian process models, the MNL surrogate does not yield a full posterior predictive distribution. Instead, it provides a parametric estimate of the latent utility together with an uncertainty proxy derived from the asymptotic covariance of the maximum-likelihood estimator. In classical GP-based BO, acquisition functions such as GP-UCB and EI are grounded in exact posterior quantities, which enables formal regret analysis [34,35]. These guarantees rely on the GP posterior being a proper probability measure, a property that is not available in the MNL setting. This distinction is fundamentally epistemological: GP-based BO operates within a Bayesian paradigm, maintaining a prior and updating a posterior over functions, whereas MNL-BO operates within a frequentist paradigm, relying on maximum-likelihood estimation and asymptotic approximations. As a result, the proposed acquisition functions cannot inherit these guarantees and should be interpreted as surrogate-based ranking criteria that balance exploitation through the estimated utility and exploration through the uncertainty proxy. Their justification is grounded in the consistency of MLE-based utility estimation and empirical performance, particularly in finite discrete domains where effective optimization can be achieved without formal posterior guarantees.

2.6.6. Relation to Other Acquisition Strategies

There are now several advanced acquisition strategies for Bayesian optimization. These include information-theoretic criteria like Predictive Entropy Search [36] and mixed-variable acquisition rules like WB2 and WB2S [37]. In many cases, these methods work very well in the real world, especially when the problem is complicated or has a lot of dimensions.

However, many of these acquisition rules rely on posterior quantities that are naturally defined for Gaussian-process surrogates but are not directly available for multinomial logit models. For this reason, the present work focuses on adapting classical acquisition strategies such as EI and confidence-bound criteria, which can be reformulated more transparently using latent utilities and delta-method uncertainty estimates. Extending information-theoretic acquisition design to discrete-choice surrogates is a promising direction for future work.

2.7. Specific Characteristics of MNL-BO

The proposed MNL-BO framework differs from traditional GP-based BO methods in several important ways:

- **Natural treatment of categorical variables.** The surrogate operates directly on discrete alternatives and therefore avoids the need for specialized categorical kernels, one-hot relaxations, or latent embeddings.
- **Relaxed smoothness assumptions.** The method does not require the objective function to vary smoothly with respect to a continuous geometry of the search space, making it suitable for non-smooth or combinatorial design domains.
- **Interpretability.** The estimated utility coefficients provide a transparent parametric description of how features contribute to the surrogate preference structure, which is often harder to obtain from kernel-based models.
- **Finite-domain acquisition optimization.** Acquisition maximization is carried out over a finite feasible set, avoiding many of the numerical difficulties associated with optimizing acquisition functions in continuous spaces.

https://doi.org/10.3390/a19050361

2.8. Optimization Process

To summarize the sequential optimization procedure, Algorithm 1 presents the proposed MNL-BO framework, while Figure 1 illustrates the overall workflow of the method.

Algorithm 1 MNL-BO via Pairwise Preferences over a Categorical Domain

Require: Objective function $f(\mathbf{x})$, feasible domain $\mathcal{X}$, evaluation budget T, initial design size n_0, acquisition rule $\alpha \in \{\text{EI}_{\text{MNL}}, \text{UCB}_{\text{MNL}}\}$
Ensure: Best observed configuration $\mathbf{x}^\star$

Initialization: sample n_0 distinct configurations $\{\mathbf{x}_1, \dots, \mathbf{x}_{n_0}\} \subset \mathcal{X}$ uniformly at random
Evaluate $y_i \leftarrow f(\mathbf{x}_i)$ for $i = 1, \dots, n_0$
Set evaluated index set $\mathcal{E} \leftarrow \{1, \dots, n_0\}$ and best-so-far value $y^+ \leftarrow \max_{i \in \mathcal{E}} y_i$
Initialize preference archive $\mathcal{P} \leftarrow \emptyset$
for $t = n_0 + 1$ to T **do**
  Construct preference data:
  Select two previously evaluated configurations $\mathbf{x}_a, \mathbf{x}_b$ from the archive
  Define winner/loser according to the observed objective values:

$$(\mathbf{x}^{(w)}, \mathbf{x}^{(l)}) = \begin{cases} (\mathbf{x}_a, \mathbf{x}_b), & y_a \geq y_b, \\ (\mathbf{x}_b, \mathbf{x}_a), & y_b > y_a. \end{cases}$$

  Update $\mathcal{P} \leftarrow \mathcal{P} \cup \{(\mathbf{x}^{(w)}, \mathbf{x}^{(l)})\}$
  Fit surrogate:
  Fit a conditional logit / pairwise MNL model on $\mathcal{P}$ so that

$$\Pr(\mathbf{x}^{(w)} \succ \mathbf{x}^{(l)}) = \frac{\exp(\mu(\mathbf{x}^{(w)}))}{\exp(\mu(\mathbf{x}^{(w)})) + \exp(\mu(\mathbf{x}^{(l)}))}$$

  Obtain coefficient estimate $\hat{\beta}$ and covariance estimate $\widehat{\Sigma}_\beta$
  Predict utility and uncertainty:
  for each $\mathbf{x} \in \mathcal{X}$ **do**
    Compute $\mu(\mathbf{x}) = \phi(\mathbf{x})^\top \hat{\beta}$
    Compute $\sigma(\mathbf{x}) = \sqrt{\phi(\mathbf{x})^\top \widehat{\Sigma}_\beta \phi(\mathbf{x})}$
  end for
  Compute probability weights

$$\pi(\mathbf{x}) = \frac{\exp(\mu(\mathbf{x}))}{\sum_{\mathbf{x}' \in \mathcal{X}} \exp(\mu(\mathbf{x}'))} \qquad \forall \mathbf{x} \in \mathcal{X}$$

  Acquisition maximization:
  Evaluate $\alpha(\mathbf{x})$ for all unevaluated $\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_i\}_{i \in \mathcal{E}}$
  Select

$$\mathbf{x}_t \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_i\}_{i \in \mathcal{E}}} \alpha(\mathbf{x})$$

  Evaluate $y_t \leftarrow f(\mathbf{x}_t)$
  Update $\mathcal{E} \leftarrow \mathcal{E} \cup \{t\}$ and $y^+ \leftarrow \max(y^+, y_t)$
end for
return $\mathbf{x}^\star \leftarrow \arg\max_{i \in \mathcal{E}} y_i$

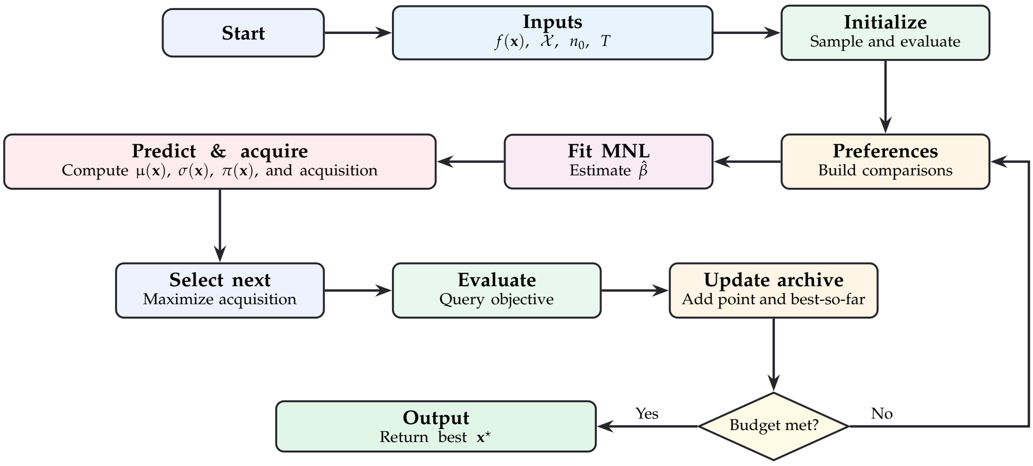

Figure 1. Flowchart of the proposed MNL-BO procedure.

2.8.1. Pairwise MNL Fitting

We briefly clarify how the preference data is used to fit the surrogate model. Given a set of pairwise comparisons $\mathcal{P} = \{(\mathbf{x}^{(w)}, \mathbf{x}^{(l)})\}$, we model the probability that $\mathbf{x}^{(w)}$ is preferred over $\mathbf{x}^{(l)}$ using a pairwise multinomial logit formulation:

$$\Pr(\mathbf{x}^{(w)} \succ \mathbf{x}^{(l)}) = \frac{\exp(\mu(\mathbf{x}^{(w)}))}{\exp(\mu(\mathbf{x}^{(w)})) + \exp(\mu(\mathbf{x}^{(l)}))}.$$

Here, $\mu(\mathbf{x}) = \phi(\mathbf{x})^\top \beta$ denotes the latent utility function. The model parameters β are estimated via maximum likelihood using the preference dataset $\mathcal{P}$, yielding the estimator $\hat{\beta}$ together with its associated covariance estimate $\widehat{\Sigma}_\beta$.

2.8.2. Computational Complexity

Let N denote the number of evaluated configurations and d the dimensionality of the feature representation. Training the multinomial logit surrogate via maximum likelihood requires solving a convex optimization problem whose cost grows approximately as $\mathcal{O}(Nd)$ per iteration. Within the Bayesian optimization loop, the main computational components consist of surrogate fitting and acquisition evaluation over the candidate set. In many practical applications of Bayesian optimization, the cost of these operations remains small compared with the cost of evaluating the underlying objective function. In the experiments reported in this work, the computational overhead of surrogate fitting remained negligible relative to the overall optimization runtime.

2.8.3. Scaling of the Preference Archive

A naïve use of all pairwise comparisons among T evaluated points would result in $\mathcal{O}(T^2)$ comparisons. To avoid this growth, the implementation uses problem-dependent subsampling strategies:

- for the synthetic categorical benchmark, one cumulative choice situation is constructed per BO iteration;
- for the TSP experiments, the conditional logit model is fit using a fixed number of random choice sets per iteration;
- for the pressure vessel experiments, the model is fit using a fixed number of randomly sampled preference pairs per iteration.

This keeps the refitting cost manageable in the larger application-oriented problems while preserving sufficient preference information for stable estimation.

In particular, for combinatorial problems with exponentially large search spaces, the effectiveness of the framework depends critically on the design of candidate generation operators, which may require problem-specific adaptation.

2.9. Theoretical Properties of the MNL Surrogate

The following section provides theoretical insight into the proposed surrogate and clarifies its connection to established preference-learning models.

2.9.1. Consistency Under Random Utility Assumptions

When the data-generating process follows a random utility model with logit choice probabilities, the maximum-likelihood estimator is consistent and asymptotically normal under standard regularity conditions for multinomial logit estimation [27,28]. In the context of MNL-BO, this indicates that when observed comparisons arise from a stable latent utility structure and the preference sample increases adequately, the estimated coefficient vector $\hat{\beta}$ approaches the parameter values that define that utility representation. So, even though the BO loop is sequential, the surrogate itself is based on a strong inferential foundation.

2.9.2. Connection to Bradley-Terry and Pairwise Ranking

When the choice set comprises precisely two alternatives, the multinomial logit model simplifies to the Bradley-Terry model for paired comparisons [29,31]. Consequently, the pairwise variant of MNL-BO can be construed as the acquisition of a parametric ranking model for candidate configurations. This link makes the proposed framework's methodological basis stronger by showing that the surrogate is not just a temporary replacement for GP regression but a principled ranking-based model built into a BO loop.

2.9.3. Regret Intuition in Finite Domains

Classical GP-based BO allows for regret bounds under smoothness assumptions that are typically not present in categorical or mixed discrete spaces. The current framework does not assert similar regret guarantees. The MNL-BO mechanism still follows the basic BO idea of using predicted utility and uncertainty to give evaluations in order. In a limited search space, repeated sampling diminishes uncertainty regarding the candidate set, while acquisition-driven exploration enhances the probability of encountering high-utility configurations. This gives us a qualitative regret intuition that is similar to preference-based BO methods, but it doesn't make the assumption that things are always smooth, which is often not true in categorical domains [32,33].

2.9.4. Scope of the Theoretical Claims

The above arguments justify the surrogate from the viewpoint of discrete choice theory and preference learning, but they do not constitute a formal global convergence proof for the complete MNL-BO algorithm. Establishing finite-time regret bounds or consistency results for the full sequential procedure remains an important direction for future work.

2.10. Implementation Details and Reproducibility

All experiments were implemented in `R` [38] using the `mlogit` package for conditional logit estimation [39]. For the TSP and pressure-vessel studies, we additionally used `data.table` for data handling [40] and `ggplot2` for visualization [41].
Inputs and outputs per BO iteration. At iteration t, the algorithm takes as input the archive of evaluated designs $\{(\mathbf{x}_i, y_i)\}_{i=1}^{t-1}$. It then: (i) constructs a preference dataset from the archive, (ii) fits a conditional logit model to obtain $\hat{\beta}$ and $\widehat{\Sigma}_\beta$, (iii) evaluates the acquisition function over a finite candidate set, and (iv) returns the next query $\mathbf{x}_t$ together with the updated best-so-far value.

Randomness control. All stochastic aspects, such as the initial design, preference subsampling, and candidate generation, are regulated by fixed random seeds to guarantee reproducibility.

Computational Overhead in Practice

Let N denote the number of evaluated configurations and d the dimensionality of the feature representation. Training the multinomial logit surrogate via maximum likelihood typically requires solving a convex optimization problem with computational complexity approximately $\mathcal{O}(Nd)$ per iteration. In the Bayesian optimization loop, the dominant computational cost arises from surrogate fitting and acquisition evaluation over the candidate set. In practice, this cost remains small relative to expensive objective evaluations in typical Bayesian optimization applications.

3. Empirical Study

This section presents the empirical evaluation of the proposed **MNL-BO (Multinomial Logit Bayesian Optimization)** framework. The experiments are designed to assess the effectiveness of the proposed surrogate model and acquisition functions on optimization problems in which categorical or mixed-variable decisions play a central role.

We consider three representative classes of problems:

1. a synthetic purely categorical benchmark designed to stress-test categorical modeling and uncertainty handling;
2. a combinatorial optimization problem, namely the Traveling Salesman Problem (TSP);
3. a mixed-variable engineering design problem involving pressure vessel optimization with material selection.

The first benchmark provides a controlled setting in which the discrete optimization behavior of the method can be isolated and analyzed in detail, whereas the latter two serve as application-oriented case studies.

3.1. Evaluation Protocol and Baselines

All methods are evaluated under a fixed query budget, defined as the number of expensive objective-function evaluations. For stochastic methods, results are reported over multiple independent runs with different random seeds, using the same initial design size and the same evaluation budget. Performance is summarized using: (i) best-so-far objective-value trajectories over iterations, and (ii) final best-so-far values reported as mean $\pm$ standard deviation across runs.

Baselines. To contextualize the performance of MNL-BO, we compare against several baseline strategies commonly used in discrete and mixed black-box optimization: (i) *Random Search*, which performs uniform sampling over the feasible domain; (ii) *Local-search heuristics*, such as the widely used 2-opt neighborhood improvement strategy for the Traveling Salesman Problem [42]; (iii) *Classical metaheuristics*, including simulated annealing [43] and restarted hill climbing, which are commonly applied to permutation-based combinatorial optimization; (iv) *SMAC-style tree-based Bayesian optimization* using a random-forest surrogate for discrete and mixed spaces [20]; and (v) *Derivative-free mixed-variable optimization* methods such as NOMAD (MADS) [16]. Where applicable, we also discuss connections with Gaussian-process-based mixed-variable BO methods available in the literature and implemented in toolboxes such as SMT 2.0 [14].

All baseline methods are evaluated under the same initialization procedure and evaluation budget to ensure a fair comparison.

Reproducibility

All experiments were implemented in the R environment using publicly available packages, primarily `mlogit`, `data.table`, and `ggplot2`. Random seeds were fixed for all multi-run experiments to ensure repeatability. To improve reproducibility, we explicitly report the main implementation settings used in each study, including initialization sizes, evaluation budgets, acquisition parameters, surrogate-training settings, and candidate-generation procedures.

For the purely categorical benchmark, we use an initial design of $n_0 = 10$ configurations and a total evaluation budget of $T = 60$, with acquisition parameter $\kappa = 2.0$ and a small observation noise level.

For the TSP experiments, the representative runs use $n_{\text{init}} = 30$ initial tours and $n_{\text{iter}} = 150$ sequential iterations. At each iteration, the surrogate is trained using 40 sampled choice sets of size $K = 5$. Candidate solutions are generated using a mixture of 90 local moves and 30 globally random tours, and the candidate pool is capped at 120 tours. Local candidate generation relies on standard permutation operators, including 2-opt, swap, and insertion moves.

For the pressure-vessel problem, we use $n_{\text{init}} = 50$ initial designs and $n_{\text{iter}} = 150$ iterations, with 80 pairwise comparisons used per surrogate refit and a candidate pool size of 120. Candidate generation combines local mutations and global random sampling, with local perturbation parameters set to $p_{\text{mat}} = 0.25$, $p_{\text{thick}} = 0.40$, and Gaussian perturbations for continuous variables with standard deviations $\sigma_R = 6$ and $\sigma_L = 15$. Constraint violations are handled using a penalty coefficient of 10^6.

For clarity, the main hyperparameters used across experiments are summarized in Table 1.

Table 1. Key hyperparameters used in the empirical study.

Experiment	**Parameter**	**Value**
Categorical	Initial design (n_0)	10
Categorical	Budget (T)	60
Categorical	Acquisition parameter (κ)	2.0
TSP	Initial tours (n_{init})	30
TSP	Iterations (n_{iter})	150
TSP	Choice sets / iteration	40
TSP	Choice set size (K)	5
TSP	Candidate pool size	120
Pressure Vessel	Initial designs (n_{init})	50
Pressure Vessel	Iterations (n_{iter})	150
Pressure Vessel	Pairwise comparisons	80
Pressure Vessel	Candidate pool size	120
Pressure Vessel	Penalty coefficient	10^6

3.2. Purely Categorical Benchmark

We use a synthetic, purely categorical revise to avoid repetition, for example, a benchmark designed to stress-test categorical surrogate modeling and uncertainty handling in MNL-BO in a controlled and completely separate setting. This benchmark lets us see how the proposed method works in the situation it is meant to work best in, without adding any fake metric structure or smoothness assumptions.

3.2.1. Problem Definition

Let the decision variable be a vector of categorical components

$$\mathbf{c} = (c_1, c_2, \ldots, c_5),$$

where each variable takes values in the finite set

$$c_j \in \{1, 2, 3, 4\}, \quad j = 1, \ldots, 5.$$

The resulting search space contains

$$|\mathcal{X}| = 4^5 = 1024$$

distinct configurations, forming a finite but nontrivial categorical domain.

The objective function is defined as

$$f(\mathbf{c}) = \sum_{j=1}^{5} \theta_j(c_j) \; + \; \sum_{j<k} \theta_{jk}(c_j, c_k) \; + \; \eta, \tag{12}$$

where $\theta_j(c_j)$ represents the main effect associated with variable c_j, $\theta_{jk}(c_j, c_k)$ represents pairwise interaction effects, and η denotes an optional zero-mean stochastic perturbation used to emulate noisy evaluations.

3.2.2. Construction of the Utility Landscape

To enforce a distinct global optimum, the coefficients $\theta_j(\cdot)$ and $\theta_{jk}(\cdot, \cdot)$ are generated randomly and then adjusted. Specifically, we select the reference configuration

$$\mathbf{c}^{\star} = (1, 2, 3, 4, 1)$$

and make it more favorable than all other configurations so that it becomes the best way to achieve the noise-free objective. To make the optimization task more difficult, several configurations that only one change component $\mathbf{c}^{\star}$ are also given high objective values. This creates a group of near-optimal competitors.

Rationale for Enforcing a Distinct Optimum

The benchmark is designed so that the global maximizer is known ahead of time, which makes it possible to use oracle-based convergence assessment with a limited evaluation budget. Adding near-optimal competitors is meant to make the problem more like real-world situations where many options perform similarly, which is a common problem in choosing materials and designs. This controlled construction lets us look at how the proposed preference-based surrogate optimizes without adding any fake smoothness or a continuous metric to the categorical domain.

The resulting landscape is rugged and completely discrete. It has a single best configuration, but the best configurations are only a few small objective gaps apart. Because the domain is limited, it is possible to find the true global optimum and its objective value by listing all 1024 configurations, which makes oracle-based evaluation possible.

3.2.3. Empirical Protocol

We used MNL-BO with two acquisition strategies that work in a categorical setting: EI_{MNL}, which is the probability-weighted expected improvement, and UCB_{MNL}, which is the probability-weighted upper confidence bound. The budget for the evaluation is set at

$T = 60$, and the first random design has $n_0 = 10$ elements. All experiments are structured as maximization problems.

Preference Construction Strategy

The MNL surrogate is trained in this benchmark by going through a series of "choice situations". At iteration t, the choice set includes all the configurations that have been looked at so far. The chosen alternative is the configuration that has the best observed objective value in that set. This creates one long-format choice block for each iteration, which means that we don't have to make all pairwise comparisons directly.

3.2.4. Single-Run Performance

Table 2 shows convergence statistics for a typical run of each acquisition strategy. In particular, it lists the best objective value at the start, the best final value after the full budget, the total improvement made, and the number of strict incumbent improvements seen during the run.

Table 2. Single-run convergence summary on the purely categorical benchmark ($T = 60$, $n_0 = 10$).

Method	T	Init. Best	Final Best	Total Improv.	# Improv.
EI_{MNL}	60	10.301	22.511	12.210	3
UCB_{MNL}	60	10.168	22.507	12.339	3

Figure 2 shows the same representative run by plotting the best objective value so far against the number of iterations. For reference, the oracle optimum, which was found by exhaustive enumeration, is shown as a dashed horizontal line.

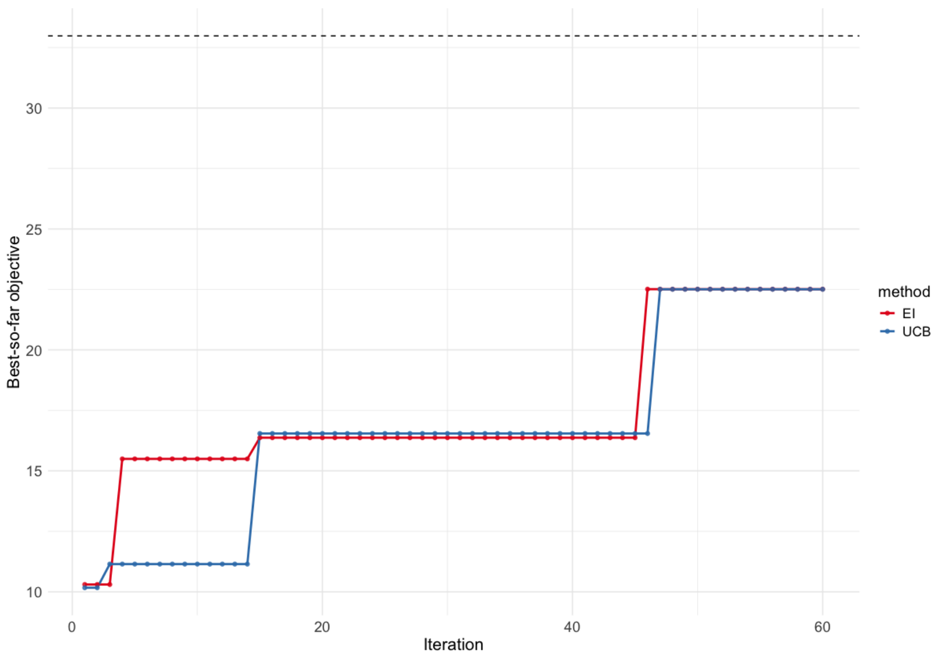

Figure 2. Best-so-far objective value over iterations for EI_{MNL} and UCB_{MNL} on the purely categorical benchmark. The dashed line indicates the oracle optimum.

3.2.5. Multi-Run Robustness

To evaluate robustness concerning random initialization, we conduct the experiment over 30 independent runs for each acquisition strategy. Table 3 shows the average and standard deviation of the initial best value, the final best value, and the total improvement over all runs.

Table 3. Multi-run performance statistics (mean ± standard deviation) over 30 independent runs.

Method	Mean Init. Best	SD Init. Best	Mean Final Best	SD Final Best	Mean Improv.
EI_{MNL}	8.777	7.152	22.810	1.111	14.033
UCB_{MNL}	6.288	6.193	23.408	2.899	17.120

Figure 3 shows the mean best-so-far objective value over iterations together with ±1 standard deviation bands, providing a compact summary of convergence speed and across-run variability.

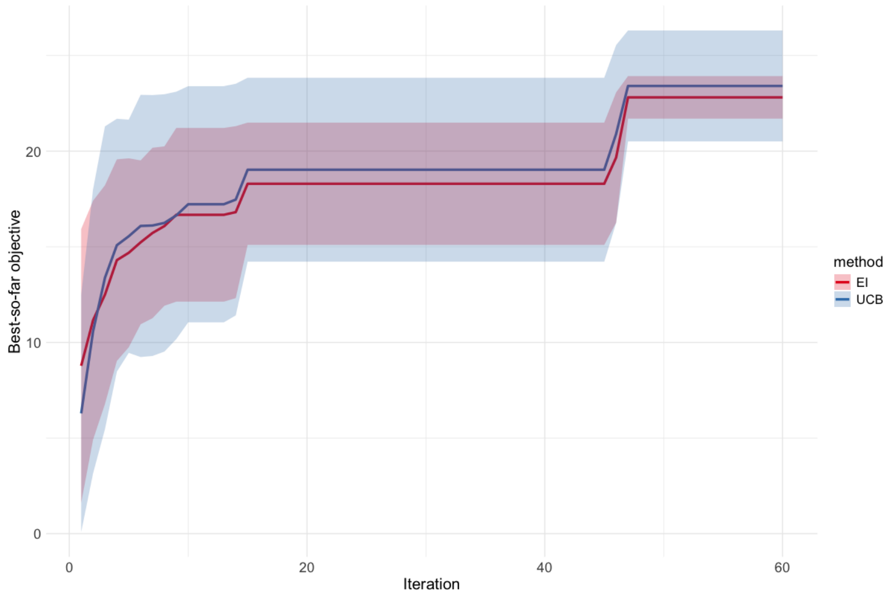

Figure 3. Mean best-so-far objective value over 30 runs for EI_{MNL} and UCB_{MNL}. Shaded regions indicate ±1 standard deviation.

3.2.6. Surrogate Calibration

Figure 4 assesses the quality of the acquired surrogate model. The predicted latent utility is used to put configurations into quantile bins. For each bin, the mean noise-free objective value is shown with error bars. This calibration-style test shows that, on average, higher predicted utilities lead to better true objective values.

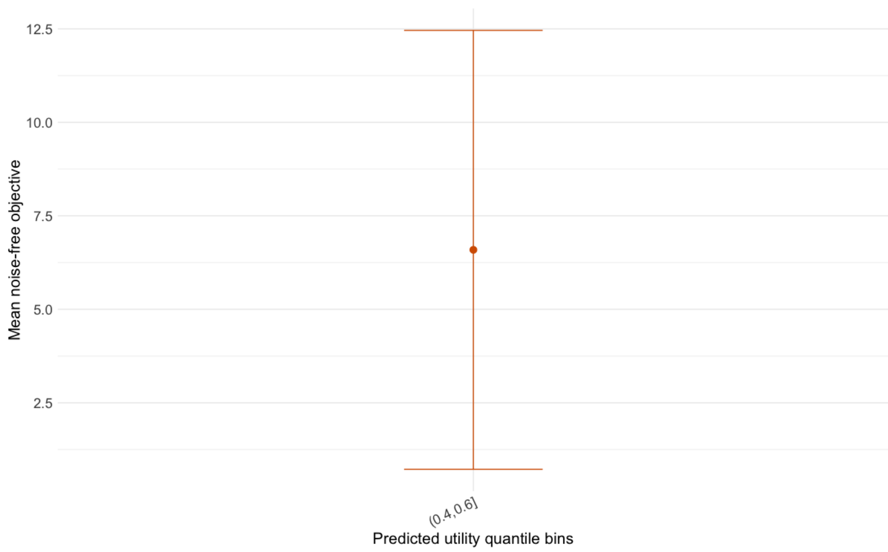

Figure 4. Surrogate calibration on the categorical benchmark. Predicted utility bins versus mean noise-free objective value with error bars.

3.2.7. Discussion

The results show that MNL-BO can successfully navigate a search space that is only made up of categories. The single-run and multi-run analyses demonstrate consistent enhancement compared to the random initial design, whereas the surrogate-calibration analysis validates that the acquired utility model effectively represents a significant preference structure across categorical configurations. These results confirm that MNL surrogates are appropriate for Bayesian optimization in completely discrete domains.

Positioning Relative to State-of-the-Art Categorical BO Methods

The results on the categorical benchmark allow us to situate MNL-BO more concretely within the broader landscape of methods for categorical and mixed-variable optimization. Methods such as CoCaBO [19] and GRYFFIN [44] employ GP-based surrogates with specialized kernels or kernel approximations designed to handle categorical variables. These approaches offer greater modeling flexibility, particularly when the objective function exhibits complex nonlinear interactions among variables, since the GP prior can capture such structure nonparametrically.

MNL-BO is expected to perform comparably or better than these alternatives in problems where the following characteristics are present:

- **Purely categorical or mixed discrete–continuous domains with low-to-moderate interaction order.** The additive-plus-pairwise structure of the synthetic benchmark in Equation (12) is a natural fit for the MNL utility model, which captures main effects and pairwise interactions efficiently. In such settings, the parametric structure of MNL avoids the kernel specification and hyperparameter estimation overhead required by GP-based methods.
- **Small-to-moderate evaluation budgets.** The MNL surrogate requires only preference comparisons rather than absolute function values, making it more data-efficient in settings where only ordinal information is reliably available. In contrast, GP-based methods such as CoCaBO rely on exact function observations to fit the posterior, which may be less robust when evaluations are noisy or only approximately comparable.
- **Constrained mixed-variable problems with categorical decisions.** As demonstrated in the pressure vessel case study, MNL-BO integrates naturally with penalty-based constraint handling and outperforms RF-SMAC baselines under the same budget. GP-based methods with mixed-variable kernels can in principle handle such settings, but the kernel must be carefully designed to respect the interaction structure between continuous and categorical components.

Conversely, MNL-BO is expected to be at a disadvantage relative to CoCaBO, GRYFFIN, or GP-based methods in the following situations:

- **High-order nonlinear interactions.** The MNL surrogate assumes a linear utility structure and satisfies the IIA property, which limits its ability to model complex higher-order dependencies. On the TSP benchmark, where the objective involves deeply nonlinear interactions among positions and cities, RF-SMAC outperforms MNL-BO, consistent with this limitation.
- **Large continuous subspaces with smooth structure.** When the search space contains high-dimensional continuous variables with smooth objective landscapes, GP-based surrogates can exploit gradient-like information through kernel smoothness assumptions. MNL-BO does not exploit such structure, and its performance may degrade relative to GP-based methods in such settings.
- **Very large candidate spaces without effective operators.** As discussed in the scalability limitations, when the feasible set is exponentially large and no effective candidate-

generation operator is available, the performance of MNL-BO is tightly coupled to the quality of the candidate pool, which may limit its applicability.

Overall, MNL-BO should be viewed as a complementary approach within the categorical BO landscape, particularly well-suited to problems with moderate interaction structure, mixed categorical-continuous variables, and settings where preference-based modeling is natural. Direct empirical comparison with CoCaBO and GRYFFIN on shared benchmarks is a valuable direction for future work.

3.3. Traveling Salesman Problem

The Traveling Salesman Problem (TSP) is a standard benchmark in combinatorial optimization with a large literature in operations research and heuristic optimization [45,46]. The classical objective is to find the shortest Hamiltonian cycle through a set of cities, visiting each city exactly once and returning to the origin. More broadly, the underlying permutation structure appears in scheduling, sequencing, layout design, and related planning problems [47].

In this work, the TSP serves as an application-oriented case study for evaluating MNL-BO when the decision variable is a *permutation*. Because the proposed surrogate is designed to handle categorical structure, each tour is represented using a position–city categorical encoding, and the MNL model is trained using preference observations derived from tour comparisons.

3.3.1. Problem Formulation

We consider $n = 15$ major U.S. cities:

$$\mathcal{V} = \{\text{Atlanta, Boston, Chicago, Dallas, Denver, Houston, Las Vegas, Los Angeles,}\\ \text{Miami, New Orleans, New York, Phoenix, San Francisco, Seattle, Washington D.C.}\}.$$

A tour is a permutation $\mathbf{x} = (x_1, \ldots, x_n)$ of the cities in $\mathcal{V}$. The cost of a tour is the total travel distance along the cycle:

$$L(\mathbf{x}) = \sum_{i=2}^{n} C_{x_{i-1},x_i} + C_{x_n,x_1}, \tag{13}$$

where $C \in \mathbb{R}_{\geq 0}^{n\times n}$ is a symmetric distance matrix and $C_{a,b}$ is the distance between cities a and b. In our instance every city is connected to every other city, so all permutations are feasible. The search space has cardinality $n!$, which for $n = 15$ exceeds 10^{12}. This combinatorial explosion makes exhaustive search infeasible and motivates the use of sample-efficient optimization strategies such as Bayesian optimization.

To keep tables and figures readable, tours are reported using short city codes. The mapping is given in Table 4.

Table 4. City abbreviations used in TSP tables and figures.

Code	City	Code	City	Code	City
ATL	Atlanta	BOS	Boston	CHI	Chicago
DAL	Dallas	DEN	Denver	HOU	Houston
LV	Las Vegas	LA	Los Angeles	MIA	Miami
NO	New Orleans	NY	New York	PHX	Phoenix
SF	San Francisco	SEA	Seattle	WDC	Washington D.C.

Distance matrix. Table 5 reports the inter-city distance matrix used in the experiments.

Table 5. Inter-city distance matrix C (miles) used for the 15-city TSP instance.

	ATL	BOS	CHI	DAL	DEN	HOU	LV	LA	MIA	NO	NY	PHX	SF	SEA	WDC
ATL	0	1095	715	805	1437	844	1920	2230	675	499	884	1832	2537	2730	657
BOS	1095	0	983	1815	1991	1886	2500	3036	1539	1541	213	2664	3179	3043	44
CHI	715	983	0	931	1050	1092	1500	2112	1390	947	840	1729	2212	2052	695
DAL	805	1815	931	0	801	242	1150	1425	1332	504	1604	1027	1765	2122	1372
DEN	1437	1991	1050	801	0	1032	885	1174	2094	1305	1780	836	1266	1373	1635
HOU	844	1886	1092	242	1032	0	1525	1556	1237	365	1675	1158	1958	2348	1443
LV	1920	2500	1500	1150	885	1525	0	289	2640	1805	2486	294	573	1188	2568
LA	2230	3036	2112	1425	1174	1556	289	0	2757	1921	2825	398	403	1150	2680
MIA	675	1539	1390	1332	2094	1237	2640	2757	0	892	1328	2359	3097	3389	1101
NO	499	1541	947	504	1305	365	1805	1921	892	0	1330	1523	2269	2626	1098
NY	884	213	840	1604	1780	1675	2486	2825	1328	1330	0	2442	3036	2900	229
PHX	1832	2664	1729	1027	836	1158	294	398	2359	1523	2442	0	800	1482	2278
SF	2537	3179	2212	1765	1266	1958	573	403	3097	2269	3036	800	0	817	2864
SEA	2730	3043	2052	2122	1373	2348	1188	1150	3389	2626	2900	1482	817	0	2755
WDC	657	44	695	1372	1635	1443	2568	2680	1101	1098	229	2278	2864	2755	0

To provide a visual representation of the spatial relationships among the cities, Figure 5 illustrates their approximate geographic positions and the inter-city connections used in the TSP instance. This visualization helps interpret the structure of the distance matrix: geographically closer cities correspond to smaller distances, whereas long-range connections reflect larger values.

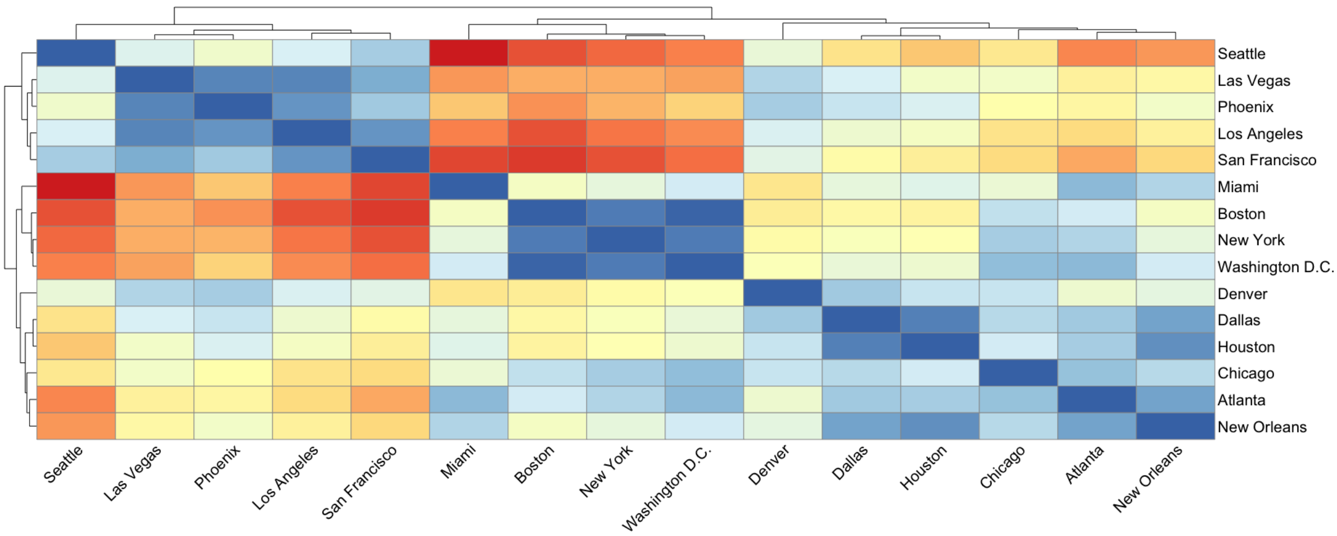

Figure 5. Visualization of inter-city connections. The spatial layout provides an intuitive interpretation of the distances reported in Table 5.

3.3.2. MNL-BO Setup for Permutations

We cast the TSP as a utility-maximization problem by defining the latent utility of a tour as

$$u(\mathbf{x}) = -L(\mathbf{x}),$$

so that shorter tours correspond to higher utility. Each tour is encoded using a position–city one-hot representation: for each position $p \in \{1, \dots, n\}$ and city $v \in \mathcal{V}$, the corresponding feature indicates whether city v is assigned to position p in the tour. This encoding is purely categorical and is therefore aligned with the surrogate assumptions.

Permutation Feasibility During Acquisition Maximization

Although the surrogate uses a position–city one-hot representation, acquisition maximization is *not* performed over a relaxed assignment space. Instead, the acquisition function is evaluated only on a finite candidate pool of valid permutations generated by permutation-preserving operators applied to the incumbent tour, together with globally random tours. Because every candidate is explicitly constructed as a permutation of the city list, infeasible tours containing repeated or missing cities cannot be produced during acquisition optimization.

The MNL surrogate is trained on preference observations generated from small *choice sets* sampled from the archive of evaluated tours. For a choice set of K candidate tours, the winner is the tour with minimum distance, equivalently maximum latent utility, and the conditional logit model is fitted to explain the winning tour as a function of the position features. Uncertainty is obtained from the asymptotic covariance of the maximum-likelihood estimator via the delta method, yielding $\sigma(\mathbf{x})$ for any candidate tour.

At each BO iteration, candidate tours are produced using a mixture of (i) local permutation moves applied to the current incumbent (2-opt, swap, and insertion), and (ii) globally random tours. Because acquisition optimization is restricted to this finite candidate batch, the scalability of the overall method depends directly on how effectively these operators cover promising regions of the exponentially large permutation space. The next tour is selected by maximizing either $\mathrm{UCB}_{\mathrm{MNL}}$ or $\mathrm{EI}_{\mathrm{MNL}}$ over the candidate batch and is then evaluated by computing its route length.

3.3.3. Experimental Setting and Reported Outputs

We initialize the archive with $n_{\mathrm{init}} = 30$ randomly generated tours and then perform $n_{\mathrm{iter}} = 150$ sequential BO iterations, for a total of 180 evaluated tours in a representative run. At each iteration, the MNL surrogate is trained using preference-style choice sets sampled from the current archive (40 sets of size $K = 5$), and the acquisition step is optimized over a candidate pool of up to 120 tours produced by a mixture of local operators and global random permutations. Because the objective is to minimize total travel distance, lower values indicate better solutions.

Table 6 demonstrates some of the main checkpoints that were used during optimization. At each checkpoint, we give the distance of the tour chosen at that point and the best distance so far. This behavior is typical of BO in combinatorial domains: the selected tour can change substantially as new options are explored, while the best-so-far value improves more gradually as better incumbents are identified, as illustrated in Figure 6.

Table 6. TSP progress table for a representative run. *Selected distance* is the tour length evaluated at the checkpoint iteration; *Best-so-far distance* is the best tour length found up to that iteration (lower is better).

Iteration	Selected Distance	Best-so-far Distance
1	17,905	16,527
5	17,426	16,527
10	25,344	16,257
25	18,068	15,984
50	14,048	13,295
75	26,492	12,732
100	14,346	12,732
125	12,940	12,614
150	16,095	12,081

https://doi.org/10.3390/a19050361

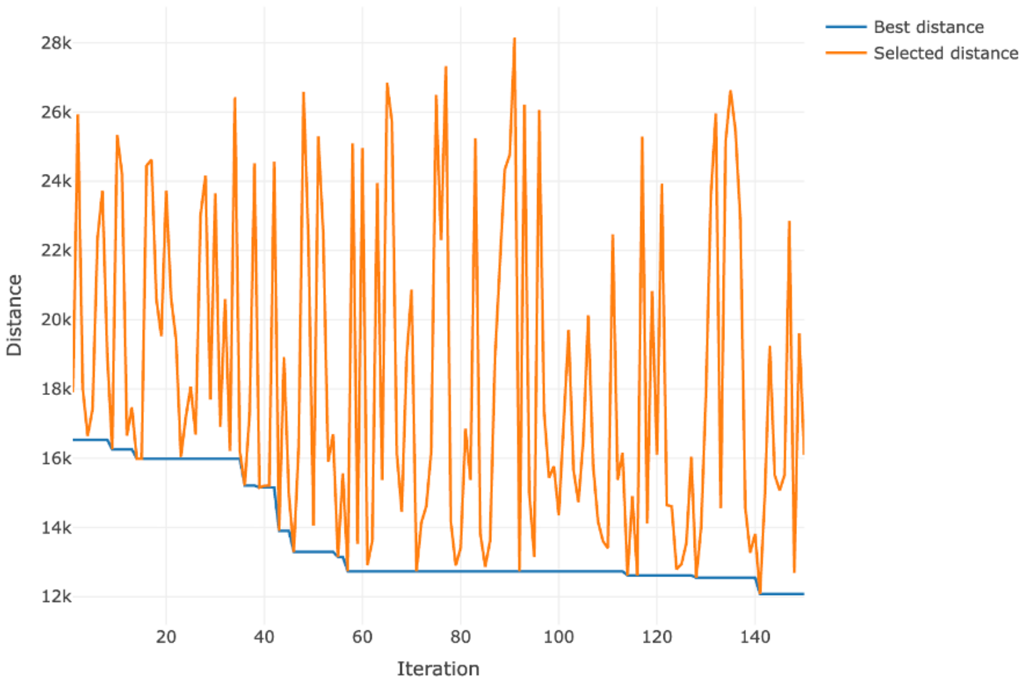

Figure 6. Representative TSP convergence behavior: selected tour distance and best-so-far distance over iterations.

Table 7 reports the best tour found in the representative run. The route is written as a closed cycle using the short city codes defined earlier.

Table 7. Best tour found in the representative run (lower is better). The route is reported as a closed cycle.

Best Distance	Tour (Cycle)
12,081	DAL → ATL → WDC → NY → BOS → CHI → DEN → LA → PHX → LV → SF → SEA → MIA → NO → HOU → DAL

Near-optimal tours. Table 8 shows the 10 best unique tours that were found during the run. This provides a clearer view of the solution landscape explored during optimization. Beyond identifying the single best solution, it illustrates that the method concentrates on a relatively small set of high-quality tours when making decisions. The structure of the best tour is visualized in Figure 7.

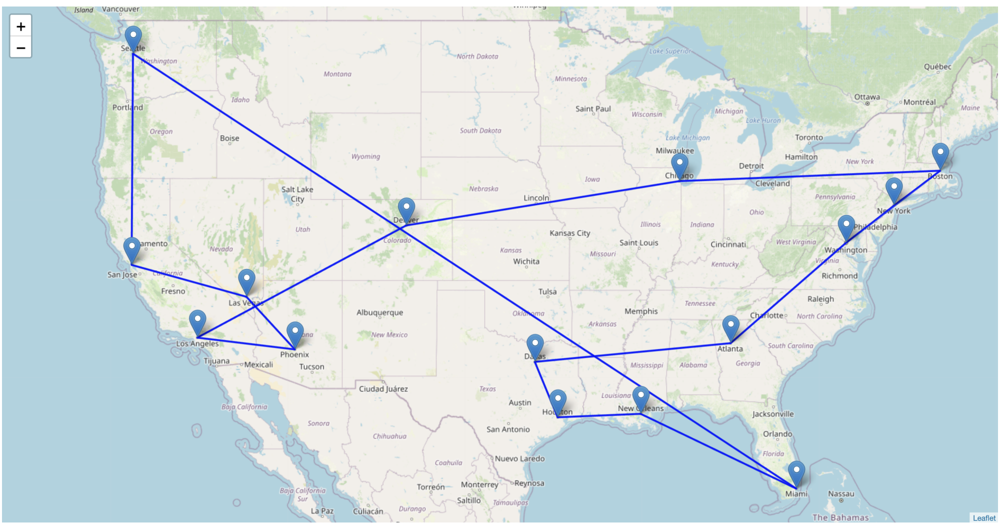

Figure 7. Map visualization of the best tour found. Cities are connected in visit order with directional edges.

Table 8. Top 10 unique tours encountered during the run, ranked by total distance (lower is better). Tours are shown using short city codes.

Rank	Distance	Tour
1	12081	DAL→ATL→WDC→NY→BOS→CHI→DEN→LA→PHX→LV→SF→SEA→MIA→NO→HOU
2	12546	DAL→ATL→WDC→SEA→SF→LV→PHX→LA→DEN→CHI→BOS→NY→MIA→NO→HOU
3	12614	DAL→ATL→WDC→SEA→SF→LV→PHX→LA→DEN→CHI→NY→BOS→MIA→NO→HOU
4	12614	DAL→ATL→BOS→NY→CHI→DEN→LA→PHX→LV→SF→SEA→WDC→MIA→NO→HOU
5	12688	MIA→SEA→SF→LV→PHX→LA→DEN→CHI→BOS→NY→WDC→ATL→DAL→NO→HOU
6	12732	DAL→ATL→WDC→SEA→SF→PHX→LV→LA→DEN→CHI→NY→BOS→MIA→NO→HOU
7	12732	DAL→ATL→WDC→CHI→DEN→LA→LV→PHX→SF→SEA→NY→BOS→MIA→NO→HOU
8	12792	DAL→NO→MIA→BOS→NY→CHI→DEN→LA→PHX→LV→SF→SEA→WDC→ATL→HOU
9	12866	DAL→MIA→BOS→NY→CHI→DEN→LA→LV→PHX→SF→SEA→WDC→ATL→NO→HOU
10	12902	DAL→CHI→DEN→LA→LV→PHX→SF→SEA→WDC→ATL→NY→BOS→MIA→NO→HOU

3.3.4. Multi-Run Robustness and Baseline Comparison

To assess robustness to random initialization and candidate generation, we repeat the TSP experiment over $R = 20$ independent runs with different random seeds. For computational efficiency, the multi-run study uses a reduced budget of $n_{\text{init}} = 30$ initial tours and $n_{\text{iter}} = 100$ sequential iterations. The representative single-run analysis in the previous subsection uses a larger budget ($n_{\text{iter}} = 150$) to illustrate longer-term convergence behavior.

In addition to the two proposed MNL-BO variants (EI and UCB), we compare against six baselines under the same evaluation budget: uniform random search, 2-opt local search, simulated annealing, restarted hill climbing, and two tree-based Bayesian optimization baselines inspired by the Sequential Model-based Algorithm Configuration (SMAC) framework using random-forest surrogates with EI and UCB acquisition criteria. For fairness, the surrogate-based methods use the same position–city feature representation and the same candidate-generation mechanism based on local permutation operators and random tours.

Table 9 summarizes the multi-run results. Among all methods considered, RF-SMAC (EI) achieves the lowest mean final tour length (10816.25), followed by RF-SMAC (UCB) and 2-opt local search. The proposed MNL-BO variants remain competitive, with MNL-BO (UCB) outperforming MNL-BO (EI) in terms of mean final performance and both substantially outperforming random search. Among the classical heuristic baselines, restarted hill climbing performs better on average than simulated annealing, although both remain less effective than the strongest surrogate-based and local-search methods on this instance.

Table 9. TSP multi-run results over $R = 20$ runs (lower is better).

Method	Mean Final Best	SD Final Best	Best (min)
MNL-BO (EI)	12,028.15	1114.61	9790.00
MNL-BO (UCB)	11,922.00	987.52	10073.00
Random Search	16,121.30	1015.47	13445.00
2-opt Local Search	11,533.20	1203.30	9181.00
Simulated Annealing	13,396.45	1519.18	10,996.00
Restarted Hill Climbing	12,521.90	1187.57	10,205.00
RF-SMAC (EI)	10,816.25	1061.57	9232.00
RF-SMAC (UCB)	11,271.05	499.68	10,230.00

These results provide several useful insights. First, the proposed MNL-BO framework consistently improves over the evaluation budget and clearly outperforms uninformed random search. Second, performance depends on problem structure: on this TSP instance, RF-SMAC and 2-opt local search achieve better average final performance than the proposed MNL-BO variants. This highlights the strength of tree-based surrogates and permutation-specific local improvement methods for route-optimization problems. At the same time, the proposed MNL-BO approach remains competitive while providing a probabilistically interpretable categorical surrogate model. Overall, the results indicate that MNL-BO is a viable and robust alternative for combinatorial Bayesian optimization, although it is not uniformly superior across all benchmark classes considered in this study. We acknowledge that the current comparison does not include dedicated categorical BO methods such as CoCaBO [19] or GRYFFIN [44]. A direct empirical evaluation against these methods on shared benchmarks remains an important direction for future work and would allow a more definitive assessment of the relative standing of MNL-BO within the broader categorical BO landscape.

The convergence curves in Figure 8 further illustrate these trends. Surrogate-based Bayesian optimization methods (MNL-BO and RF-SMAC) exhibit steady improvement throughout the optimization process, while random search quickly stagnates after the initial evaluations. Heuristic methods such as 2-opt local search show strong early improvements but tend to plateau as the search progresses. Overall, the results indicate that surrogate-based approaches can provide steady improvement over the evaluation budget, although their relative performance depends on the underlying problem structure and on the availability of strong problem-specific heuristics.

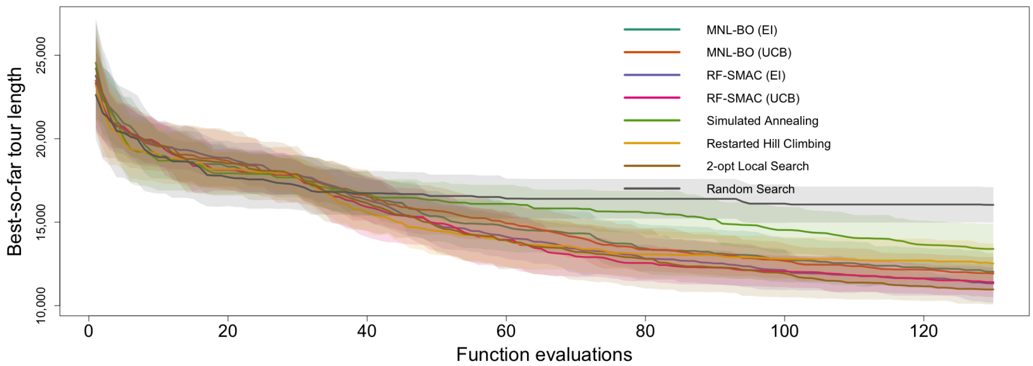

Figure 8. TSP multi-run convergence comparison across MNL-BO, RF-SMAC, and heuristic baselines. Curves show mean best-so-far tour length with ±1 standard deviation bands over $R = 20$ runs.

Role of Candidate Generation and Scalability

In the TSP experiments, candidate solutions are generated using standard permutation-preserving local search operators such as 2-opt, swap, and insertion, together with a smaller number of globally random tours. These operators define the neighborhood structure explored at each BO iteration and are commonly used in combinatorial optimization.

It is important to distinguish between candidate generation and candidate selection in the proposed framework. The role of the heuristic operators is to produce a finite batch of feasible candidate tours, while the MNL-based surrogate and acquisition function determine which candidate is evaluated next. Thus, the heuristics and the surrogate model serve complementary purposes: the former controls which part of the search space is exposed to the optimizer, whereas the latter guides exploration–exploitation within that exposed subset.

This distinction is also central to scalability. Because acquisition maximization is performed over a finite candidate set rather than over the full permutation space, the effectiveness of MNL-BO in very large combinatorial problems depends strongly on the quality and diversity of the candidate-generation mechanism. As the search space grows exponentially, a weak or overly local operator may fail to expose promising regions, in which case the surrogate cannot recover them regardless of how informative the acquisition function is. In this sense, scalability is limited not only by surrogate quality but also by the ability of the candidate generator to propose informative and diverse feasible solutions.

Empirically, baseline methods such as random search and local search, which rely on similar move operators but do not use a surrogate-driven selection strategy, show inferior performance. This indicates that the improvement achieved by MNL-BO is not solely due to the use of heuristic operators but also to the surrogate-guided exploration–exploitation mechanism. However, designing effective candidate-generation operators for new discrete domains remains a domain-dependent challenge and an important limitation of the current framework. More scalable variants could incorporate adaptive neighborhoods, learned proposal mechanisms, or hybrid search strategies; a detailed ablation of candidate-generation strategies is left for future work.

3.4. Pressure Vessel Design with Material Selection

We conclude the empirical evaluation with a real-world engineering optimization problem involving pressure vessel design with material selection. This example is representative of mixed-variable optimization tasks in which categorical decisions interact with continuous and discrete design parameters. The problem is a modified version of the classical pressure vessel benchmark originally proposed by [48], augmented here to explicitly include categorical material variables.

3.4.1. Problem Description

The objective is to minimize the total manufacturing cost of a cylindrical pressure vessel with hemispherical heads. The vessel components are joined using a single 60° weld, as illustrated in Figure 9. The total cost includes material costs for the shell and heads as well as welding costs, which depend on the material selected for the shell.

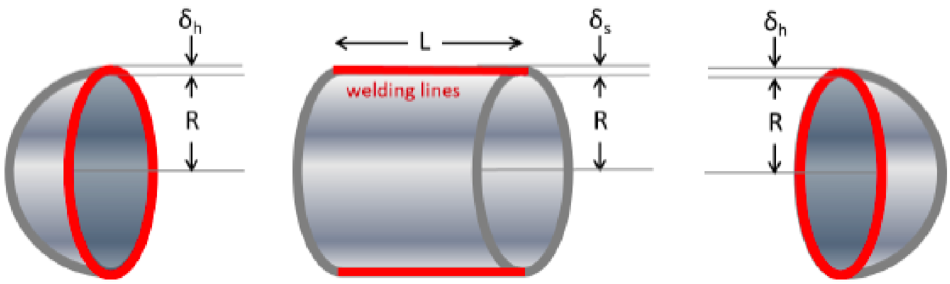

Figure 9. Components and dimensions of the pressure vessel.

The design variables include:

- discrete thickness indices for the shell and heads,
- continuous geometric parameters (radius and length),
- categorical material selections for the shell and heads.

3.4.2. Mathematical Formulation

The optimization problem is defined as

$$\begin{aligned}
\min_{n_s,n_h,R,L,y_s,y_h} \quad & 2\pi R(\rho_s C_s \delta_s L + \rho_h C_h \delta_h R) + \frac{4\pi\delta_s^2 \rho_s C_w}{9}(L + 2\pi R) \\
\text{s.t.} \quad & 0.0193R - \delta_s \le 0, \\
& 0.00954R - \delta_h \le 0, \\
& 1.296 \times 10^6 - \pi R^2 L - \frac{4}{3}\pi R^3 \le 0, \\
& \delta_s = 0.0625\, n_s, \quad \delta_h = 0.0625\, n_h, \\
& n_s, n_h \in \mathbb{N}, \\
& 0 < L \le 240, \quad R > 0, \\
& y_s \in \{\text{Steel A}, \text{Steel B}, \text{Steel C}\}, \\
& y_h \in \{\text{Steel A}, \text{Steel B}, \text{Steel C}\}.
\end{aligned} \tag{14}$$

Here, ρ denotes material density, C_s and C_h denote the costs of rolled and forged plates, respectively, and C_w denotes the welding cost. The categorical variables y_s and y_h specify the materials used for the shell and heads.

Mixed-Variable Feature Representation in the Logit Surrogate

For the pressure vessel problem, the surrogate is a conditional logit model whose feature map $\phi(\mathbf{x})$ includes both continuous and categorical covariates. Specifically, we use normalized numeric features for the thickness indices and geometric variables (e.g., $n_s/n_s^{\max}$, $n_h/n_h^{\max}$, $R/R^{\max}$, $L/L^{\max}$), together with one-hot indicators for the categorical material variables (y_s, y_h) and their interaction terms. This formulation is fully compatible with the logit likelihood and does not require discretizing (R, L) into bins.

3.4.3. Material Properties

The candidate construction materials and their associated properties are reported in Table 10. These material-dependent parameters directly affect both the objective function and feasibility through the cost and density terms.

Table 10. Material properties used in the pressure vessel design problem.

Material	Density ρ [lb/in^3]	Rolled Plate C_s [\$/lb]	Forged Plate C_h [\$/lb]	Welding C_w [\$/lb]
Steel A	0.2830	0.35	1.00	8.00
Steel B	0.2836	0.30	1.05	8.50
Steel C	0.2844	0.40	0.95	7.50

3.4.4. Optimization Results and Analysis

The proposed preference-based MNL-BO framework was applied to the pressure vessel problem using a fixed evaluation budget. Both single-run behavior and multi-run robustness were analyzed to assess convergence characteristics, feasibility handling, and final solution quality.

Single-Run Convergence Behavior

Table 11 reports representative checkpoints from a single optimization run, showing the selected design cost and the incumbent best-so-far cost. Owing to the presence of hard geometric and structural constraints, a substantial fraction of candidate designs sampled during early iterations are infeasible and therefore incur a large penalty cost. Nevertheless, the algorithm progressively identifies feasible regions and consistently improves the incumbent solution over time.

Table 11. Pressure vessel optimization progress at selected iterations for a representative run.

Iteration	Selected Cost	Best-so-far Cost
1	1,220,961.32	11,833.32
5	3,166,923.02	11,833.32
10	17,203.16	11,833.32
25	3,738,115.26	11,276.36
50	1,191,713.59	11,276.36
75	1,237,795.64	11,276.36
100	1,459,423.32	11,097.77
125	1,066,798.12	11,097.77
150	1,403,842.67	10,955.68

The decreasing trend in the best-so-far cost indicates that the MNL surrogate is able to balance exploration of new regions with exploitation of promising feasible designs despite the highly constrained and mixed-variable nature of the problem.

Constraint Handling Within the MNL-BO Loop

Constraints are incorporated into the MNL-BO framework through a penalty-based approach that operates at the preference-construction stage. Specifically, before any pairwise or choice-based preferences are assembled, each evaluated configuration is assigned a penalized objective value according to

$$\tilde{f}(\mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } \mathbf{x} \text{ is feasible,} \\ f_{\max} + \lambda \cdot v(\mathbf{x}) & \text{otherwise,} \end{cases} \tag{15}$$

where $f_{\max}$ is the worst feasible cost observed so far, $\lambda = 10^6$ is the penalty coefficient, and $v(\mathbf{x}) \geq 0$ is the total constraint violation. These penalized values $\tilde{f}(\mathbf{x})$ are then used in place of raw objective values when constructing preference comparisons. As a result, in any choice set containing both feasible and infeasible configurations, the feasible option is always preferred, regardless of its absolute cost.

This design ensures that the preference-learning model implicitly learns to favor the feasible region without requiring explicit constraint encoding in the surrogate. The MNL surrogate is fitted entirely on the penalized preference data and therefore treats feasibility as an intrinsic component of the utility structure.

Potential bias and mitigation. The penalty may introduce a bias in the surrogate's preference structure: when the archive contains very few feasible configurations relative to infeasible ones, the surrogate utility landscape is dominated by the penalty signal rather than by the fine-grained ranking among feasible designs. In the pressure vessel experiments, this effect is limited in practice because the feasible and infeasible regions are well separated geometrically. As shown in Table 11, the algorithm progressively identifies feasible solutions and transitions from infeasibility-driven exploration to feasibility-aware optimization after approximately 25–50 iterations. A more principled alternative, such as constrained acquisition functions or separate feasibility surrogates, is left for future work.

Best design obtained. The best pressure vessel configuration identified in the representative run is summarized in Table 12. The solution corresponds to a shell made from Steel A and heads made from Steel C, with moderate radius and length values and discrete thickness levels satisfying all structural constraints.

Table 12. Best pressure vessel design obtained in a representative run.

Best Cost	n_s	n_h	δ_s	δ_h	R	L	**Materials** (y_s/y_h)
10,955.68	18	22	1.125	1.375	49.98	99.21	Steel A / Steel C

This result illustrates MNL-BO's capability to concurrently optimize continuous geometric dimensions, discrete thickness variables, and categorical material selections within a cohesive framework.

Multi-run robustness. To check for robustness, the optimization was done again with the same settings but with 10 different random seeds. Table 13 shows the average and standard deviation of the best feasible costs at the beginning and end of the runs, as well as the average improvement made.

Table 13. Pressure vessel optimization results over 10 independent runs (mean ± standard deviation).

Init. Best	**Final Best**	**Improvement**	**Infeasible Rate**	**Best (min)**
12,505.8 ± 2683.0	8037.1 ± 1332.0	4468.7 ± 2534.6	0.758 ± 0.055	6503.80

Feasibility reporting. We provide feasibility-aware statistics because infeasible designs get a big penalty. We keep track of the "best feasible" objective that comes up in each run separately. The infeasible evaluation rate is the number of designs that break at least one rule out of all the designs that were tested.

Interaction with the preference learning model. The penalty approach interacts with the MNL surrogate in two ways. First, during early iterations when most evaluated configurations are infeasible, the surrogate learns primarily to distinguish between levels of infeasibility rather than to rank feasible designs by quality. This is reflected in the high infeasibility rate of 0.758 ± 0.055 reported in Table 13. Second, once feasible configurations begin to appear in the archive, the large utility gap induced by the penalty causes the surrogate to sharply concentrate probability mass on feasible regions, effectively biasing candidate selection toward the feasible part of the design space.

Convergence and variability across runs. Figure 10 displays the optimal cost trajectories for all runs, accompanied by their average trend. Even though each run has a different speed of convergence, all of the trajectories show steady improvement and eventually reach the same level of solution quality.

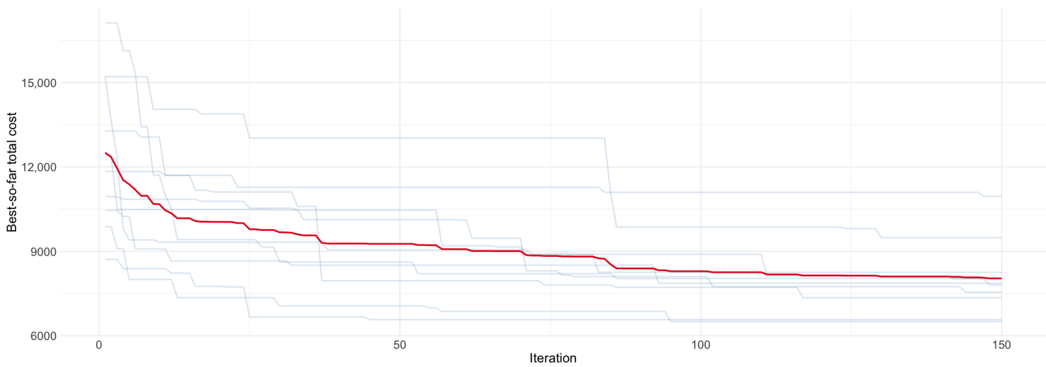

Figure 10. Best-so-far cost trajectories over 10 independent runs. Thin lines indicate individual runs, while the bold line shows the mean trend.

Figure 11 summarizes the distribution of final best costs across runs using a violin-plus-boxplot representation, highlighting both central tendency and variability.

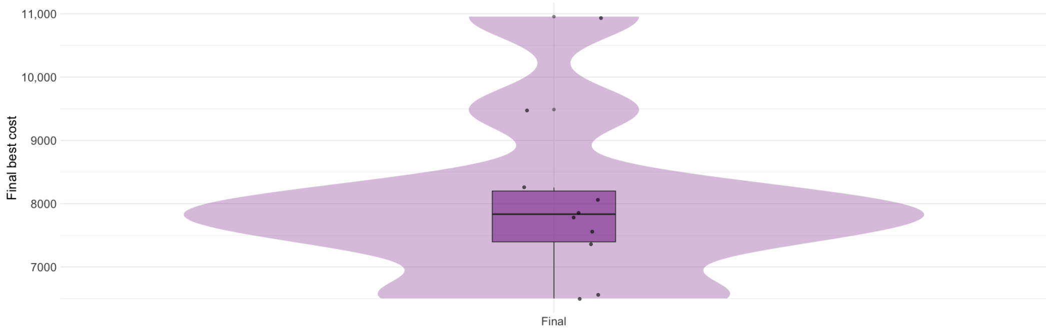

Figure 11. Distribution of final best costs across 10 runs.

3.4.5. Comparison with Tree-Based Bayesian Optimization (RF-SMAC)

To establish an additional baseline consistent with established mixed-variable optimization methodologies, we compare MNL-BO to a SMAC-style tree-based surrogate method. In this baseline, a random-forest surrogate is trained to predict the penalized objective cost from the mixed-variable feature representation, and the next design is selected using standard acquisition criteria for minimization, namely expected improvement and a lower confidence bound.

All methods are evaluated under the same budget and experimental protocol as in the previous subsection: 10 independent runs with different random seeds, $n_{\text{init}} = 50$ initial designs, $n_{\text{iter}} = 150$ sequential iterations, identical candidate-generation procedures, and the same constraint-handling penalty.

Table 14 summarizes the multi-run results. MNL-BO (UCB) achieves the best overall performance on this benchmark, with the lowest mean final best cost (8037.05) and the best minimum value observed across runs (6503.80). Both RF-SMAC variants are competitive, but converge to higher mean final costs, with RF-SMAC (EI) reaching 9450.75 on average and RF-SMAC (LCB) reaching 10,410.98.

Table 14. Pressure vessel multi-run results over $R = 10$ runs (lower is better).

Method	Mean Final Best	SD Final Best	Best (min)
MNL-BO (UCB)	8037.05	1332.02	6503.80
RF-SMAC (EI)	9450.75	1445.84	7456.89
RF-SMAC (LCB)	10,410.98	1306.79	8200.60

These results suggest that, for this constrained mixed-variable engineering-design task, the preference-based MNL surrogate can effectively guide the search toward high-quality feasible regions and can outperform a strong tree-based surrogate baseline under the same evaluation budget.

Overall, these results show that the proposed preference-based MNL-BO framework is effective for this mixed-variable engineering design problem, exhibiting strong convergence behavior, robustness in the presence of infeasible regions, and consistent identification of high-quality feasible designs.

3.5. Computational Overhead

We also report the computational overhead of the proposed method by timing (i) surrogate fitting (conditional logit via `mlogit`), (ii) acquisition evaluation on the candidate set, and (iii) candidate generation. Timings are measured in R (version 4.6.0) using `the system.time` and averaged across iterations and runs.

This overhead is compared against representative baselines that do not require repeated logit fitting, such as random search and local-search heuristics, and against tree-based BO baselines when available. These measurements help contextualize the additional computational cost incurred by repeated preference-model training.

Table 15 shows the average computational cost of each main algorithmic component per iteration, such as surrogate fitting, acquisition evaluation, and candidate generation. Overall, the cost of fitting the MNL surrogate is still low compared to the total optimization runtime and is modest compared to the expensive objective evaluations that are common in real engineering design problems. These findings indicate that the suggested preference-based surrogate model is computationally feasible for the moderate evaluation budgets typically employed in Bayesian optimization research.

Table 15. Average computational overhead per iteration.

Component	Mean Time (s)	SD (s)
MNL surrogate fitting	0.18	0.04
Acquisition evaluation	0.05	0.01
Candidate generation	0.02	0.01
Total overhead	0.25	0.05

4. Concluding Remarks

This research examined a preference-oriented Bayesian optimization framework referred to as MNL-BO (Multinomial Logit Bayesian Optimization) for optimization challenges involving categorical and mixed decision variables. The suggested method uses a multinomial logit model that was trained on pairwise preference comparisons instead of Gaussian process surrogates. This provides a natural way to model decision structures that are categorical or combinatorial.

We tested the proposed method on three optimization problems of increasing difficulty: a purely categorical benchmark (Infinity77), a combinatorial Traveling Salesman problem, and a constrained mixed-variable engineering design problem that required material selection for a pressure vessel. The framework consistently optimized well, and the optimal objective value steadily improved with each new evaluation in these case studies.

The categorical Infinity77 benchmark demonstrated that the preference-based surrogate could acquire useful representations of categorical decision spaces without the necessity for random numerical encodings. The Traveling Salesman Problem showed that the method works well in large combinatorial search spaces, where the algorithm found shorter tours and performed consistently across many independent runs. The proposed method consistently yielded designs that were both feasible and cost-effective in the pressure vessel design problem, characterized by nonlinear constraints and a mix of continuous, discrete, and categorical variables. The proposed MNL-BO method demonstrates its competitiveness against a tree-based surrogate baseline derived from the Sequential Model-based Algorithm Configuration (SMAC) framework, utilizing the same evaluation budget.

The findings demonstrate that the proposed framework has a number of benefits. First, using pairwise comparisons to model preferences gives us a flexible substitute representation that can handle many decision variables. Second, the multinomial logit model's estimates of predictive utility and uncertainty make it easier to come up with acquisition methods that strike a good balance between exploration and exploitation. Third, the method moves smoothly from purely categorical problems to more realistic engineering optimization tasks that involve a mix of factors. This shows that it works well for a lot of different kinds of problems.

When optimization problems have categorical or mixed decision variables, the proposed preference-based Bayesian optimization framework is a clear and useful alternative to traditional surrogate models. The method preserves interpretability by directly modeling relative preferences instead of depending on arbitrary numerical representations of categorical variables, while also keeping the sequential decision-making benefits of Bayesian optimization.

4.1. Limitations

The proposed framework also has several limitations. The method depends on pairwise preference comparisons based on objective evaluations. If there is a lot of noise, this can lead to inconsistent preference relationships, which can make maximum-likelihood estimation less stable and uncertainty estimates less reliable. The multinomial logit model also assumes independence of irrelevant alternatives (IIA), but this may not be true in all cases. Finally, acquisition optimization is performed over a finite candidate set rather than over the full combinatorial search space. A key limitation of the current framework is that, in large combinatorial spaces, surrogate-based acquisition can only be as effective as the candidate-generation mechanism over which it is optimized. This represents a fundamental limitation of the current framework: for any new combinatorial domain, designing effective candidate-generation operators is a non-trivial, domain-specific task that requires substantial expertise and may significantly constrain the method's applicability in novel settings. Designing good operators is therefore essential and may require substantial domain knowledge in new application settings.

4.2. Sensitivity to Hyperparameters

The proposed framework involves several design choices, including the exploration parameter κ, the number of preference pairs or choice sets used per surrogate refit, and the size of the candidate pool used during acquisition optimization. In the present study, these parameters are fixed across experiments based on commonly used defaults and preliminary tuning, as summarized in Table 1.

Based on the structure of the acquisition functions and the preference-learning surrogate, we can offer the following qualitative guidance on the sensitivity of the framework to these choices. The exploration parameter $\kappa = 2.0$ follows the convention established in GP-UCB [34] and controls the balance between exploitation of high-utility regions and exploration of uncertain ones. Larger values of κ encourage broader exploration and may be beneficial in highly multimodal landscapes, while smaller values concentrate search near the current best estimate. The candidate pool size determines the diversity of configurations exposed to the acquisition function at each iteration: pools that are too small may miss promising regions, while very large pools increase computational cost without proportional benefit. The number of preference pairs or choice sets per surrogate refit controls the amount of comparative information available to the MNL estimator at each iteration: too few comparisons may yield unreliable utility estimates, particularly early in the search, while too many increase the cost of surrogate training. In our experiments, the chosen values struck a practical balance across the problem classes considered.

A formal sensitivity analysis and adaptive or problem-specific strategies for hyperparameter selection represent important directions for future work. Such an analysis would clarify the robustness of MNL-BO across a wider range of problem settings and budget regimes.

4.3. Comparison with Additional BO Baselines

While the experimental evaluation includes several classical and state-of-the-art baselines, including RF-SMAC, we note that other widely used Bayesian optimization methods for mixed-variable problems, such as GP-based frameworks (e.g., BoTorch), Tree-structured Parzen Estimators (TPE), and GRYFFIN, are not included in the current comparison.

The primary objective of this work is to introduce and study a new surrogate modeling approach based on multinomial logit models, rather than to provide a fully exhaustive benchmarking study. In this context, RF-SMAC serves as a strong reference baseline for discrete and mixed-variable optimization.

We observe that RF-SMAC can outperform the proposed method in certain settings. However, the MNL-BO framework offers complementary advantages, including interpretability, a natural treatment of categorical variables, and a preference-based formulation. A more comprehensive empirical comparison with additional BO methods is an important direction for future work.

4.4. Future Research Directions

The current study presents multiple avenues for future research. First, this work includes comparisons with a tree-based surrogate baseline inspired by the SMAC framework. However, broader benchmarking against other discrete Bayesian optimization methods would help to better understand the strengths of the proposed method. Specifically, comparisons with Tree-Parzen estimators, latent embedding models for categorical variables, and categorical kernel-based Gaussian process methods would yield a more thorough empirical assessment.

The existing framework focuses on single-objective optimization. Extending the MNL-based surrogate model to multi-objective Bayesian optimization appears to be a prudent strategy. Integrating preference-based multinomial logit models with Pareto-front learning and multi-objective acquisition procedures may expedite the resolution of issues involving conflicting objectives and diverse decision factors.

Third, scaling the proposed framework to larger combinatorial and high-dimensional categorical search spaces remains an important direction for future research. In such settings, exhaustive evaluation of all candidate configurations during acquisition optimization becomes infeasible. To address this, future work may explore candidate subsampling strategies, adaptive neighborhood generation, and hybrid search mechanisms that restrict acquisition evaluation to a subset of promising configurations.

A further direction concerns the automated design and adaptation of candidate-generation operators. In the current framework, operators such as 2-opt moves and insertion heuristics are hand-crafted for each problem domain, requiring substantial domain expertise. Learning these operators from data—for instance through reinforcement learning, surrogate-guided mutation strategies, or meta-learning across related problem instances—could substantially reduce the domain expertise required to deploy MNL-BO in new combinatorial settings and improve its scalability to larger search spaces. This direction is closely related to the broader challenge of automated algorithm configuration and represents an important step toward a more general and self-contained optimization framework.

In addition, scalable preference data construction and more efficient surrogate updates will be important for maintaining computational efficiency. The parametric nature of the MNL surrogate provides a favorable foundation for such extensions, as it avoids some of the computational overhead associated with nonparametric models. These developments would enable the framework to be applied to more complex large-scale optimization problems.

A comprehensive theoretical analysis of preference-based Bayesian optimization using discrete-choice surrogates would strengthen the methodological foundations of the methodology. Investigating convergence properties, regret behavior, and robustness for noisy or inconsistent preferences represents a crucial avenue for future research.

Author Contributions: Conceptualization, methodology, formal analysis, software, investigation, data curation, writing—original draft preparation, and writing—review and editing, M.A.S.; supervision, methodological guidance, validation, critical revision of the manuscript, and overall scientific oversight, A.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research did not receive any additional funding.

Data Availability Statement: All calculations for the application examples are included in the manuscript. The complete source code and supporting data for the empirical study will be made publicly available through the authors' GitHub repository: https://github.com/MuhammadAmirSaeed66/MNL-BO, accessed on 28 April 2026.

Acknowledgments: The authors gratefully acknowledge the University of Milano-Bicocca and the Italian Government for providing a fully funded Ph.D. scholarship that made this research possible.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Garnett, R. *Bayesian Optimization*; Cambridge University Press: Cambridge, MA, USA, 2023.
2. Frazier, P.I. Bayesian optimization. In *Recent Advances in Optimization and Modeling of Contemporary Problems*; INFORMS: Catonsville, MD, USA, 2018; pp. 255–278.
3. Baratchi, M.; Wang, C.; Limmer, S.; Van Rijn, J.N.; Hoos, H.; Bäck, T.; Olhofer, M. Automated machine learning: Past, present and future. *Artif. Intell. Rev.* **2024**, *57*, 122.
4. Hutter, F.; Kotthoff, L.; Vanschoren, J. *Automated Machine Learning: Methods, Systems, Challenges*; Springer Nature: Cham, Switzerland, 2019.
5. Bergstra, J.; Bardenet, R.; Bengio, Y.; Kégl, B. Algorithms for Hyper-Parameter Optimization. In Proceedings of the Advances in Neural Information Processing Systems, Granada, Spain, 12–17 December 2011.
6. Mockus, J.; Tiesis, V.; Zilinskas, A. The application of Bayesian methods for seeking the extremum. In *Towards Global Optimization*; Elsevier: Amsterdam, The Netherlands, 1978; Volume 2, pp. 117–129.
7. Osborne, M.A.; Garnett, R.; Roberts, S.J. Gaussian processes for global optimization. In Proceedings of the Third International Conference on Learning and Intelligent Optimization (LION 3), Trento, Italy, 14–18 January 2009.
8. Pelamatti, J.; Brevault, L.; Balesdent, M.; Talbi, E.; Guerin, Y. Efficient global optimization of constrained mixed variable problems. *J. Glob. Optim.* **2019**, *73*, 583–613.
9. Cuesta-Ramirez, J.; Le Riche, R.; Roustant, O.; Perrin, G.; Durantin, C.; Gliere, A. A comparison of mixed-variables Bayesian optimization approaches. *Adv. Model. Simul. Eng. Sci.* **2021**, *9*, 6.
10. Garrido-Merchán, E.C.; Hernández-Lobato, D. Dealing with categorical and integer-valued variables in bayesian optimization with gaussian processes. *Neurocomputing* **2020**, *380*, 20–35.
11. Zhang, Y.; Tao, S.; Chen, W.; Apley, D. A latent variable approach to Gaussian process modeling with qualitative and quantitative factors. *Technometrics* **2020**, *62*, 291–302.
12. Roustant, O.; Padonou, E.; Deville, Y.; Clément, A.; Perrin, G.; Giorla, J.; Wynn, H. Group kernels for Gaussian process metamodels with categorical inputs. *SIAM J. Uncertain. Quantif.* **2020**, *8*, 775–806.
13. Deng, X.; Lin, C.; Liu, K.; Rowe, R. Additive Gaussian process for computer models with qualitative and quantitative factors. *Technometrics* **2017**, *59*, 283–292.
14. Saves, P.; Lafage, R.; Bartoli, N.; Diouane, Y.; Bussemaker, J.; Lefebvre, T.; Hwang, J.T.; Morlier, J.; Martins, J.R. SMT 2.0: A surrogate modeling toolbox with a focus on hierarchical and mixed variables Gaussian processes. *Adv. Eng. Softw.* **2024**, *188*, 103571.
15. Balandat, M.; Karrer, B.; Jiang, D.; Daulton, S.; Letham, B.; Wilson, A.; Bakshy, E. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–12 December 2020; Volume 33, pp. 21524–21538.
16. Audet, C.; Le Digabel, S.; Rochon Montplaisir, V.; Tribes, C. Algorithm 1027: NOMAD Version 4: Nonlinear Optimization with the MADS Algorithm. *ACM Trans. Math. Softw.* **2022**, *48*, 35.
17. Luce, R.D. *Individual Choice Behavior: A Theoretical Analysis*; John Wiley & Sons: Hoboken, NJ, USA, 1959.
18. So, Y.; Kuhfeld, W.F. Multinomial logit models. In Proceedings of the SUGI 20, Orlando, FL, USA, 2–5 April 1995; Volume 1995, pp. 1227–1234.
19. Ru, B.; Alvi, A.; Nguyen, V.; Osborne, M.A.; Roberts, S. Bayesian optimisation over multiple continuous and categorical inputs. In Proceedings of the International Conference on Machine Learning, Virtual, 13–18 July 2020; PMLR: London, UK, 2020; pp. 8276–8285.
20. Hutter, F.; Hoos, H.H.; Leyton-Brown, K. Sequential model-based optimization for general algorithm configuration. In Proceedings of the Learning and Intelligent Optimization Conference (LION), Rome, Italy, 17–21 January 2011; Springer: Berlin/Heidelberg, Germany, 2011; Volume 6683, pp. 507–523.
21. Bergstra, J.; Yamins, D.; Cox, D.D. Hyperopt: A Python Library for Optimizing the Hyperparameters of Machine Learning Algorithms. In Proceedings of the SciPy Conference, Austin, TX, USA, 24–29 June 2013; Volume 13, p. 20.

22. Abramson, M.A.; Audet, C.; Dennis, J. Filter pattern search algorithms for mixed variable constrained optimization problems. *Pac. J. Optim.* **2007**, *3*, 477–500.
23. Liu, D.; Zhang, H. Residuals and diagnostics for ordinal regression models: A surrogate approach. *J. Am. Stat. Assoc.* **2018**, *113*, 845–854.
24. Li, S.; Fan, Z.; Liu, I.; Morrison, P.S.; Liu, D. Surrogate method for partial association between mixed data with application to well-being survey analysis. *Ann. Appl. Stat.* **2024**, *18*, 2254–2276.
25. Archetti, F.; Candelieri, A. *Bayesian Optimization and Data Science*; Springer: Cham, Switzerland, 2019; Volume 849.
26. Shahriari, B.; Snoek, J.; Adams, R.P.; de Freitas, N. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE* **2016**, *104*, 148–175. https://doi.org/10.1109/JPROC.2015.2494218.
27. McFadden, D. Conditional Logit Analysis of Qualitative Choice Behavior. In *Frontiers in Econometrics*; Zarembka, P., Ed.; Academic Press: New York, NY, USA, 1974; pp. 105–142.
28. Train, K.E. *Discrete Choice Methods with Simulation*, 2nd ed.; Cambridge University Press: Cambridge, MA, USA, 2009.
29. Bradley, R.A.; Terry, M.E. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika* **1952**, *39*, 324–345.
30. Agresti, A. *Categorical Data Analysis*, 3rd ed.; Wiley: Hoboken, NJ, USA, 2013.
31. Cattelan, M. Models for Paired Comparison Data: A Review with Emphasis on Dependent Data. *Stat. Sci.* **2012**, *27*, 412–433.
32. Chu, W.; Ghahramani, Z. Preference Learning with Gaussian Processes. In Proceedings of the 22nd International Conference on Machine Learning, Bonn, Germany, 7–11 August 2005; pp. 137–144.
33. Brochu, E.; Cora, V.M.; de Freitas, N. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. *arXiv* **2010**, arXiv:1012.2599.
34. Srinivas, N.; Krause, A.; Kakade, S.M.; Seeger, M.W. Gaussian process optimization in the bandit setting: No regret and experimental design. In Proceedings of the 27th International Conference on Machine Learning, Haifa, Israel, 21–25 June 2010.
35. Bull, A.D. Convergence rates of efficient global optimization algorithms. *J. Mach. Learn. Res.* **2011**, *12*, 2879–2904.
36. Hernández-Lobato, J.M.; Hoffman, M.W.; Ghahramani, Z. Predictive Entropy Search for Efficient Global Optimization of Black-box Functions. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 8–13 December 2014; Volume 27.
37. Roy, S.; Crossley, W.A.; Stanford, B.; Moore, K.T.; Gray, J.S. A Mixed Integer Efficient Global Optimization Algorithm with Multiple Infill Strategy Applied to Wing Topology Optimization Problem. In Proceedings of the AIAA Scitech Forum, San Diego, CA, USA, 7–11 January 2019.
38. R Core Team. *R: A Language and Environment for Statistical Computing*; R Foundation for Statistical Computing: Vienna, Austria, 2024.
39. Croissant, Y. Estimation of Multinomial Logit Models in R: The `mlogit` Package. *J. Stat. Softw.* **2013**, *50*, 1–24.
40. Dowle, M.; Srinivasan, A. data.table: Extension of 'data.frame'. In *R Package Documentation*; R Foundation for Statistical Computing: Vienna, Austria, 2024.
41. Wickham, H. *ggplot2: Elegant Graphics for Data Analysis*; Springer: Berlin/Heidelberg, Germany, 2016.
42. Lawler, E.L.; Lenstra, J.K.; Rinnooy Kan, A.H.G.; Shmoys, D.B. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*; Wiley: Hoboken, NJ, USA, 1985.
43. Kirkpatrick, S.; Gelatt, C.D.; Vecchi, M.P. Optimization by simulated annealing. *Science* **1983**, *220*, 671–680.
44. Häse, F.; Aldeghi, M.; Hickman, R.J.; Roch, L.M.; Aspuru-Guzik, A. Gryffin: An algorithm for Bayesian optimization of categorical variables informed by expert knowledge. *Appl. Phys. Rev.* **2021**, *8*, 031406.
45. Laporte, G. A concise guide to the traveling salesman problem. *J. Oper. Res. Soc.* **2010**, *61*, 35–40. https://doi.org/10.1057/jors.2009.76.
46. Toth, P.; Vigo, D. An overview of vehicle routing problems. In *The Vehicle Routing Problem*; Toth, P., Vigo, D., Eds.; Society for Industrial and Applied Mathematics: Philadelphia, PA, USA, 2002; pp. 1–26.
47. Georgiadis, G.P.; Elekidis, A.P.; Georgiadis, M.C. Optimization-based scheduling for the process industries: From theory to real-life industrial applications. *Processes* **2019**, *7*, 438.
48. Sandgren, E. Nonlinear Integer and Discrete Programming in Mechanical Design Optimization. *J. Mech. Des.* **1990**, *112*, 223–229. https://doi.org/10.1115/1.2912596.