
Modeling Multi-Objective Tradeoffs with Monotonic Utility Functions

Edward Chen¹

Natalie Dullerud¹

Thomas Niedermayr¹

Elizabeth Kidd¹

Ransalu Senanayake³

Pang Wei Koh²

Sanmi Koyejo¹

Carlos Guestrin¹

¹Stanford University, ²University of Washington, ³Arizona State University

Abstract

Countless science and engineering applications in multi-objective optimization (MOO) necessitate that decision-makers (DMs) select a Pareto-optimal (PO) solution which aligns with their preferences. Evaluating individual solutions is often expensive, and the high-dimensional trade-off space makes exhaustive exploration of the full Pareto frontier (PF) infeasible. We introduce a novel, principled two-step process for obtaining a compact set of PO points that aligns with user preferences, which are specified *a priori* as general *monotonic utility functions* (MFs). Our process (1) densely samples the user’s region of interest on the PF, then (2) sparsifies the results into a small, diverse set for the DM. We instantiate this framework with *soft-hard functions* (SHFs), an intuitive class of MFs that operationalizes the common expert heuristic of imposing soft and hard bounds. We provide extensive empirical validation of our framework instantiated with SHFs on diverse domains, including brachytherapy, engineering design, and large language models. For brachytherapy, our approach returns a compact set of points with over 3% greater SHF-defined utility than the next best approach. Among the other domains, our approach consistently leads in utility, as a final compact set of just 5 points captures over 99% of the utility offered by the entire dense set.

1 INTRODUCTION

Optimizing across multiple competing objectives — multi-objective optimization (MOO) — is a fundamental challenge in critical real-world applications, from engineering design and drug discovery to deep learning and healthcare (Fromer and Coley, 2023; Luukkonen et al., 2023; Xie et al., 2021; Yu et al., 2000; Papadimitriou and Yannakakis, 2001). In these scenarios, objectives often conflict, and evaluating potential solutions can be computationally expensive or require expert knowledge. The ultimate goal is typically not to identify the entire set of optimal trade-offs (the Pareto frontier, PF), but rather to efficiently present a *compact and relevant set* of high-quality solutions from which a decision-maker (DM) can select one that best aligns with their latent preferences for the ideal balance among objectives. The complexity of the trade-off space and the cognitive burden of sifting through numerous options make exhaustive PF exploration impractical. DMs often possess valuable prior knowledge to guide the search, yet effectively incorporating these priors to identify a manageable and useful subset of the PF remains a significant hurdle, especially when solution evaluations are costly.

This challenge is particularly acute in high-stakes domains, where practitioners frame objectives in terms of utility rather than simple maximization. A decision-maker’s satisfaction with an outcome is often a monotonic and bounded function of the objective value (Huber et al., 2025; Bassett, 1987; Sobrie et al., 2018). This functional class flexibly models diverse expert priors, from diminishing returns to sharp satisfaction thresholds. In healthcare, for instance, improving a clinical metric from a dangerous to an acceptable level yields high utility, while further improvements may offer only marginal gains (Mandal and Sutradhar, 2019). We argue that a powerful way to incorporate such expert priors is to have DMs specify their preferences *a priori* via a monotonic utility function (MF) for each objective. This approach directly translates domain knowledge

into the optimization process. As another example, in brachytherapy, clinicians must balance tumor coverage against healthy organ exposure (Deufel et al., 2020), a task where they naturally reason with *desired targets* ("soft bounds") and *strict limits* ("hard bounds") (Viswanathan et al., 2012). This dual-level thinking is a utility structure explicitly modeled by a specific class of MFs we term *soft-hard functions* (SHFs).

Current MOO techniques, while advanced, often do not directly incorporate this intuitive and flexible class of MFs to guide the search for a compact, decision-relevant set of Pareto-optimal (PO) points. While it is well-understood that exploring the entire PF can be intractable and unnecessary (Paria et al., 2019; Zuluaga et al., 2016), and many methods aim to find representative PF subsets or elicit preferences through mechanisms like pairwise comparisons or utility function modeling (Paria et al., 2019; Abdolshah et al., 2019; Suzuki et al., 2020; Zuluaga et al., 2016; While et al., 2012), there remains a gap in methods that allow DMs to *directly articulate their preferences using these explicit monotonic utilities for each objective.* This direct articulation can constrain the search space effectively and align the presented solutions more closely with the DM’s operational thinking, especially for domain experts who already possess such priors. Our work aims to fill this gap. We assume the DM possesses underlying (unknown) trade-off preferences, which are implicitly constrained by their specified MFs. Given that evaluating each PO point is costly, we employ a two-step process: (1) efficient dense sampling of the PF region consistent with current bounds using Bayesian optimization, and (2) sparsification of these points using robust submodular optimization *to present a small, diverse, and high-utility set to the DM.* We demonstrate that this sparsification step maintains theoretical guarantees on near-optimality with respect to the DM’s unknown preferences encoded by the MFs.

Although a rich body of work exists on incorporating preferences into MOO (Malkomes et al., 2021; Abdolshah et al., 2019; Zuluaga et al., 2016), many methods do not allow for the *a priori* specification of flexible, nonlinear utility functions. Existing preference elicitation schemes often rely on iterative feedback, such as pairwise comparisons (Paria et al., 2019; Suzuki et al., 2020), which may be sample-inefficient or not fully leverage an expert’s pre-existing, well-defined knowledge. In contrast, our framework uniquely focuses on *directly using the DM’s specified utility functions as the primary mechanism for guiding exploration and obtaining a decision-relevant set of PO solutions.* Our contributions are:

1. A two-step process for finding a compact set of

PO points that aligns with a DM’s priors, specified via general monotonic utility functions (MFs). The process (1) uses multi-objective Bayesian optimization to densely sample the relevant PF region, then (2) employs robust submodular optimization to sparsify these points into a small, diverse set. This theoretically-grounded approach maintains guarantees on both search coverage and the near-optimality of the final set.

2. To demonstrate our framework’s practical utility, we introduce *soft-hard functions* (SHFs), a concrete class of MFs that models the common expert heuristic of imposing soft and hard bounds. We provide **extensive empirical validation** of our framework with SHFs on diverse domains, including brachytherapy, engineering design, and large language model personalization.

2 MULTI-OBJECTIVE OPTIMIZATION WITH MONOTONIC UTILITY FUNCTIONS

In this section, we outline the optimization setting we consider and introduce notation used throughout the paper in §2.1. We then introduce our concept of modeling preferences via general monotonic utility functions in §2.2. We thus arrive at the problem definition through our mini-max formulation in §2.3.

2.1 Multi-Objective Optimization Background

As the name suggests, a multi-objective optimization (MOO) problem is an optimization problem that concerns multiple objective functions. Any MOO problem can be written as the joint maximization of L objective functions over some input space $X \subseteq \mathbb{R}^d$,

$$\max_{x \in X} (f_1(x), \dots, f_L(x))$$

in which each f_ℓ , $\ell \in [L]$, defines a function $f_\ell : X \rightarrow \mathbb{R}$. Broadly speaking, there does not typically exist a *feasible* solution that marginally optimizes each objective function simultaneously. Therefore, work in MOO generally focuses on Pareto-optimal (PO) solutions. A feasible solution x^\dagger is considered *Pareto-optimal* if no objective can improve without degrading another; in other words, if x^\dagger is not *Pareto-dominated* by any other solution (defined in Appendix D).

A common approach to multi-objective optimization is to convert the L -dimensional objective to a scalar in order to utilize standard optimization methods via

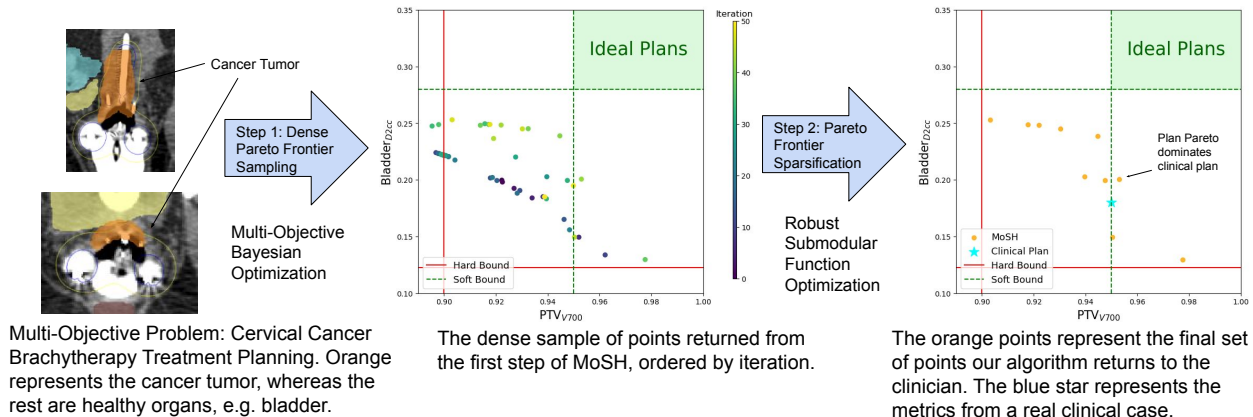


Figure 1: Our two-step pipeline, which we refer to as MoSH, demonstrated on a real cervical cancer brachytherapy clinical case. This figure instantiates our general framework using *soft-hard functions* (SHFs), a class of MFs that captures the natural way clinicians reason with "soft" targets and "hard" limits (§5.1). The bounds shown reflect those used in real clinical settings. The MOO task involves balancing tumor radiation dosage against minimizing exposure to nearby healthy organs; two objectives (to maximize) are shown. Since the clinician’s precise preferences are unknown, Step 1 densely samples the "Ideal Plans" region defined by the soft bounds. Step 2 then sparsifies this set into a small, diverse collection of high-utility options for the DM. Notably, our final set includes a plan that Pareto dominates the metrics from an actual clinical case, demonstrating the practical value of our approach.

a scalarization function. Scalarization functions typically take the form $s_{\lambda} : \mathbb{R}^L \rightarrow \mathbb{R}$, parameterized by λ from some set Λ in L -dimensional space (Roijs et al., 2013; Paria et al., 2019). For instance, the general class of linear scalarization functions $s_{\Lambda}(\mathbf{y}) := \{\lambda^T \mathbf{y} \mid \lambda \in \Lambda\}$ constitutes all convex combinations of the objectives in Λ . The parameters $\lambda \in \Lambda$ can be viewed as weights, or relative preferences, on the objective functions in the scalarized optimization objective $\max_{x \in X} s_{\lambda}([f_1(x), \dots, f_L(x)])$. Then, the advantage of using scalarization functions is that the solution to maximizing $s_{\lambda}([f_1(x), \dots, f_L(x)])$, for a fixed value of λ , may lead to a solution along the PF.

2.2 Monotonic Utility Functions

Many practical MOO problems, particularly in domains like engineering and healthcare, involve more nuanced desiderata than simply maximizing raw objective values (Cui et al., 2018; Grosman et al., 2010; Cairoli et al., 2019). For instance, improving an objective from a poor value to an acceptable one may yield high utility, while further improving it from a good value to an excellent one might offer only marginal gains (Deufel et al., 2020). Our framework aims to operationalize such nuanced priors by allowing a DM to specify their utilities for each objective *a priori* through bounded, monotonic utility functions (MFs).

An MF, $u_{f_{\ell}} : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty\}$, transforms the raw

output of an objective function $f_{\ell}(x)$ into a numerical representation of its usefulness. We require two fundamental properties for our framework. First, the function must be **monotonic**. A utility function is monotonically increasing if for any two objective values v_1 and v_2 , an improvement ($v_1 > v_2$) implies an equal or greater utility ($u_{f_{\ell}}(v_1) \geq u_{f_{\ell}}(v_2)$). This is a cornerstone of rational choice theory, ensuring that better objective performance is always preferred (Huber et al., 2025; Sobrie et al., 2018). Second, the utility function must be **bounded** from above, which models the realistic notion that the utility of any single objective does not grow infinitely (Bassett, 1987). This boundedness prevents a single objective from completely dominating all others in the optimization process.

By enabling DMs to specify these bounded, monotonic functions directly, our framework empowers them to encode their domain-specific knowledge into the optimization process from the outset. As these are the only strict requirements, effective utility functions can also model other notable priors, such as strongly penalizing infeasible solutions and capturing diminishing returns through concavity. A concrete and intuitive instantiation of a function possessing these properties, the *soft-hard function* (SHF), is detailed in §5.1 and used in our empirical validations.

2.3 Problem Definition

Given a selected class of scalarization functions s_{λ} parameterized by $\lambda \in \Lambda$ and a set of MFs as defined above, we wish to elucidate a set of useful points along the PF in the optimization problem

$$\max_{x \in X} s_{\lambda}(u_f(x))$$

where $u_f := [u_{f_1}, \dots, u_{f_L}]$. Since we only want to select points from the PF, we follow Roijers et al. (2013) and Paria et al. (2019) in assuming that $s_{\lambda}(u_f(x))$ is monotonically increasing in all coordinates $u_{f_{\ell}}(x)$, which leads to the solution to $\arg \max_{x \in X} s_{\lambda}(u_f(x))$ residing on the PF, for some λ . We assume the DM has some hidden set of preferences, λ^* , for which the solution to the optimization problem $\max_{x \in X} s_{\lambda^*}(u_f(x))$ represents the ideal solution on the PF. Therefore, we wish to return a set of Pareto optimal points, C , which contains the unknown $c^* = \arg \max_{x \in X} s_{\lambda^*}(u_f(x))$, i.e. the ideal solution that aligns with the DM preferences. Since λ^* is unknown to us, we want a set C which is robust against any potential value of λ^* , or weight (preference) from the DM. Furthermore, $|C| \leq k$ for some small integer value k to avoid overwhelming the DM with too many choices, which has been shown to decrease choice quality (Diehl, 2005).

As a result, we formulate our general problem of selecting a set of points from the regions defined by the MFs as:

$$\max_{C \subseteq X, |C| \leq k} \min_{\lambda \in \Lambda} \left[\frac{\max_{x \in C} s_{\lambda}(u_f(x))}{\max_{x \in X} s_{\lambda}(u_f(x))} \right] \quad (1)$$

We refer to the right term, $\max_{x \in C} s_{\lambda}(u_f(x)) / \max_{x \in X} s_{\lambda}(u_f(x))$, as the MF utility ratio. Intuitively, the MF utility ratio is maximized when the points in C are Pareto optimal and span the high utility regions of the PF, as defined by the MFs. We use an instantiation of the MF utility ratio as an evaluation metric for some of our experiments in Section 5. Given the unknown λ^* and the expensive cost of evaluating points, we solve Equation 1 via a principled two-step process, allowing us to address the distinct challenges of continuous exploration and discrete set selection with methods that each carry strong theoretical guarantees: (1) We first densely sample the PF to obtain a high-quality discrete representation that provides robust coverage against uncertainty in λ^* . (2) Then, we sparsify this dense set to find a small, diverse, and near-optimal set of points for the decision-maker. This approach ensures our final solution is supported by guarantees on both the initial search and the final set, a property difficult to achieve with a single, monolithic method.

3 STEP 1: DENSE PARETO FRONTIER SAMPLING WITH BAYESIAN OPTIMIZATION

In this section, we consider the goal of obtaining a dense set of Pareto optimal points. As described earlier, since the DM's preferences, λ^* , are unknown to us, we wish to obtain a set D which is diverse, high-coverage, and is modeled after the DM-defined MFs. As is typical in various science and engineering applications, we assume access to some noisy and expensive black-box function – often modeled with a Gaussian process (GP) (Williams and Rasmussen, 1995) as the surrogate function. To achieve that goal, we extend our formulation (1) into a Bayesian setting, assuming a prior $p(\lambda)$ with support Λ imposed on the set of Pareto optimal values and using the notion of random scalarizations (Paria et al., 2019). In this continuous and stochastic setting, we assume that each of the $\ell \in [L]$ objectives are sampled from known GP priors with a common domain, and produce noisy observations, e.g. $y_{\ell} = f_{\ell}(x) + \epsilon_{\ell}$, where $\epsilon \sim N(0, \sigma_{\ell}^2)$, $\forall \ell \in [L]$. We optimize over a set of scalarizations weighted by the prior $p(\lambda)$.

Our overall aim is still to obtain a set of points C on the PF which are robust against the worst-case potential λ^* , within the user-defined MF. For computational feasibility, however, we convert the worst-case into an average-case maximization (analyzed further in §5.2). Taking into consideration the aforementioned set of assumptions, we modify the formulation (1) to be the following:

$$\max_{D \subseteq X, |D| \leq k_D} \mathbb{E}_{\lambda \sim p(\lambda)} \left[\frac{\max_{x \in D} s_{\lambda}(u_f(x))}{\max_{x \in X} s_{\lambda}(u_f(x))} \right] \quad (2)$$

Random Scalarizations. We now describe our sampling-based algorithm to optimize Equation 2. Similar to Paria et al. (2019), we use the notion of random scalarizations to sample a λ from $p(\lambda)$ at each iteration which is then used to compute a multi-objective acquisition function based on s_{λ} and the MF. The maximizer of the multi-objective acquisition function is then chosen as the next sample input value to be evaluated with the expensive black-box function. For the acquisition function, we use $\text{acq}(u, \lambda_t, x) = s_{\lambda_t}(u_{\varphi(x)})$ where $\varphi(x) = \mu_t(x) + \sqrt{\beta_t} \sigma_t(x)$ and $\beta_t = \sqrt{0.125 \times \log(2 \times t + 1)}$ (Appendix B). For scalarization, we use the augmented Chebyshev scalarization, allowing for non-convex PF sampling (Wierzbicki, 1982). We note, however, that *our algorithm is agnostic to the scalarization and acquisition functions*. The full algorithm is described in Algorithm 1. Finally, we provide formal guarantees proving the lower bound of the MF utility ratio, which goes to one as $T \rightarrow \infty$, and

thus providing the dense set D (in Appendix D).

Algorithm 1 Step 1: Dense Pareto Frontier Sampling

- 1: Initialize MFs $\forall \ell \in [L]$
 - 2: Initialize $D^{(0)} = \emptyset, GP_\ell^{(0)} = GP(0, \kappa) \forall \ell \in [L]$
 - 3: **for** $t = 1 \rightarrow T$ **do**
 - 4: Obtain $\lambda_t \sim p(\lambda)$
 - 5: $x_t = \arg \max_{x \in X} \text{acq}(u, \lambda_t, x)$
 - 6: Obtain $y = f(x_t)$
 - 7: $D^{(t)} = D^{(t-1)} \cup \{(x_t, y)\}$
 - 8: $GP_\ell^{(t)} = \text{post}(GP_\ell^{(t-1)} | (x_t, y)) \forall \ell \in [L]$
 - 9: **end for**
 - 10: Return $D^{(T)}$
-

4 STEP 2: PARETO FRONTIER SPARSIFICATION

We now assume there already exists a dense set of points on the PF, sampled from step 1 with MFs. As DM validation of such a dense set would be costly, the goal for step 2 is now to sparsify that set of points to then present to the DM a more navigable set which still maintains as much utility as the dense set. We do so by leveraging the notion of diminishing returns in utility for each additional point the DM validates. This notion is encapsulated by the property of submodularity, which further allows us to design optimization algorithms with strong theoretical guarantees. We use the definition of submodularity first developed in Nemhauser et al. (1978) (see Appendix D for definition).

As a result, we adapt Equation 1 and formulate the sparsification problem as a robust submodular observation selection (RSOS) problem (Krause et al., 2008):

$$\max_{C \subseteq D, |C| \leq k} \min_{\lambda \in \Lambda} \underbrace{\left[\frac{\max_{x \in C} s_\lambda(u_f(x))}{\max_{x \in D} s_\lambda(u_f(x))} \right]}_{F_\lambda} \quad (3)$$

where D represents the PO dense set obtained from step 1 (Section 3, Algorithm 1) and C is the sparse set of MF-defined PO points returned to the DM. Equation 3 is not submodular, but can be viewed as a maximin over submodular functions F_λ (Lemma 1), studied by Krause et al. (2008).

Lemma 1. *For some DM preference value λ , the set function, which takes as input some MF-defined set C : $\max_{x \in C} s_\lambda(u_f(x)) / \max_{x \in D} s_\lambda(u_f(x))$ is normalized ($F_\lambda(\emptyset) = 0$), monotonic (for all $A \subseteq C \subseteq D, F_\lambda(A) \leq F_\lambda(C)$), and submodular. (See proof in Appendix D).*

Algorithm. If Equation 3 were submodular, the simple greedy algorithm would provide a near optimal so-

lution (Nemhauser et al., 1978) (theorem in Appendix C). Since it is not, the greedy algorithm performs arbitrarily worse than the optimal solution when solving Equation 3, or more generally, problems formulated as RSOS (Krause et al., 2008) - often defined as such:

$$\max_{C \subseteq X, |C| \leq k} \min_i F_i(C) \quad (4)$$

where the goal is to find a set C of observations which is robust against the worst possible objective, $\min_i F_i$, from a set of submodular objectives. As a result, we solve Equation 3 using the Submodular Saturation algorithm, or SATURATE, which provides us with strong theoretical guarantees on the optimality of set C (Krause et al., 2008).

Theorem 2. (Krause et al., 2008) *For any integer k , SATURATE finds a solution C_S such that $\min_i F_i(C_S) \geq \max_{|C| \leq k} \min_i F_i(C)$ and $|C_S| \leq \psi k$, for $\psi = 1 + \log(\max_{x \in D} \sum_i F_i(\{x\}))$. The total number of submodular function evaluations is $\mathcal{O}(|D|^2 m \log(m \min_i F_i(D)))$, where $m = |\Lambda|$.*

At a high level, SATURATE first defines a relaxed version of the original RSOS problem, which contains a superset of feasible solutions, and is guaranteed to find solutions to that relaxed version which are at least as informative as the optimal solution, only at a slightly higher cost. Within our context, this enables us to select a subset of points C , from the dense set obtained from Step 1, D , which achieves the optimal coverage of $\lambda \in \Lambda$ albeit with a slightly greater number of points. The complete algorithm is shown in Algorithm 2.

Algorithm 2 Step 2: Pareto Frontier Sparsification

- 1: **procedure** SPARSIFY($F_1, \dots, F_{|\Lambda|}, k, \psi$)
 - 2: $q_{min} \leftarrow 0; q_{max} \leftarrow \min_i F_i(D); C_{best} \leftarrow \emptyset$
 - 3: **while** $(q_{max} - q_{min}) \geq 1/|\Lambda|$ **do**
 - 4: $q \leftarrow (q_{min} + q_{max})/2$
 - 5: Define $\bar{F}_q(C) = \frac{1}{|\Lambda|} \sum_i \min\{F_i(C), q\}$
 - 6: $\hat{C} \leftarrow \text{GPC}(\bar{F}_q, q) \triangleright$ Appendix: Algorithm 3
 - 7: **if** $|\hat{C}| > \psi k$ **then**
 - 8: $q_{max} \leftarrow q$
 - 9: **else**
 - 10: $q_{min} \leftarrow q; C_{best} \leftarrow \hat{C}$
 - 11: **end if**
 - 12: **end while**
 - 13: **return** C_{best}
 - 14: **end procedure**
-

5 EXPERIMENTAL RESULTS

To empirically validate our framework, we instantiate it with a class of MFs designed for expert-driven, constrained domains. In this section, we introduce these functions and outline the details for our experiments.

5.1 Concrete MF: Soft-Hard Functions

To instantiate our general framework for empirical validation, we introduce a concrete and intuitive class of MFs called *soft-hard functions* (SHFs). These functions are designed to operationalize the dual-level preferences common in expert-driven domains like brachytherapy and diabetes care, where practitioners reason with both strict requirements ("hard bounds") and desirable targets ("soft bounds") (Deufel et al., 2020; Grosman et al., 2010; Cairoli et al., 2019). Traditional methods that only enforce hard constraints fail to capture this nuanced structure. SHFs explicitly model this prevalent DM heuristic by defining utility over these regions.

By construction, an SHF satisfies four key characteristics. First, objective values $f_\ell(x)$ below the hard bound $\alpha_{\ell,H}$ yield $-\infty$ utility, strongly penalizing infeasible solutions. Second, for $f_\ell(x) \geq \alpha_{\ell,S}$ (the soft bound), the utility function becomes concave, modeling diminishing returns once a satisfactory level is met. Third, beyond a saturation point $\alpha_{\ell,\tau}$ ($\geq \alpha_{\ell,S}$), utility becomes constant, preventing a single objective from dominating others. Finally, the function is monotonically increasing in $f_\ell(x)$, ensuring better objective values are preferred.

We note that SHFs maintain generality in the landscape of optimization problems, as the function can be understood as an instantiation of fairly standard components of an optimization problem: feasibility constraints on objectives (hard constraints), and constraints with slack variables (soft constraints). While optimization problems frequently do not include constraints on *every* objective function, our experiments and implementation consider the most constrained setting (soft and hard bounds on every objective). Nonetheless, flexibility is maintained in the general case for SHFs if the defined soft or hard bounds are effectively vacuous (i.e. no soft or hard constraint is enforced via the transformation) – for instance, $\alpha_{\ell,H} = -\infty$, $\alpha_{\ell,S} = \alpha_{\ell,\tau} = \infty$.

For simplicity and clear interpretability, we present a specific piecewise-linear functional form for SHFs in Figure 2. Additional exposition and the precise mathematical form are described in Appendix E.1. When our general two-step framework is instantiated with these SHFs, we refer to the overall method as **MoSH**. The first step, dense PF sampling, is called **MoSH-Dense**, and the second step, PF sparsification, is called **MoSH-Sparse**. Accordingly, we refer to the specialized MF utility ratio as the **SHF utility ratio**.

5.2 Baseline Methodologies

Step 1: Dense Pareto Frontier Sampling: We experiment with both synthetic and real-world appli-

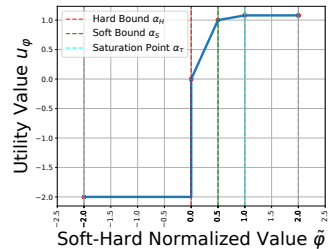


Figure 2: Example of a normalized soft-hard function (SHF) u . Dashed lines indicate hard, soft, and saturation point regions. Utility below the hard bound is $-\infty$ (shown as -2 for illustration).

cations and compare against other similar Bayesian multi-objective optimization approaches: Expected hypervolume improvement (EHVI) (Emmerich, 2008), ParEGO (Knowles, 2006), Multi-objective Bayesian optimization Using Random Scalarizations (MOBO-RS) (Paria et al., 2019), and random sampling. We compare against MOBO-RS with variations on the scalarization function, Chebyshev and linear, and acquisition function, UCB and Thompson sampling. Additional experimental results comparing against preference-based evolutionary multi-objective (PBEMO) algorithms may be found in Appendix E. *We also use SHFs to conduct empirical analyses of the worst-case to average-case conversion of Equation 2 in Appendix E.5.*

Step 2: Pareto Frontier Sparsification: We compare our method, MoSH-Sparse, against greedy and random algorithms. The greedy baseline starts with the empty set, and iteratively adds the element $c = \arg \max_{x \in D \setminus C} H(C \cup \{x\})$, where $H = \min_{\lambda \in \Lambda} F_\lambda$ for the F_λ in Equation 3, until some stopping point. For all experiments, additional details and figures are provided in the Appendix.

5.3 Performance Evaluation

5.3.1 Evaluation of Step 1: Dense Pareto Frontier Sampling

As detailed in Section 3, our goal for the dense point set D is to achieve diversity, high coverage, and alignment with the DM’s MFs. To conduct a concrete empirical analysis, we instantiate our general framework using SHFs from §5.1, allowing us to assess our method’s ability to handle this practical and intuitive class of expert priors. We introduce four soft-hard metrics specifically designed to measure performance with respect to the key structural properties of SHFs: their distinct soft (A_S) and hard (A_H) regions. These metrics quantify how well an algorithm achieves diversity, high coverage, and adherence to these user-defined regions.

Soft-Hard Fill Distance: Measures sample diversity across A_S and A_H . This adapts the fill distance by Malkomes et al. (2021): $\text{FILL}(C, D) = \sup_{x' \in D} \min_{x \in C} \kappa(f(x), f(x'))$, where C is the set of sampled points, D is a comprehensive reference set of points in the target region, and $\kappa(\cdot)$ is a distance metric (typically Euclidean). Our metric is a v -weighted average of such fill distances calculated separately for soft and hard regions, encouraging exploration in both. Lower values indicate better, more uniform coverage.

Soft-Hard Positive Samples Ratio: Measures adherence to SHF constraints via the v -weighted ratio of sampled points within soft (C_S) and hard (C_H) regions: $v(|C_S|/|C|) + (1 - v)(|C_H|/|C|)$.

Soft-Hard Hypervolume: Assesses sample coverage by calculating the hypervolume occupied by points within the PF regions defined by the soft bounds (e.g., relative to $r_S = (\alpha_{1,S}, \dots, \alpha_{L,S})$) and, separately, the hard bounds. The final metric is a v -weighted average of these two hypervolumes.

Soft Region Distance-Weighted Score: We seek to explicitly measure faithfulness to the high-utility regions of SHFs, the intersection of the soft bounds, by measuring the density of points. We calculate this using the following: $\sum_{x \in C} 1/\kappa(r_S, f(x))$, where κ is a measure of distance, typically Euclidean distance and r_S is defined above.

5.3.2 Evaluation of Step 2: Pareto Frontier Sparsification

To evaluate the set of points, C , returned to the DM, we simulate a λ^* using the heuristic $\lambda = \mathbf{u}/\|\mathbf{u}\|_1$, where $\mathbf{u}_\ell \sim \mathcal{N}(\alpha_{\ell,S}, |\alpha_{\ell,H} - \alpha_{\ell,S}|/3)$ (Paria et al., 2019). We then compute the SHF utility ratio, right term in Equation 3, for the set of points C_t after each iteration of greedy, random, or MoSH-Dense. The denominator is calculated using λ^* with a full set of points, D , computed offline.

5.4 Step 1 Experiments: Synthetic and Real-World Applications

We provide additional experimental details and results on a synthetic problem, Branin-Currin, an engineering design problem, Four Bar Truss, and deep learning model selection in Appendix E.

5.4.1 Large Language Model (LLM) Personalization

We aim to personalize an LLM to generate outputs that are both concise and informative — two competing objectives. We employ proxy tuning (Liu et al., 2024; Mitchell et al., 2023; Shi et al., 2024), steering a base

model M using expert (M^+) and anti-expert (M^-) models for each objective. Following Liu et al. (2024), the proxy-tuned output distribution $p_{\tilde{M}}(X_t|x_{<t})$ is a function of the base model logits and weighted differences between expert/anti-expert logits for conciseness (M_0^+, M_0^-) and informativeness (M_1^+, M_1^-). The weights θ_0, θ_1 are our decision variables, controlling the trade-off. MoSH is used to find optimal θ_i values based on soft-hard bounds defined for conciseness and informativeness. As shown in Figure 3, MoSH performs consistently well across metrics, demonstrating its general applicability (more details in Appendix E.10).

5.4.2 Real Clinical Case: Cervical Cancer Brachytherapy Treatment Planning

We evaluate our method on treatment planning for a real cervical cancer brachytherapy clinical case. This problem consists of four objectives and three continuous decision variables. The objectives are (1) maximize the radiation dosage level to the cancer tumor, and minimize the radiation dosage levels to the (2) bladder, (3) rectum, and (4) bowel. We converted objectives (2)-(4) into maximization objectives. The decision variables are used as inputs to a linear program formulated as an epsilon-constraint method (Deufel et al., 2020). Figure 3 illustrates the plots with the soft-hard performance metrics. We notice that our proposed method surpasses the baselines by a greater amount in this many-objective setting, notably the soft-hard hypervolume and soft-hard positive ratio metrics.

5.5 Step 2 and End-to-End Experiments

To evaluate step (2), we evaluate our baselines on the sparsification of the dense set of points from MoSH-Dense. Figure 4 displays the SHF utility ratio values for successive points that the DM views, across two applications. We observe that in all five experimental settings, MoSH-Sparse matches or exceeds the baselines in achieving the overall highest SHF utility ratio. In the brachytherapy setting, MoSH-Sparse achieves the highest SHF utility ratio at the fastest pace, showcasing the effectiveness of our approach in distilling the dense set of points into a useful smaller set.

We further holistically evaluate our entire two-step process by comparing against all baselines for step (1), using MoSH-Sparse (Figure 4). We show that our method consistently leads in providing DM utility in most experiments (more in Appendix E.13).

6 RELATED WORKS

Populating the Pareto Frontier. Most existing MOO works aim to approximate the entire PF, using

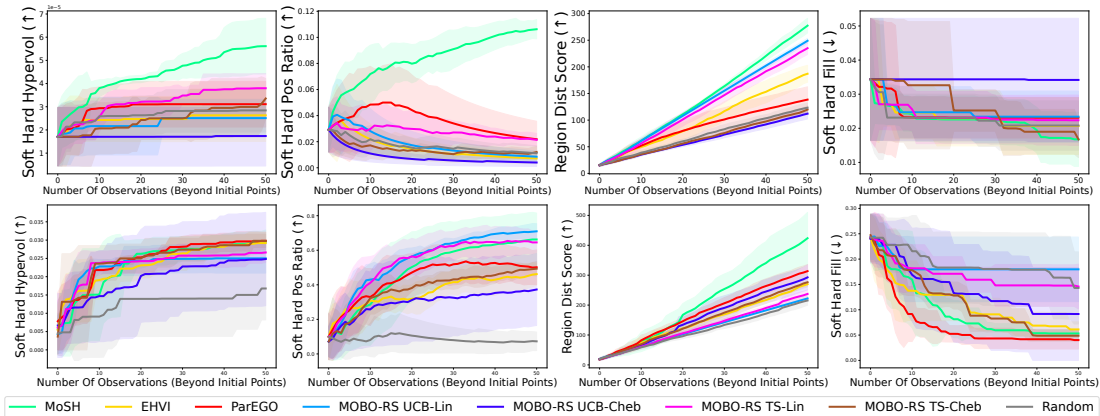


Figure 3: Top row: Brachytherapy treatment planning. Bottom row: LLM personalization. Plots show the metrics defined in Section 5.3, illustrating MoSH’s consistent lead in providing a dense set of points which is (1) diverse, (2) high-coverage, and (3) modeled after the SHFs. Mean \pm std. were computed over 6 independent runs.

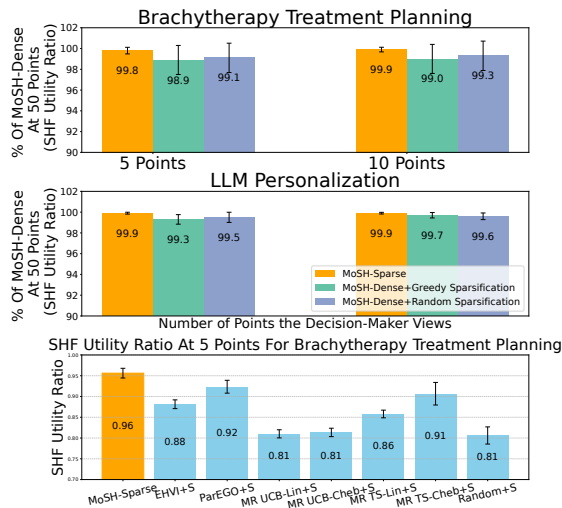


Figure 4: Evaluation for step (2) of MoSH, illustrating the SHF utility ratio (and std. dev) for the various methods relative to MoSH-Dense’s at 50 points, compared when the DM views 5 and 10 points. (Bottom) End-to-end evaluation, using the SHF utility ratio. Only the dense set, from step (1), changes for each bar. MR = MOBO-RS. S = SATURATE.

heuristics such as the maximum hypervolume improvement (Campigotto et al., 2014; Ponweiser et al., 2008; Emmerich, 2008; Picheny, 2015; Hernández-Lobato et al., 2016; Zhang et al., 2009). Others leverage RS of the objective values to attempt to recover the PF (Knowles, 2006; Paria et al., 2019; Zhang and Li, 2007; Zhang et al., 2010). Additional works also place a greater emphasis on sparse, diversified PF coverage (Zuluaga et al., 2016; Konakovik Lukovic et al., 2020) – more recently, in hard-constrained regions (Malkomes

et al., 2021). In contrast, we employ RS in a novel setting which aims to obtain a diverse, compact subset of the PF according to flexible user-defined priors (MFs).

MOO Feedback Mechanisms. Many feedback mechanisms, such as pairwise feedback, have been proposed – but not all are designed for MOO (Zintgraf et al., 2018; Roijers et al., 2018; Astudillo and Frazier, 2020; Ip et al., 2025). Hakanen and Knowles (2017) enables the DM to specify per-objective numerical ranges for MOO. Some allow for DM to order objectives by importance (Abdolshah et al., 2019). Ozaki et al. (2023) introduced MOO improvement request feedback type. These works enable fine-grained MOO control, but our work uniquely accounts for multiple levels of preferences without needing to specify exact numerical values, which is often psychologically more difficult (Qian et al., 2015).

Utility Function Specification. Rather than embed the utilities directly into the optimization process, works such as Huber et al. (2025) and Tomczyk and Kadziński (2020) propose to estimate monotonic utilities with pairwise preferences. Such assumptions of monotonic utilities are also often implicit, rather than explicit (Karasakal and Köksalan, 2013; Sobrie et al., 2018; Zintgraf et al., 2018), decreasing their flexibility. The formulation of objective utilities as inequality constraints, where the DMs aim to find inputs which satisfy thresholds on the objectives, is related to the topic of level set estimation (Gotovos, 2013; Zanette et al., 2018; Malkomes et al., 2021). In contrast, our work is native to the MOO setting and may additionally incorporate a flexible set of constraints (e.g. SHFs), which directly leverage the DM’s priors. *Several works related to MOBO may be found in Appendix G.*

7 CONCLUSION

We introduced a principled, two-step framework for MOO guided by general *monotonic utility functions* (MFs) in expensive black-box settings. Our approach returns a small, robust set of high-utility PO points by first performing dense PF sampling with Bayesian optimization, followed by a theoretically-grounded sparsification step using robust submodular optimization. As an instantiation of MFs, we introduced *soft-hard functions* (SHFs), an intuitive class of utility functions that models the common expert heuristic of imposing soft and hard bounds. We show that our framework has broad applicability to areas in brachytherapy, engineering design, and LLM personalization. For cervical cancer brachytherapy, our approach provides a compact set of treatment plans offering over 3% greater utility than the next best method. Across all domains, our method consistently allows the DM to capture over 99% of the utility from the dense set by validating just 5 points. We describe limitations in Appendix H.

References

- Majid Abdolshah, Alistair Shilton, Santu Rana, Sunil Gupta, and Svetha Venkatesh. Multi-objective Bayesian optimisation with preferences over objectives. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://papers.nips.cc/paper_files/paper/2019/hash/a7b7e4b27722574c611fe91476a50238-Abstract.html.
- Raul Astudillo and Peter Frazier. Multi-attribute Bayesian optimization with interactive preference learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pages 4496–4507. PMLR, June 2020. URL <https://proceedings.mlr.press/v108/astudillo20a.html>. ISSN: 2640-3498.
- Peter Auer, Chao-Kai Chiang, Ronald Ortner, and Madalina Drugan. Pareto Front Identification from Stochastic Bandit Feedback. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 939–947. PMLR, May 2016. URL <https://proceedings.mlr.press/v51/auer16.html>. ISSN: 1938-7228.
- Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020. URL <http://arxiv.org/abs/1910.06403>.
- Gilbert W. Bassett. The St. Petersburg Paradox and Bounded Utility. *History of Political Economy*, 19(4): 517–523, November 1987. ISSN 0018-2702, 1527-1919. doi: 10.1215/00182702-19-4-517.
- Julian Blank and Kalyanmoy Deb. Pymoo: Multi-Objective Optimization in Python. *IEEE Access*, 8: 89497–89509, 2020. ISSN 2169-3536. doi: 10.1109/ACCESS.2020.2990567. URL <https://ieeexplore.ieee.org/document/9078759/>.
- Brent Bryan and Jeff Schneider. Actively learning level-sets of composite functions. In *Proceedings of the 25th international conference on Machine learning - ICML '08*, pages 80–87, Helsinki, Finland, 2008. ACM Press. ISBN 978-1-60558-205-4. doi: 10.1145/1390156.1390167. URL <http://portal.acm.org/citation.cfm?doid=1390156.1390167>.
- Brent Bryan, Robert C. Nichol, Christopher R. Genovese, Jeff Schneider, Christopher J. Miller, and Larry Wasserman. Active Learning For Identifying Function Threshold Boundaries. In *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2005. URL <https://proceedings.neurips.cc/paper/2005/hash/8e930496927757aac0dbd2438cb3f4f6-Abstract.html>.
- Cédric Bélanger, Songye Cui, Yunzhi Ma, Philippe Després, J Adam M Cunha, and Luc Beaulieu. A GPU-based multi-criteria optimization algorithm for HDR brachytherapy. *Physics in Medicine & Biology*, 64(10):105005, May 2019. ISSN 1361-6560. doi: 10.1088/1361-6560/ab1817. URL <https://iopscience.iop.org/article/10.1088/1361-6560/ab1817>.
- Francesca Cairoli, Gianfranco Fenu, Felice Andrea Pellegrino, and Erica Salvato. Model Predictive Control of glucose concentration based on Signal Temporal Logic specifications. In *2019 6th International Conference on Control, Decision and Information Technologies (CoDIT)*, pages 714–719, April 2019. doi: 10.1109/CoDIT.2019.8820492. URL <https://ieeexplore.ieee.org/abstract/document/8820492>. ISSN: 2576-3555.
- Paolo Campigotto, Andrea Passerini, and Roberto Battiti. Active Learning of Pareto Fronts. *IEEE Transactions on Neural Networks and Learning Systems*, 25(3):506–519, March 2014. ISSN 2162-2388. doi: 10.1109/TNNLS.2013.2275918. URL <https://ieeexplore.ieee.org/document/6606803>. Conference Name: IEEE Transactions on Neural Networks and Learning Systems.
- Guillermo D. Canas and Lorenzo Rosasco. Learning Probability Measures with respect to Optimal

Transport Metrics, September 2012. URL <http://arxiv.org/abs/1209.1077>. arXiv:1209.1077.

Lisha Chen, A. F. M. Saif, Yanning Shen, and Tianyi Chen. FERERO: A Flexible Framework for Preference-Guided Multi-Objective Learning, December 2024. URL <http://arxiv.org/abs/2412.01773>. arXiv:2412.01773 [cs].

F. Y. CHENG and X. S. LI. Generalized Center Method for Multiobjective Engineering Optimization. *Engineering Optimization*, 31(5):641–661, May 1999. ISSN 0305-215X. doi: 10.1080/03052159908941390. URL <https://doi.org/10.1080/03052159908941390>. Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/03052159908941390>.

Tinkle Chugh. Scalarizing Functions in Bayesian Multiobjective Optimization, April 2019. URL <http://arxiv.org/abs/1904.05760>. arXiv:1904.05760 [cs, stat].

Songye Cui, Philippe Després, and Luc Beaulieu. A multi-criteria optimization approach for HDR prostate brachytherapy: II. Benchmark against clinical plans. *Physics in Medicine & Biology*, 63(20):205005, October 2018. ISSN 1361-6560. doi: 10.1088/1361-6560/aae24f. URL <https://iopscience.iop.org/article/10.1088/1361-6560/aae24f>.

John M. Danskin. The Theory of Max-Min, with Applications. *SIAM Journal on Applied Mathematics*, 14(4):641–664, July 1966. ISSN 0036-1399. doi: 10.1137/0114053. URL <https://epubs.siam.org/doi/abs/10.1137/0114053>. Publisher: Society for Industrial and Applied Mathematics.

Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Multi-Objective Bayesian Optimization over High-Dimensional Search Spaces, June 2022. URL <http://arxiv.org/abs/2109.10964>. arXiv:2109.10964 [cs].

Kalyanmoy Deb and J. Sundar. Reference point based multi-objective optimization using evolutionary algorithms. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, GECCO '06, pages 635–642, New York, NY, USA, July 2006. Association for Computing Machinery. ISBN 978-1-59593-186-3. doi: 10.1145/1143997.1144112. URL <https://dl.acm.org/doi/10.1145/1143997.1144112>.

Li Deng. The MNIST Database of Handwritten Digit Images for Machine Learning Research [Best of the Web]. *IEEE Signal Processing Magazine*, 29(6):141–142, November 2012. ISSN 1558-0792. doi: 10.1109/MSP.2012.2211477. URL <https://ieeexplore.ieee.org/document/6296535>.

Christopher L. Deufel, Marina A. Epelman, Kalyan S. Pasupathy, Mustafa Y. Sir, Victor W. Wu, and Michael G. Herman. PNaV: A tool for generating a high-dose-rate brachytherapy treatment plan by navigating the Pareto surface guided by the visualization of multidimensional trade-offs. *Brachytherapy*, 19(4):518–531, July 2020. ISSN 15384721. doi: 10.1016/j.brachy.2020.02.013. URL <https://linkinghub.elsevier.com/retrieve/pii/S1538472120300325>.

Kristin Diehl. When Two Rights Make a Wrong: Searching Too Much in Ordered Environments. *Journal of Marketing Research*, 42(3):313–322, August 2005. ISSN 0022-2437. doi: 10.1509/jmkr.2005.42.3.313. URL <https://doi.org/10.1509/jmkr.2005.42.3.313>. Publisher: SAGE Publications Inc.

Karim El Moutaouakil, Abdellatif El Ouissari, Vasile Palade, Anas Charroud, Adrian Olaru, Hicham Baïzri, Saliha Chellak, and Mouna Cheggour. Multi-Objective Optimization for Controlling the Dynamics of the Diabetic Population. *Mathematics*, 11(13):2957, January 2023. ISSN 2227-7390. doi: 10.3390/math11132957. URL <https://www.mdpi.com/2227-7390/11/13/2957>. Publisher: Multidisciplinary Digital Publishing Institute.

Michael Emmerich. The computation of the expected improvement in dominated hypervolume of Pareto front approximations. January 2008.

Peppino Fazio and Mauro Tropea. Advanced Resources Reservation in Mobile Cellular Networks: Static vs. Dynamic Approaches under Vehicular Mobility Model. *Telecom*, 2(4):302–327, September 2021. ISSN 2673-4001. doi: 10.3390/telecom2040020. URL <https://www.mdpi.com/2673-4001/2/4/20>. Company: Multidisciplinary Digital Publishing Institute Distributor: Multidisciplinary Digital Publishing Institute Institution: Multidisciplinary Digital Publishing Institute Label: Multidisciplinary Digital Publishing Institute Publisher: publisher.

Jenna C. Fromer and Connor W. Coley. Computer-aided multi-objective optimization in small molecule discovery. *Patterns*, 4(2):100678, February 2023. ISSN 26663899. doi: 10.1016/j.patter.2023.100678. URL <https://linkinghub.elsevier.com/retrieve/pii/S2666389923000016>.

Alkis Gotovos. Active Learning for Level Set Estimation. 2013. doi: 10.3929/ETHZ-A-009767767. URL <http://hdl.handle.net/20.500.11850/153931>. Medium: application/pdf, Online-Resource Publisher: ETH Zurich.

Benyamin Grosman, Eyal Dassau, Howard C. Zisser, Lois Jovanović, and Francis J. Doyle. Zone

- Model Predictive Control: A Strategy to Minimize Hyper- and Hypoglycemic Events. *Journal of Diabetes Science and Technology*, 4(4):961–975, July 2010. ISSN 1932-2968. doi: 10.1177/193229681000400428. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2909531/>.
- Sajad Haghaniifar, Michael McCourt, Bolong Cheng, Jeffrey Wuenschell, Paul Ohodnicki, and Paul W. Leu. Discovering high-performance broadband and broad angle antireflection surfaces by machine learning. *Optica*, 7(7):784–789, July 2020. ISSN 2334-2536. doi: 10.1364/OPTICA.387938. URL <https://opg.optica.org/optica/abstract.cfm?uri=optica-7-7-784>. Publisher: Optica Publishing Group.
- Jussi Hakanen and Joshua D. Knowles. On Using Decision Maker Preferences with ParEGO. In Heike Trautmann, Günter Rudolph, Kathrin Klammroth, Oliver Schütze, Margaret Wiecek, Yaochu Jin, and Christian Grimme, editors, *Evolutionary Multi-Criterion Optimization*, volume 10173, pages 282–297. Springer International Publishing, Cham, 2017. ISBN 978-3-319-54156-3 978-3-319-54157-0. doi: 10.1007/978-3-319-54157-0_20. URL http://link.springer.com/10.1007/978-3-319-54157-0_20. Series Title: Lecture Notes in Computer Science.
- Daniel Hernández-Lobato, José Miguel Hernández-Lobato, Amar Shah, and Ryan P. Adams. Predictive Entropy Search for Multi-objective Bayesian Optimization, February 2016. URL <http://arxiv.org/abs/1511.05467>. arXiv:1511.05467 [stat].
- Felix Huber, Sebastian Rojas Gonzalez, and Raul Astudillo. Bayesian preference elicitation for decision support in multiobjective optimization, July 2025. URL <http://arxiv.org/abs/2507.16999>. arXiv:2507.16999 [stat].
- Joshua Hang Sai Ip, Ankush Chakrabarty, Ali Mesbah, and Diego Romeres. User Preference Meets Pareto-Optimality in Multi-Objective Bayesian Optimization, March 2025. URL <http://arxiv.org/abs/2502.06971>. arXiv:2502.06971 [cs].
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. Camels in a Changing Climate: Enhancing LM Adaptation with Tulu 2, November 2023. URL <http://arxiv.org/abs/2311.10702>. arXiv:2311.10702 [cs].
- Shogo Iwazaki, Yu Inatsu, and Ichiro Takeuchi. Bayesian Experimental Design for Finding Reliable Level Set Under Input Uncertainty. *IEEE Access*, 8:203982–203993, 2020. ISSN 2169-3536. doi: 10.1109/ACCESS.2020.3036863. URL <https://ieeexplore.ieee.org/document/9252941/?arnumber=9252941>. Conference Name: IEEE Access.
- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. Personalized Soups: Personalized Large Language Model Alignment via Post-hoc Parameter Merging, October 2023. URL <http://arxiv.org/abs/2310.11564>. arXiv:2310.11564 [cs].
- Fuhao Ji, Auralee Edelen, Ryan Roussel, Xiaozhe Shen, Sara Miskovich, Stephen Weathersby, Duan Luo, Mianzhen Mo, Patrick Kramer, Christopher Mayes, Mohamed A. K. Othman, Emilio Nanni, Xijie Wang, Alexander Reid, Michael Minitti, and Robert Joel England. Multi-objective Bayesian active learning for MeV-ultrafast electron diffraction. *Nature Communications*, 15(1):4726, June 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-48923-9. URL <https://www.nature.com/articles/s41467-024-48923-9>. Publisher: Nature Publishing Group.
- Esra Karasakal and Murat Köksalan. An interactive partitioning approach for multiobjective decision making under a general monotone utility function. *International Journal of Information Technology & Decision Making*, 12(05):969–997, September 2013. ISSN 0219-6220. doi: 10.1142/S0219622013400051. URL <https://www.worldscientific.com/doi/10.1142/S0219622013400051>. Publisher: World Scientific Publishing Co.
- J. Knowles. ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, February 2006. ISSN 1941-0026. doi: 10.1109/TEVC.2005.851274. URL <https://ieeexplore.ieee.org/document/1583627>. Conference Name: IEEE Transactions on Evolutionary Computation.
- Mina Konakovic Lukovic, Yunsheng Tian, and Wojciech Matusik. Diversity-Guided Multi-Objective Bayesian Optimization With Batch Evaluations. In *Advances in Neural Information Processing Systems*, volume 33, pages 17708–17720. Curran Associates, Inc., 2020. URL <https://papers.nips.cc/paper/2020/hash/cd3109c63bf4323e6b987a5923becb96-Abstract.html>.
- Andreas Krause, H. Brendan McMahan, Carlos Guestrin, and Anupam Gupta. Robust Submod-

- ular Observation Selection. *Journal of Machine Learning Research*, 9(93):2761–2801, 2008. ISSN 1533-7928. URL <http://jmlr.org/papers/v9/krause08b.html>.
- Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia Tsvetkov, Yejin Choi, and Noah A. Smith. Tuning Language Models by Proxy, August 2024. URL <http://arxiv.org/abs/2401.08565>. arXiv:2401.08565 [cs].
- Sohvi Luukkonen, Helle W. van den Maagdenberg, Michael T. M. Emmerich, and Gerard J. P. van Westen. Artificial intelligence in multi-objective drug design. *Current Opinion in Structural Biology*, 79: 102537, April 2023. ISSN 0959-440X. doi: 10.1016/j.sbi.2023.102537. URL <https://www.sciencedirect.com/science/article/pii/S0959440X23000118>.
- Gustavo Malkomes, Bolong Cheng, Eric H. Lee, and Mike Mccourt. Beyond the Pareto Efficient Frontier: Constraint Active Search for Multiobjective Experimental Design. In *Proceedings of the 38th International Conference on Machine Learning*, pages 7423–7434. PMLR, July 2021. URL <https://proceedings.mlr.press/v139/malkomes21a.html>. ISSN: 2640-3498.
- Sharmistha Mandal and Ashoke Sutradhar. Robust multi-objective blood glucose control in Type-1 diabetic patient. *IET Systems Biology*, 13(3):136–146, June 2019. ISSN 1751-8849. doi: 10.1049/iet-syb.2018.5093. URL <https://pmc.ncbi.nlm.nih.gov/articles/PMC8687400/>.
- Eric Mitchell, Rafael Rafailov, Archit Sharma, Chelsea Finn, and Christopher D. Manning. An Emulator for Fine-Tuning Large Language Models using Small Language Models, October 2023. URL <http://arxiv.org/abs/2310.12962>. arXiv:2310.12962 [cs].
- Björn Morén, Torbjörn Larsson, and Åsa Carlsson Tedgren. Optimization in treatment planning of high dose-rate brachytherapy — Review and analysis of mathematical models. *Medical Physics*, 48(5): 2057–2082, 2021. ISSN 2473-4209. doi: 10.1002/mp.14762. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mp.14762>.
- John Weirstrass Muteba Mwamba, Leon Mishindo Mbucici, and Jules Clement Mba. Multi-Objective Portfolio Optimization: An Application of the Non-Dominated Sorting Genetic Algorithm III. *International Journal of Financial Studies*, 13(1):15, March 2025. ISSN 2227-7072. doi: 10.3390/ijfs13010015. URL <https://www.mdpi.com/2227-7072/13/1/15>. Publisher: Multidisciplinary Digital Publishing Institute.
- G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14(1):265–294, December 1978. ISSN 0025-5610, 1436-4646. doi: 10.1007/BF01588971. URL <http://link.springer.com/10.1007/BF01588971>.
- Ryota Ozaki, Kazuki Ishikawa, Youhei Kanzaki, Shinya Suzuki, Shion Takeno, Ichiro Takeuchi, and Masayuki Karasuyama. Multi-Objective Bayesian Optimization with Active Preference Learning, November 2023. URL <http://arxiv.org/abs/2311.13460>. arXiv:2311.13460 [cs].
- Christos H. Papadimitriou and Mihalis Yannakakis. Multiobjective query optimization. In *Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, PODS '01, pages 52–59, New York, NY, USA, May 2001. Association for Computing Machinery. ISBN 978-1-58113-361-5. doi: 10.1145/375551.375560. URL <https://dl.acm.org/doi/10.1145/375551.375560>.
- Biswajit Paria, Kirthevasan Kandasamy, and Barnabás Póczos. A Flexible Framework for Multi-Objective Bayesian Optimization using Random Scalarizations, June 2019. URL <http://arxiv.org/abs/1805.12168>. arXiv:1805.12168 [cs, stat].
- Victor Picheny. Multiobjective optimization using Gaussian process emulators via stepwise uncertainty reduction. *Statistics and Computing*, 25(6): 1265–1280, November 2015. ISSN 1573-1375. doi: 10.1007/s11222-014-9477-x. URL <https://doi.org/10.1007/s11222-014-9477-x>.
- Wolfgang Ponweiser, Tobias Wagner, Dirk Biermann, and Markus Vincze. Multiobjective Optimization on a Limited Budget of Evaluations Using Model-Assisted S -Metric Selection. In Günter Rudolph, Thomas Jansen, Nicola Beume, Simon Lucas, and Carlo Poloni, editors, *Parallel Problem Solving from Nature – PPSN X*, pages 784–794, Berlin, Heidelberg, 2008. Springer. ISBN 978-3-540-87700-4. doi: 10.1007/978-3-540-87700-4_78.
- Li Qian, Jinyang Gao, and H. V. Jagadish. Learning user preferences by adaptive pairwise comparison. *Proc. VLDB Endow.*, 8(11):1322–1333, July 2015. ISSN 2150-8097. doi: 10.14778/2809974.2809992. URL <https://dl.acm.org/doi/10.14778/2809974.2809992>.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct Preference Optimization: Your Language Model is Secretly a Reward Model, July 2024. URL <http://arxiv.org/abs/2305.18290>. arXiv:2305.18290 [cs].

D. M. Roijers, P. Vamplew, S. Whiteson, and R. Dazeley. A Survey of Multi-Objective Sequential Decision-Making. *Journal of Artificial Intelligence Research*, 48:67–113, October 2013. ISSN 1076-9757. doi: 10.1613/jair.3987. URL <https://www.jair.org/index.php/jair/article/view/10836>.

Diederik M Roijers, Luisa M Zintgraf, Pieter Libin, and Ann Nowé. Interactive Multi-Objective Reinforcement Learning in Multi-Armed Bandits for Any Utility Function. 2018.

Daniel Russo and Benjamin Van Roy. Learning to Optimize via Posterior Sampling. *Mathematics of Operations Research*, 39(4):1221–1243, November 2014. ISSN 0364-765X. doi: 10.1287/moor.2014.0650. URL <https://pubsonline.informs.org/doi/abs/10.1287/moor.2014.0650>. Publisher: INFORMS.

Roxana Rădulescu, Patrick Mannion, Diederik M. Roijers, and Ann Nowé. Multi-objective multi-agent decision making: a utility-based analysis and survey. *Autonomous Agents and Multi-Agent Systems*, 34(1):10, December 2019. ISSN 1573-7454. doi: 10.1007/s10458-019-09433-x. URL <https://doi.org/10.1007/s10458-019-09433-x>.

Pedram Samadi, Amir-Hamed Mohsenian-Rad, Robert Schober, Vincent W. S. Wong, and Juri Jatskevich. Optimal Real-Time Pricing Algorithm Based on Utility Maximization for Smart Grid. In *2010 First IEEE International Conference on Smart Grid Communications*, pages 415–420, October 2010. doi: 10.1109/SMARTGRID.2010.5622077. URL <https://ieeexplore.ieee.org/abstract/document/5622077>.

Ahmed A. Shaier, Mahmoud M. Elymany, Mohamed A. Enany, and Nadia A. Elsonbaty. Multi-objective optimization and algorithmic evaluation for EMS in a HRES integrating PV, wind, and backup storage. *Scientific Reports*, 15(1):1147, January 2025. ISSN 2045-2322. doi: 10.1038/s41598-024-84227-0. URL <https://www.nature.com/articles/s41598-024-84227-0>. Publisher: Nature Publishing Group.

Ruizhe Shi, Yifang Chen, Yushi Hu, Alisa Liu, Hananeh Hajishirzi, Noah A. Smith, and Simon Du. Decoding-Time Language Model Alignment with Multiple Objectives, June 2024. URL <http://arxiv.org/abs/2406.18853>. arXiv:2406.18853 [cs].

Thiago F. Soares, Maurício C. Escarpinati, and Paulo H. R. Gabriel. MULTI-OBJECTIVE OPTIMIZATION APPLIED TO DIET PLANNING FOR PEOPLE WITH DIABETES. *Pesquisa Operacional*, 44: e285156, 2024. ISSN 0101-7438, 1678-5142. doi: <https://doi.org/10.1590/0101-7438.2023.043.00285156>.

URL <https://www.scielo.br/j/pope/a/PSbGK4NMqmp96s6zdYrm4xF/?format=html&lang=en>. Publisher: Sociedade Brasileira de Pesquisa Operacional.

Olivier Sobrie, Nicolas Gillis, Vincent Mousseau, and Marc Pirlot. UTA-poly and UTA-splines: additive value functions with polynomial marginals. *European Journal of Operational Research*, 264(2):405–418, January 2018. ISSN 03772217. doi: 10.1016/j.ejor.2017.03.021. URL <http://arxiv.org/abs/1603.02626>. arXiv:1603.02626 [math].

Shinya Suzuki, Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Multi-objective Bayesian Optimization using Pareto-frontier Entropy. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9279–9288. PMLR, November 2020. URL <https://proceedings.mlr.press/v119/suzuki20a.html>. ISSN: 2640-3498.

Ryoji Tanabe and Hisao Ishibuchi. An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89:106078, April 2020. ISSN 1568-4946. doi: 10.1016/j.asoc.2020.106078. URL <https://www.sciencedirect.com/science/article/pii/S1568494620300181>.

Michał K. Tomczyk and Miłosz Kadziński. Decomposition-Based Interactive Evolutionary Algorithm for Multiple Objective Optimization. *IEEE Transactions on Evolutionary Computation*, 24(2):320–334, April 2020. ISSN 1941-0026. doi: 10.1109/TEVC.2019.2915767. URL <https://ieeexplore.ieee.org/document/8710313>.

Marjolein C. van der Meer, Peter A.N. Bosman, Yury Niatsetski, Tanja Alderliesten, Niek van Wieringen, Bradley R. Pieters, and Arjan Bel. Bi-objective optimization of catheter positions for high-dose-rate prostate brachytherapy. *Medical Physics*, 47(12):6077–6086, 2020. ISSN 2473-4209. doi: 10.1002/mp.14505. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mp.14505>.

Yash Vesikar, Kalyanmoy Deb, and Julian Blank. Reference Point Based NSGA-III for Preferred Solutions. In *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1587–1594, Bangalore, India, November 2018. IEEE. ISBN 978-1-5386-9276-9. doi: 10.1109/SSCI.2018.8628819. URL <https://ieeexplore.ieee.org/document/8628819/>.

Akila N. Viswanathan, Sushil Beriwal, Jennifer F. De Los Santos, D. Jeffrey Demanes, David Gaffney, Jorgen Hansen, Ellen Jones, Christian Kirisits, Bruce

Thomadsen, and Beth Erickson. American Brachytherapy Society consensus guidelines for locally advanced carcinoma of the cervix. Part II: High-dose-rate brachytherapy. *Brachytherapy*, 11(1):47–52, January 2012. ISSN 1538-4721. doi: 10.1016/j.brachy.2011.07.002. URL <https://www.sciencedirect.com/science/article/pii/S1538472111003515>.

Lyndon While, Lucas Bradstreet, and Luigi Barone. A Fast Way of Calculating Exact Hypervolumes. *IEEE Transactions on Evolutionary Computation*, 16(1):86–95, February 2012. ISSN 1941-0026. doi: 10.1109/TEVC.2010.2077298. Conference Name: IEEE Transactions on Evolutionary Computation.

Andrzej P. Wierzbicki. A mathematical basis for satisficing decision making. *Mathematical Modelling*, 3(5):391–405, January 1982. ISSN 0270-0255. doi: 10.1016/0270-0255(82)90038-0. URL <https://www.sciencedirect.com/science/article/pii/0270025582900380>.

Christopher Williams and Carl Rasmussen. Gaussian Processes for Regression. In *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1995. URL https://proceedings.neurips.cc/paper_files/paper/1995/hash/7cce53cf90577442771720a370c3c723-Abstract.html.

Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. MARS: Markov Molecular Sampling for Multi-objective Drug Discovery, March 2021. URL <http://arxiv.org/abs/2103.10432>. arXiv:2103.10432 [cs, q-bio].

Yan Yu, J. B Zhang, Gang Cheng, M. C Schell, and Paul Okunieff. Multi-objective optimization in radiotherapy: applications to stereotactic radiosurgery and prostate brachytherapy. *Artificial Intelligence in Medicine*, 19(1):39–51, May 2000. ISSN 0933-3657. doi: 10.1016/S0933-3657(99)00049-4. URL <https://www.sciencedirect.com/science/article/pii/S0933365799000494>.

Andrea Zanette, Junzi Zhang, and Mykel J. Kochenderfer. Robust Super-Level Set Estimation using Gaussian Processes, November 2018. URL <http://arxiv.org/abs/1811.09977>. arXiv:1811.09977 [cs, stat].

Chao Zhang, Samson Lasaulce, Li Wang, Lucas Saludjian, and H. Vincent Poor. A refined consumer behavior model for energy systems: Application to the pricing and energy-efficiency problems. *Applied Energy*, 308:118239, February 2022. ISSN 0306-2619. doi: 10.1016/j.apenergy.2021.118239. URL <https://www.sciencedirect.com/science/article/pii/S0306261921015026>.

Qingfu Zhang and Hui Li. MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, December 2007. ISSN 1941-0026. doi: 10.1109/TEVC.2007.892759. URL <https://ieeexplore.ieee.org/document/4358754>. Conference Name: IEEE Transactions on Evolutionary Computation.

Qingfu Zhang, Aimin Zhou, Shizheng Zhao, Ponnuthurai Nagaratnam Suganthan, Wudong Liu, and Santosh Tiwari. Multiobjective optimization Test Instances for the CEC 2009 Special Session and Competition. 2009.

Qingfu Zhang, Wudong Liu, Edward Tsang, and Botond Virginas. Expensive Multiobjective Optimization by MOEA/D With Gaussian Process Model. *IEEE Transactions on Evolutionary Computation*, 14(3):456–474, June 2010. ISSN 1941-0026. doi: 10.1109/TEVC.2009.2033671. URL <https://ieeexplore.ieee.org/document/5353656>. Conference Name: IEEE Transactions on Evolutionary Computation.

Luisa M. Zintgraf, Diederik M. Roijers, Sjoerd Linders, Catholijn M. Jonker, and Ann Nowé. Ordered Preference Elicitation Strategies for Supporting Multi-Objective Decision Making, February 2018. URL <http://arxiv.org/abs/1802.07606>. arXiv:1802.07606 [cs].

Marcela Zuluaga, Andreas Krause, and Markus Püschel. e-PAL: An Active Learning Approach to the Multi-Objective Optimization Problem. *Journal of Machine Learning Research*, 17(104):1–32, 2016. ISSN 1533-7928. URL <http://jmlr.org/papers/v17/15-047.html>.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. Yes. Comprehensive descriptions of the setting, assumptions, and algorithms may be found in Sections 2, 3, and 4. Additional details are presented in Appendix B and C.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. Yes. For Step 1, Section 3 and Appendices B and E.5 provide extensive analyses of the properties and motivation. For Step 2, Section 4 and Appendix C provide details regarding the properties and complexity of the algorithm.

- (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. Yes, provided as part of the supplemental material.
2. For any theoretical claim, check if you include:
- (a) Statements of the full set of assumptions of all theoretical results. Yes. Appendix D explicitly mentions the assumptions made for the theoretical results.
 - (b) Complete proofs of all theoretical results. Yes. All proofs are provided in Appendix D.
 - (c) Clear explanations of any assumptions. Yes. Appendix D provides clear explanations of any assumptions.
3. For all figures and tables that present empirical results, check if you include:
- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). Yes, provided as part of the supplemental material.
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). Yes. Extensive training details for all of the experiments are provided in Appendix E. Additional details regarding number of trials and performance evaluation are listed in Appendix F.
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). Yes. Such information is presented in each of the figures throughout this manuscript. Additional definitions are provided in Appendix F.1.
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). Yes, these are provided in Appendix F.2.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. Yes. The models used for the LLM personalization experiment are properly cited in Appendix E.10. Usage of the benchmarks and datasets for the other experiments are all properly cited in Appendix E.
 - (b) The license information of the assets, if applicable. Not Applicable
 - (c) New assets either in the supplemental material or as a URL, if applicable. Not Applicable
- (d) Information about consent from data providers/curators. Yes. For the brachytherapy treatment planning data, consent is explicitly described in Appendix E.8.
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. Not Applicable
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. Not Applicable
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. Not Applicable
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. Not Applicable

Supplementary Material

A MODELING MULTI-OBJECTIVE TRADEOFFS WITH MONOTONIC UTILITY FUNCTIONS

We articulate the differences between MoSH and prior methods in Table 1 below.

Table 1: Descriptive Table Of Differences Comparing MoSH To Existing Methods.

Method	Prior Specification	Cognitive Load for DM	Theoretical Guarantees?
MoSH (Ours)	Explicit (via MFs)	Low (High utility within just 5 points)	Yes (Coverage & Optimality)
Pairwise Feedback (Qian et al., 2015) (Astudillo and Frazier, 2020) (Ozaki et al., 2023)	Indirect (Iterative)	High (Requires many queries)	No
Hypervolume-Based (Ponweiser et al., 2008)	No	Very High (Returns large sets)	Yes
Random Scalarizations (Paria et al., 2019)	Yes (Not intuitive for DM)	High (Large sets, unintuitive prior)	Yes

B STEP 1: DENSE PARETO FRONTIER SAMPLING WITH BAYESIAN OPTIMIZATION ADDITIONAL DETAILS

Here we describe the acquisition function used in line 7 of step 1, Algorithm 1. For our experiments, we use the Upper Confidence Bound (UCB) heuristic. We define $\text{acq}(u, \lambda_t, x) = s_{\lambda_t}(u_{\varphi(x)})$ where $\varphi(x) = \mu_t(x) + \sqrt{\beta_t}\sigma_t(x)$ and $\beta_t = \sqrt{0.125 \times \log(2 \times t + 1)}$. For β_t , we followed the optimal suggestion in Paria et al. (2019). We emphasize that *our algorithm is agnostic to the scalarization and acquisition functions*.

C STEP 2: PARETO FRONTIER SPARSIFICATION ADDITIONAL DETAILS

Theorem 3. (Nemhauser et al., 1978) *In the case of any normalized, nonotonic submodular function F , the set A_G obtained by the greedy algorithm achieves at least a constant fraction $(1 - \frac{1}{e})$ of the objective value obtained by the optimal solution, that is,*

$$F(A_G) \geq (1 - \frac{1}{e}) \max_{|A| \leq k} F(A)$$

Algorithm 3 describes the greedy submodular partial cover (GPC) algorithm in detail, used as a subroutine in Algorithm 2.

Algorithm 3 Greedy Submodular Partial Cover (GPC) Algorithm (Krause et al., 2008)

```

1: procedure GPC( $\bar{F}_q, q$ )
2:    $C \leftarrow \emptyset$ 
3:   while  $\bar{F}_q(C) < q$  do
4:     foreach  $c \in D \setminus C$  do  $\delta_c = \bar{F}_q(C \cup \{c\}) - \bar{F}_q(C)$ 
5:      $C \leftarrow C \cup \{\operatorname{argmax}_c \delta_c\}$ 
6:   end while
7:   return  $C$ 
8: end procedure

```

D REQUIRED PROOFS AND DEFINITIONS

Definition 4 (Pareto dominant). *A solution $x_1 \in X$ is Pareto dominated by another point $x_2 \in X$ if and only if $f_\ell(x_1) \leq f_\ell(x_2) \forall \ell \in [L]$ and $\exists \ell \in [L]$ s.t. $f_\ell(x_1) < f_\ell(x_2)$ (Paria et al., 2019).*

Definition 5 (Submodular). *F is submodular if and only if for all $A \subseteq B \subseteq V$ and $s \in V \setminus B$ it holds that $F(A \cup \{s\}) - F(A) \geq F(B \cup \{s\}) - F(B)$*

Theorem 6. *Consider finite sets Ω, Λ and function $f : \Omega \times \Lambda \rightarrow \mathbb{R}$. Fix $\lambda \in \Lambda$. Then, given $C \subset \Omega$, the set function*

$$F(C) := \frac{\max_{c \in C} f(c, \lambda)}{\max_{c \in \Omega} f(c, \lambda)}$$

is submodular.

Proof. Let $X \subseteq Y \subseteq \Omega$ and $x \in \Omega \setminus Y$ (note: it clearly follows that $x \notin X$). Consider $x^* := \max_{c \in \Omega} f(c, \lambda)$. We obtain the following mutually exclusive and exhaustive cases, in which we use the fact that for any sets A, B such that $A \subseteq B$, $\max_{a \in A} f(a) \leq \max_{b \in B} f(b)$ for any function f .

- $x^* \in X \subseteq Y$. By definition, $F(X \cup \{x\}) - F(X) = F(Y \cup \{x\}) - F(Y) = 0$.
- $x^* \in Y \setminus X$. By definition, $F(Y \cup \{x\}) - F(Y) = 0$. Clearly $X \subseteq X \cup \{x\}$, thus $F(X \cup \{x\}) - F(X) \geq 0 = F(Y \cup \{x\}) - F(Y)$.
- $x^* \in \Omega \setminus Y$ and $x^* \neq x$.
 - If $F(x, \lambda) \geq \max_{c \in Y} f(c, \lambda)$, then $F(X \cup \{x\}) = F(Y \cup \{x\}) = F(x, \lambda)$.
 - If $F(x, \lambda) \leq \max_{c \in X} f(c, \lambda)$, $F(X \cup \{x\}) = F(X)$ and $F(Y \cup \{x\}) = F(Y)$.
 - Finally, if $\max_{c \in X} f(c, \lambda) \leq F(x, \lambda) \leq \max_{c \in Y} f(c, \lambda)$, then $F(Y \cup \{x\}) = F(Y)$ and $F(X \cup \{x\}) \geq F(X)$.

Combining with the fact since $X \subseteq Y$, $F(X) \leq F(Y)$, in all of the above sub-cases, we indeed have $F(X \cup \{x\}) - F(X) \geq F(Y \cup \{x\}) - F(Y)$.

- $x^* = x$. As $F(X \cup \{x\}) = F(Y \cup \{x\})$, and $F(X) \leq F(Y)$, we automatically get $F(X \cup \{x\}) - F(X) \geq F(Y \cup \{x\}) - F(Y)$.

Therefore, we satisfy the definition of submodularity. □

Definition 7 (Instantaneous MF Regret). $r(x_t, \lambda_t) = 1 - \frac{s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x_t))}$

Definition 8 (Cumulative MF Regret). $R_C(T) = \sum_{t=1}^T r(x_t, \lambda_t)$

Definition 9 (Bayes MF Regret and Utility Ratio). $R_B(T) = \mathbb{E}_{\lambda \sim p(\lambda)} [1 - \frac{\max_{x \in D_T} s_{\lambda}(u_f(x))}{\max_{x \in X} s_{\lambda}(u_f(x))}]$ where $D_T = \{x_t\}_{t=1}^T$. Additionally, $U_B(T) = \mathbb{E}_{\lambda \sim p(\lambda)} [\frac{\max_{x \in D_T} s_{\lambda}(u_f(x))}{\max_{x \in X} s_{\lambda}(u_f(x))}]$, where $U_B(T)$ is the Bayes MF Utility Ratio after T iterations.

Definition 10 (Expected Bayes MF Regret and Expected Cumulative MF Regret). *Similar to (Paria et al., 2019), $\mathbb{E}R_B(T)$ is the expected Bayes MF Regret, with the expectation being taken over f , noise ϵ , and other sources of randomness. Likewise, $\mathbb{E}R_C(T)$ is the expected Cumulative MF Regret, with the expectation being taken over f , noise ϵ , and λ_t .*

Definition 11 (Maximum Information Gain). *We leverage this definition from (Paria et al., 2019). The maximum information gain after T observations measures the notion of information gained about random process f after observing some set of points A , and is defined as:*

$$\gamma_T = \max_{A \subset X: |A|=T} I(y_A; f) \quad (5)$$

Definition 12 (Lipschitz Condition). *We assume the following is M_λ -Lipschitz in the ℓ_1 -norm for all $\lambda \in \Lambda$,*

$$\left| \frac{s_\lambda(y_1)}{\max_{y' \in \text{Im}(u_f)} s_\lambda(y')} - \frac{s_\lambda(y_2)}{\max_{y' \in \text{Im}(u_f)} s_\lambda(y')} \right| \leq M_\lambda \|y_1 - y_2\|_1 \quad (6)$$

where $y \in \mathbb{R}^L$ corresponds to $u_f(x)$.

Definition 13 (Lipschitz Condition). *We assume the following is J -Lipschitz in λ for all $y \in \mathbb{R}^L$.*

$$\left| \frac{s_{\lambda_1}(y)}{\max_{y' \in \text{Im}(u_f)} s_{\lambda_1}(y')} - \frac{s_{\lambda_2}(y)}{\max_{y' \in \text{Im}(u_f)} s_{\lambda_2}(y')} \right| \leq J \|\lambda_1 - \lambda_2\|_1 \quad (7)$$

Definition 14 (Regret Bounds). *We follow similar notation as in (Paria et al., 2019). Assume that $\forall \ell \in [L], t \in [T], x \in \mathcal{X}$, each objective $f_\ell(x)$ follows a Gaussian distribution with marginal variances upper bounded by 1, and the observation noise $\epsilon_{t\ell} \sim \mathcal{N}(0, \sigma_\ell^2)$ is drawn independently of everything else. We assume upper bounds $M_\lambda \leq M$, $\sigma_\ell^2 \leq \sigma^2$, $\gamma_{T\ell} \leq \gamma_T$, where $\gamma_{T\ell}$ is the maximum information gain for the ℓ th objective. We assume $\mathcal{X} \subseteq [0, 1]^d$. Furthermore, let $x_t^* = \arg \max_{x \in X} s_{\lambda_t}(u_f(x))$. We denote by $U_t(\lambda, x) = s_\lambda(u_{\varphi(x)})$ where $\varphi(x) = \mu_t(x) + \sqrt{\beta_t} \sigma_t(x)$. Finally, the history until $T-1$ is denoted as \mathcal{H}_t , i.e. $\{(x_t, y_t, \lambda_t)\}_{t=1}^{T-1}$.*

Theorem 15. *The expected cumulative MF regret for dense sampling after T observations can be upper bounded for both Upper Confidence Bound and Thompson Sampling as,*

$$\mathbb{E}R_C(T) = O\left(M \left[\frac{L^2 T d \gamma_T \ln T}{\ln(1 + \sigma^{-2})} \right]^{1/2}\right)$$

Proof.

$$\begin{aligned} \mathbb{E}R_C(T) &= \mathbb{E} \left[\sum_{t=1}^T \left(1 - \frac{s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \left(\frac{\max_{x \in X} s_{\lambda_t}(u_f(x))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} - \frac{s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right) \right] \end{aligned}$$

Using Lemma 5 from Paria et al. (2019), we have the following decomposition for UCB:

$$\begin{aligned} \mathbb{E}R_C(T) &\leq \underbrace{\mathbb{E} \left[\sum_{t=1}^T \frac{U_t(\lambda_t, x_t)}{\max_{x \in X} s_{\lambda_t}(u_f(x))} - \frac{s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right]}_{B1} + \\ &\quad \underbrace{\mathbb{E} \left[\sum_{t=1}^T \frac{s_{\lambda_t}(u_f([x_t^*]_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} - \frac{U_t(\lambda_t, [x_t^*]_t)}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right]}_{B2} + \\ &\quad \underbrace{\mathbb{E} \left[\sum_{t=1}^T \frac{s_{\lambda_t}(u_f(x_t^*))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} - \frac{s_{\lambda_t}(u_f([x_t^*]_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right]}_{B3} \end{aligned}$$

By using Definition 12, we extend Lemma 3 from Paria et al. (2019) and obtain the following:

Lemma 16.

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T \frac{U_t(\lambda_t, x_t)}{\max_{x \in X} s_{\lambda_t}(u_f(x))} - \frac{s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right] \\ & \leq \mathbb{E} \left[\left(L\beta_T \sum_{t=1}^T M_{\lambda_t}^2 \right)^{1/2} \left(\sum_{\ell=1}^L \frac{\gamma T \ell}{\ln(1 + \sigma_\ell^{-2})} \right)^{1/2} \right] \\ & \quad + \frac{\pi^2 L \mathbb{E}[M_{\lambda}]}{6 |X|} \end{aligned}$$

By using Definition 12, we extend Lemma 2 from (Paria et al., 2019) and obtain the following:

Lemma 17. $\mathbb{E} \left[\sum_{t=1}^T \frac{s_{\lambda_t}(u_f(x_t^*))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} - \frac{U_t(\lambda_t, x_t^*)}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right] \leq \frac{\pi^2}{6} \mathbb{E}[M_{\lambda}] L$

By using Definition 12, we extend Equation (17) from (Paria et al., 2019) and obtain the following:

Lemma 18. $\mathbb{E} \left[\left| \frac{s_{\lambda}(u_f(x))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} - \frac{s_{\lambda}(u_f([x]_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right| \right] \leq L \mathbb{E}[M_{\lambda}] \frac{1}{t^2}$

Finally, we use Lemma 16 to bound B1, Lemma 17 to bound B2, Lemma 18 to bound B3, and obtain:

$$\mathbb{E} R_C(T) \leq C_1 L \mathbb{E}[M_{\lambda}] + C_2 \bar{M}_{\lambda} \left(LT(d \ln T + d \ln d) \sum_{\ell=1}^L \frac{\gamma T \ell}{\ln(1 + \sigma_\ell^{-2})} \right)^{1/2} \quad (8)$$

which converges to 0 as $T \rightarrow \infty$. □

Theorem 19. *The expected Bayes MF regret can be upper bounded as:*

$$\mathbb{E} R_B(T) \leq \frac{1}{T} \mathbb{E} R_C(T) + o(1)$$

As a result, showing that the MF utility ratio converges to 1 as $T \rightarrow \infty$ during dense sampling.

Proof. We assume that Λ is a bounded subset of a normed linear space. Following the approach in Paria et al. (2019), we begin by relating the sampling distribution to the empirical distribution. Let \hat{p} denote the empirical distribution corresponding to the samples $\{\lambda_t\}_{t=1}^T$. Consider the Wasserstein (Earth Mover's) distance between the sampling distribution $p(\lambda)$ and \hat{p} :

$$W_1(p, \hat{p}) = \inf_q \{ \mathbb{E}_q \|X - Y\|_1 : q(X) = p, q(Y) = \hat{p} \} \quad (9)$$

where q is a joint distribution on random variables X, Y with marginals p and \hat{p} respectively.

We can then use Definition 13 to bound the following in the MF setting:

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T \left(1 - \frac{s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right) - \mathbb{E} \left[1 - \frac{\max_{x \in D} s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right] \\ & \geq \frac{1}{T} \sum_{t=1}^T \left(1 - \frac{\max_{x \in D} s_{\lambda_t}(u_f(x))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right) - \mathbb{E} \left[1 - \frac{\max_{x \in D} s_{\lambda_t}(u_f(x))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right] \\ & \geq \mathbb{E}_{q(Z, Y)} \left[\left(1 - \frac{\max_{x \in D} s_Y(u_f(x))}{\max_{x \in X} s_Y(u_f(x))} \right) - \left(1 - \frac{\max_{x \in D} s_Z(u_f(x))}{\max_{x \in X} s_Z(u_f(x))} \right) \right] \\ & \geq -\mathbb{E}_{q(Z, Y)} \{ J \|Z - Y\|_1 \} \end{aligned}$$

Taking the infimum with respect to q and expectation with respect to the history \mathcal{H}_t :

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \left(1 - \frac{s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right) \right] - \mathbb{E} \left[1 - \frac{\max_{x \in D} s_{\lambda_t}(u_f(x_t))}{\max_{x \in X} s_{\lambda_t}(u_f(x))} \right] \geq -JEW_1(p, \hat{p}) \quad (10)$$

Using the fact that $\mathbb{E}[\max_{x \in \mathcal{X}} s_{\lambda}(u_f(x))] = \mathbb{E}[\max_{x \in \mathcal{X}} s_{\lambda_t}(u_f(x))]$, we obtain:

$$\mathbb{E}R_B(T) \leq \frac{1}{T} \mathbb{E}R_C(T) + JEW_1(p, \hat{p}) \quad (11)$$

By Theorem 15 and results from Paria et al. (2019), the first term converges to zero at rate $O^*(T^{-1/2})^4$ and the second term converges to zero at rate $O^*(T^{-1/D})$ for $D \geq 2$ (Canas and Rosasco, 2012), where D is the dimension of Λ . As a result, $\mathbb{E}R_B(T) \rightarrow 0$ as $T \rightarrow \infty$. Since $R_B(T)$ is the inverse of $U_B(T)$, from Definition 9, $\mathbb{E}U_B(T) \rightarrow 1$ as $T \rightarrow \infty$. \square

Lemma 20. *For a fixed $\lambda \in \Lambda$, the augmented Chebyshev scalarization function $s_{\lambda}(y) = -\max_{\ell \in [L]} \{\lambda_{\ell} |y_{\ell} - z_{\ell}^*|\} - \gamma \sum_{\ell=1}^L |y_{\ell} - z_{\ell}^*|$, as described in Section F.1, satisfies the assumption in Definition 12.*

Proof. Let $\lambda \in \Lambda$. Recall $\Lambda := \Delta^L$, thus is bounded. Furthermore, $\text{Im}(u_f)$ is bounded. First, we demonstrate that $s_{\lambda}(y)$ is Lipschitz w.r.t. y . For $\ell \in [L]$,

$$\begin{aligned} \frac{\partial s_{\lambda}(y)}{\partial y_{\ell}} &= \begin{cases} -\lambda_{\ell^*}, & y_{\ell^*} > z_{\ell^*}^*, \ell = \ell^* \\ \lambda_{\ell^*}, & y_{\ell^*} < z_{\ell^*}^*, \ell = \ell^* \\ 0, & \ell \neq \ell^* \end{cases} \Bigg|_{\ell^* := \arg \max_{\ell} \{\lambda_{\ell} |y_{\ell} - z_{\ell}^*|\}} - \gamma \begin{cases} 1 & y_{\ell} > z_{\ell}^* \\ -1 & y_{\ell} < z_{\ell}^* \end{cases} \\ \implies \left\| \frac{\partial s_{\lambda}(y)}{\partial y} \right\| &\leq \lambda_{\ell^*} \Bigg|_{\ell^* := \arg \max_{\ell} \{\lambda_{\ell} |y_{\ell} - z_{\ell}^*|\}} + L \end{aligned}$$

As the partial derivative is bounded w.r.t. y , by Mean Value Theorem (MVT) the scalarization function is Lipschitz w.r.t. y . Thus, there exists some constant C_{λ} , such that for all $y_1, y_2 \in \text{Im}(u_f)$,

$$|s_{\lambda}(y_1) - s_{\lambda}(y_2)| \leq C_{\lambda} \|y_1 - y_2\|$$

Note that for fixed λ , $\max_{y \in \text{Im}(u_f)} s_{\lambda}(y)$ is a constant. Using that fact, and the equation above, Definition 12 follows. \square

Lemma 21. *The augmented Chebyshev scalarization function $s_{\lambda}(y) = -\max_{\ell \in [L]} \{\lambda_{\ell} |y_{\ell} - z_{\ell}^*|\} - \gamma \sum_{\ell=1}^L |y_{\ell} - z_{\ell}^*|$, as described in Section F.1, satisfies the assumption in Definition 13 when $y \in S$ such that S bounded and $S \subseteq \text{Im}(u_f)$.*

Proof. We assume the following: $y \in S \setminus \{-\infty\}$ and $\exists \ell \in [L]$ s.t. $y_{\ell} = M$, where M is the upper bound of the MF. First, we apply Danskin's Theorem (Danskin, 1966) to $\max_{y \in \text{Im}(u_f)} s_{\lambda}(y)$. If $s_{\lambda}(y)$ is convex in λ for all $y \in \text{Im}(u_f)$, and at a given $\lambda_o \in \Lambda$, \exists a unique maximizer $y_o \in \text{Im}(u_f)$ then $s_{\lambda}(y)$ is differentiable w.r.t. λ at λ_o

with derivative $\left. \frac{\partial s_{\lambda}(y_o)}{\partial \lambda} \right|_{\lambda=\lambda_o}$.

As a result, if y_o is the unique maximizer at λ_o , then

$$\left. \frac{\partial s_{\lambda}(y_o)}{\partial \lambda} \right|_{\lambda=\lambda_o} = (|y_{o,\ell^*} - z_{\ell^*}^*|) \Bigg|_{\ell^* := \arg \max_{\ell} [\lambda_{\ell} |y_{\ell} - z_{\ell}^*|]}$$

Then, we demonstrate that $s_{\lambda}(y') / \max_{y \in \text{Im}(u_f)} s_{\lambda}(y)$ is Lipschitz w.r.t. λ . To conserve space, we denote

$\max_{y \in \text{Im}(u_f)} s_{\lambda}(y)$ with A .

$$\begin{aligned}
 \frac{\partial}{\partial \lambda} \left[\frac{s_{\lambda}(y')}{A} \right] \Big|_{\lambda=\lambda_o} &= \frac{\left[(A)(|y'_{\ell^*} - z_{\ell^*}^*|) \Big|_{\ell^*=\arg \max_{\ell} [\lambda_{\ell} |y'_{\ell} - z_{\ell}^*|]} \right]}{[A]^2} \\
 \frac{\left[(s_{\lambda}(y'))(|y_{o,\ell^*} - z_{\ell^*}^*|) \Big|_{\ell^*=\arg \max_{\ell} [\lambda_{\ell} |y'_{\ell} - z_{\ell}^*|]} \right]}{[A]^2} &= \underbrace{\frac{\left[(|y'_{\ell^*} - z_{\ell^*}^*|) \Big|_{\ell^*=\arg \max_{\ell} [\lambda_{\ell} |y'_{\ell} - z_{\ell}^*|]} \right]}{[A]}}_{C1} \\
 \underbrace{\frac{\left[(s_{\lambda}(y'))(|y_{o,\ell^*} - z_{\ell^*}^*|) \Big|_{\ell^*=\arg \max_{\ell} [\lambda_{\ell} |y'_{\ell} - z_{\ell}^*|]} \right]}{[A]^2}}_{C2} &
 \end{aligned}$$

The numerator for term C2 is bounded since s_{λ} is in a bounded space. The numerator for C1 is upper bounded since y , our MF, is an upper-bounded function.

The denominator values for C1 and C2 are both lower-bounded: since λ lies on the probability simplex, there must be a $\lambda_{\ell^\dagger} \geq 1/L$ for some $\ell^\dagger \in L$, by the pigeonhole principle. Since we are maximizing over $y \in \text{Im}(u_f)$, there must be a y^\dagger at the upper bound of $\text{Im}(u_f)$, which we denote as M . As a result, $\max_{y \in \text{Im}(u_f)} s_{\lambda}(y) \geq M/L$. Since the denominators of both C1 and C2 are lower-bounded, the overall partial derivative is bounded w.r.t. λ .

Since the overall partial derivative is bounded w.r.t. λ , by MVT $s_{\lambda}(y') / \max_{y \in \text{Im}(u_f)} s_{\lambda}(y)$ is Lipschitz w.r.t. λ . Thus, there exists some constant J such that for all $\lambda_1, \lambda_2 \in \Lambda$

$$\left| \frac{s_{\lambda_1}(y)}{\max_{y' \in \text{Im}(u_f)} s_{\lambda_1}(y')} - \frac{s_{\lambda_2}(y)}{\max_{y' \in \text{Im}(u_f)} s_{\lambda_2}(y')} \right| \leq J \|\lambda_1 - \lambda_2\|_1$$

□

E ADDITIONAL EXPERIMENTAL RESULTS

E.1 Concrete MF: Soft-Hard Functions

Effectively leveraging expert knowledge is a central challenge in multi-objective optimization, particularly when function evaluations are costly. This hurdle is especially acute in high-stakes domains like healthcare, where practitioners must balance treatment efficacy against harmful side-effects. Consider brachytherapy for cancer treatment (Deufel et al., 2020). Clinicians must design a radiation plan that maximizes dose to the tumor while minimizing exposure to nearby healthy organs. Each potential plan represents a complex trade-off, and devising or evaluating plans is time-consuming, often occurring under clinical pressure (Bélanger et al., 2019; Cui et al., 2018; van der Meer et al., 2020). Crucially, clinicians and other domain experts approach such tasks with practical knowledge, often conceptualizing their objectives not just as values to be purely minimized or maximized, but in terms of both *desired targets* (which we term "soft bounds" or aims) and *strict, non-negotiable boundaries* (which we term "hard bounds" or limits). For instance, a clinician might *aim* for >95% tumor coverage (a soft bound, indicating a desirable region) but require *at least* 90% (a hard bound, defining a feasibility constraint), while *aiming* for <513 centigrays (cGy) bladder dose (soft bound) with a strict *limit* of 601 cGy (hard bound) (Viswanathan et al., 2012). This natural way of specifying preferences—defining both acceptable ranges and ideal targets for each objective—is prevalent yet not fully leveraged by existing MOO frameworks for specifying

tradeoff regions (Paria et al., 2019; Abdolshah et al., 2019; Zuluaga et al., 2016; While et al., 2012; Suzuki et al., 2020). Examples of real-world applications where such priors already exist include brachytherapy, diabetes, and vitals monitoring (Grosman et al., 2010; Cairoli et al., 2019; Cui et al., 2018). Our explicit usage of MFs into the optimization process aims to fill this gap by operationalizing these prevalent soft and hard bounds using *soft-hard functions* (SHFs).

Let φ denote a generic objective function (e.g., some f_ℓ) with its associated soft bound α_S and hard bound α_H . We define its SHF as follows:

$$u_\varphi(x) = \begin{cases} 1 + 2\beta \times (\tilde{\alpha}_\tau - \tilde{\alpha}_S) & \varphi(x) \geq \alpha_\tau \\ 1 + 2\beta \times (\tilde{\varphi}(x) - \tilde{\alpha}_S) & \alpha_S < \varphi(x) < \alpha_\tau \\ 1 & \varphi(x) = \alpha_S \\ 2 \times \tilde{\varphi}(x) & \alpha_H < \varphi(x) < \alpha_S \\ 0 & \varphi(x) = \alpha_H \\ -\infty & \varphi(x) < \alpha_H \end{cases} \quad (12)$$

where $\tilde{\varphi}(x)$ and $\tilde{\alpha}$ are the soft-hard bound normalized values ¹, α_τ , the saturation point, determines where the utility values begin to saturate (as previously described, to prevent exploding utility values) ², and $\beta \in [0, 1]$ determines the fraction of the original rate of utility, in $[\alpha_H, \alpha_S]$, obtained within $[\alpha_S, \alpha_\tau]$.

E.2 Other Monotonic Utility Functions

We conducted several experiments to illustrate our method’s flexibility in adopting alternative monotonic functions. We provide brief motivation from other disciplines including additive manufacturing, energy systems and consumer behavior (business-related), and video streaming quality and evaluate on Branin-Currin.

Application Domain(s): Additive Manufacturing (Engineering & Science)

Motivation: Haghanifar et al. (2020) studied a numerical simulation to design an additive manufacturing strategy for minimizing the reflection of light at different angles. Malkomes et al. (2021) proposed to observe a diverse set of samples satisfying a set of hard boundary conditions.

Monotonic Utility Function: Hard-bounds only function. In this setting, the soft bound is effectively vacuous and the utility function is upper bounded at a user-specified value since we require upper-bounded utility functions.

Application Domain(s): Energy Systems and Consumer Behavior (Business Studies)

Motivation: In the optimization of real-time pricing for energy systems, monotonic utility functions are often used to model consumer satisfaction, to minimize electricity costs, etc. (Zhang et al., 2022; Samadi et al., 2010). Zhang et al. (2022) and Samadi et al. (2010) proposed using sigmoid utility functions and functions which resemble soft-hard bounds (to denote preferred levels of utility).

Monotonic Utility Function: Sigmoid utility functions and soft-bounds only functions may satisfy the assumptions we have in our paper. We implement the sigmoid utility function to be bounded between 0 and 1 while also allowing for the decision-maker to specify where the inflection/midpoint of the sigmoid function should be (intuitively representing the region of highest rate of increase in utility). The sigmoid function also smoothly models diminishing returns which is beneficial for the above domains.

Additional Potential Application Domains: Satisfaction Levels in Video Streaming Quality, Forest Management

Motivation: In the domain of video streaming services, Fazio and Tropea (2021) utilizes multi-step utility functions to model user satisfaction, where each discrete step corresponds to a specific video encoding layer and

¹Normalization, for value z , is performed according to the soft and hard bounds, α_S and α_H , respectively, using: $\tilde{z} = ((z - \alpha_H) / (\alpha_S - \alpha_H)) * 0.5$

²In our experiments, we determine α_τ to be $\alpha_H + \zeta(\alpha_S - \alpha_H)$, for $\zeta > 1.0$. We set $\zeta = 2.0$.

its associated utility. Similarly, in forest management, step utility functions are employed to balance the trade-off between maximizing economic profit from harvesting and minimizing ecological impact (Rădulescu et al., 2019).

Other Potential Monotonic Utility Functions: Piecewise step functions, polynomial splines

We implemented the above utility function and observed the following end-to-end results (over 6 trials) for Branin-Currin in Figure 5.

End-to-End Performance with Different Monotonic Utility Functions
(Branin-Currin, Complete-Mid)

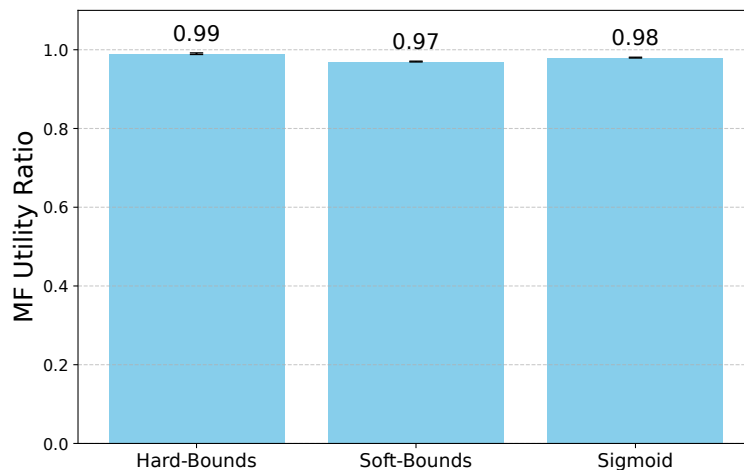


Figure 5: End-to-End MF Utility Ratio at 5 Points for different forms of Monotonic Utility Functions.

From a qualitative standpoint, the final subset from each of the above utility functions often exhibited higher density around the high-utility regions. The high MF utility ratio values quantitatively indicate that as well.

E.3 Importance of Two-Step Approach

Our central challenge is solving the max-min objective (Equation 1) to find a small set of points robust to unknown user weights. We decompose this into two steps because current literature lacks a principled single-step approach for this specific goal.

By splitting the problem, we maintain strong theoretical guarantees for both phases:

Dense Sampling (Step 1): We prove convergence to the high-utility region (Appendix D) using Bayesian Optimization, effectively discretizing the continuous search space.

Sparsification (Step 2): We leverage submodularity to provide optimality guarantees on the final, smaller set (Section 4).

A monolithic approach would sacrifice these guarantees. Relying solely on Step 1 risks overwhelming the DM with redundant points, while Step 2 requires a high-quality discrete candidate set to function. Our decomposition ensures we achieve both dense coverage (Step 1) and optimal sparsity (Step 2) with formal backing for each stage.

To demonstrate this from an empirical perspective, we implemented and ran several one-step approaches: (1) a simulated annealing approach where we directly optimize for Equation 1 of our paper, rather than split it into two separate steps, and (2) ϵ -PAL (Zuluaga et al., 2016), which is a one-step approach that has a component that acts as a form of regularization on the sparsity of the solutions. We display end-to-end results for Branin-Currin in Figure 6.

The one-step simulated annealing approach typically takes many more function evaluations to converge compared to MoSH, especially since MoSH is designed for expensive function evaluations (a critical assumption in our setting, due to the time-critical nature of applications such as Brachytherapy cancer treatment). As a result, a budget of 50 function evaluations for simulated annealing is not nearly enough for it to obtain points well-distributed in the

region defined by the soft-hard bounds. This is especially evident for the real-world Brachytherapy application, which is of much higher dimensionality and would require significantly more function evaluations to converge for simulated annealing.

ϵ -PAL (one-step), while designed for expensive black-box functions, does not take into consideration any expert-defined priors (e.g. our soft-hard bounds). As a result, ϵ -PAL spends a lot of time obtaining points along the entire Pareto frontier, as opposed to the region the decision-maker truly cares about. As a result, it is difficult for the decision-maker to make high-utility decisions in their preferred tradeoff regions.

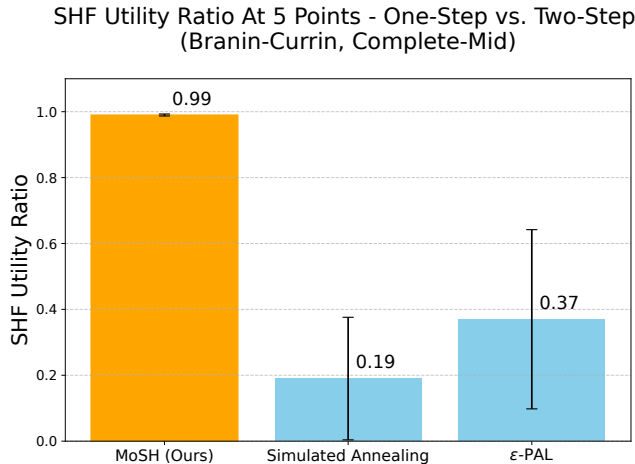


Figure 6: End-to-End SHF Utility Ratio at 5 points results comparing our two-step approach (MoSH) against two other one-step approaches.

E.4 Alternative Acquisition Functions

To illustrate the performance of MoSH across acquisition functions, we conducted several ablations comparing UCB to Thompson Sampling on the Branin-Currin, Four Bar Truss, and Brachytherapy applications. We present the end-to-end results in Figure 7 (over 6 trials):

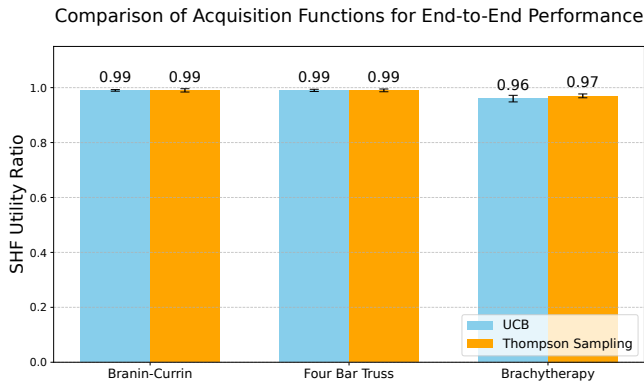


Figure 7: End-to-End SHF Utility Ratio results at 5 points comparing the Upper Confidence Bound (UCB) and Thompson Sampling (TS) acquisition functions.

We did not observe statistically significant differences in the end-to-end results between UCB and Thompson Sampling. Both methods converged similarly to the high-utility regions defined by the Soft-Hard bounds. Given the empirical similarity in performance, we believe it could be helpful to have heuristics based on secondary practical considerations rather than raw performance. Below are several high-level ones (among others):

UCB: When explicit control over the exploration-exploitation trade-off is desired or when deterministic reproducibility is desired (Russo and Van Roy, 2014).

Thompson Sampling: To simplify the design and implementation process, especially when UCB algorithms are computationally challenging (Russo and Van Roy, 2014).

E.5 Computational Feasibility

Equation 2 of the dense sampling step uses a conversion from the worst-case to an average-case minimization problem. To observe the difficulty of the worst-case minimization problem and how the results may be affected, we performed several experiments which directly solve the following:

$$\max_{D \subseteq X, |D| \leq k_D} \min_{\lambda \in \Lambda} \left[\frac{\max_{x \in D} s_{\lambda}(u_f(x))}{\max_{x \in X} s_{\lambda}(u_f(x))} \right] \tag{13}$$

To conduct the experiments, we instantiate the MFs with the soft-hard functions (SHFs) from Section 5.1 and Appendix E.1. When using SHFs, we refer to the dense sampling step of our method as MoSH-Dense.

We solve Equation 13 using a greedy algorithm over discretized input space X and weight space Λ and compare the results with both MoSH-Dense and its discretized variations.

Discretization. To create the finite approximation of the continuous input space, we discretize the domain using a uniform grid. Given input space bounds $[\mathbf{l}, \mathbf{u}] \subset \mathbb{R}^d$ where $\mathbf{l} = [l_1, \dots, l_d]$ and $\mathbf{u} = [u_1, \dots, u_d]$ are the lower and upper bounds respectively, we construct a discretized grid as follows:

For each dimension $i \in [d]$, we create an evenly-spaced sequence:

$$\mathcal{X}_i = \{l_i, l_i + \delta, l_i + 2\delta, \dots, u_i\} \tag{14}$$

where δ is the step size parameter controlling the granularity of the discretization.

The complete discretized input space $\mathcal{X}_{\text{disc}}$ is then constructed as the Cartesian product:

$$\mathcal{X}_{\text{disc}} = \mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_d \tag{15}$$

This results in a finite grid of points where each point $\mathbf{x} \in \mathcal{X}_{\text{disc}}$ represents a candidate solution in the original continuous space. The total number of points in the discretized space is $\prod_{i=1}^d \lceil \frac{u_i - l_i}{\delta} \rceil$. The same discretization method is done for Λ , except, from this grid, we select only those points that lie on the probability simplex, i.e., whose components sum to 1.

Experiments and Discussions. We term the baseline algorithms as *discrete-greedy- δ* and *discrete-average- δ* and primarily experiment with the Branin-Currin objective function (described in Section E.6). We show numerical results in Figure 8. In general, computationally, *discrete-greedy- δ* and *discrete-average- δ* take longer as $\delta \rightarrow 0$. Specifically, *discrete-greedy-0.05* and *discrete-greedy-0.10* take, on average, 20x and 3x longer (in seconds) than MoSH-Dense, respectively, to select the next sample x_t at iteration t . This computational disparity is further exacerbated as the size of the input space dimensionality increases. However, as expected, as $\delta \rightarrow 0$, the metrics improve and move closer to those of MoSH-Dense. Finally, we note that, as described in Section 4, the greedy algorithm may perform arbitrarily bad when solving Equation 13. We leave for future work exploration of other algorithms which are designed for our setting described in Section 3; notably, access to some noisy and expensive black-box function. As these results show, our formulation for MoSH-Dense leads to superior results while also being more computationally efficient and simple.

E.6 Synthetic Two-Objective Function: Branin-Currin

We leverage the Branin-Currin synthetic two-objective optimization problem provided in the BoTorch framework (Balandat et al., 2020), which has a mapping of $[0, 1]^2 \rightarrow \mathbb{R}^2$. To demonstrate our method’s flexibility in accommodating various configurations, we sample from the following variations : (1) *complete-mid*, where the hard bounds cover the *complete* Pareto frontier and the soft bounds are in the *middle* of the hard region (2)

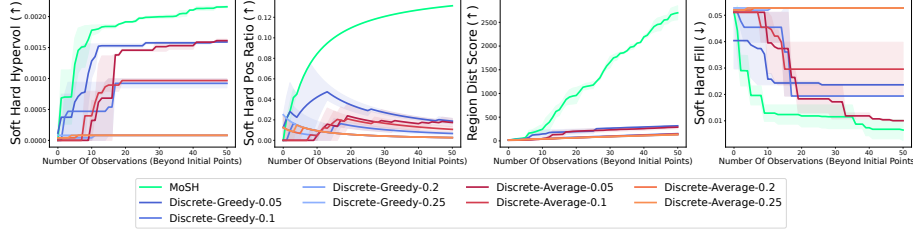


Figure 8: Step (1) results with discrete greedy algorithms which explicitly solve for the max-min problem for MoSH-Dense (Equation 13). The plots compare our proposed method, MoSH-Dense, against several variations of *discrete-greedy*- δ and *discrete-average*- δ using the metrics defined in Section 5.3.

complete-top, (3) *complete-bot*, (4) *top-mid*, and (5) *bot-mid*. Figures 9, 10, 11, 12, 13 illustrate the plots, for step 1, with the performance metrics defined in Section 5.3 (computed over 6 independent trials). Figure 14 additionally compares against preference-based evolutionary multi-objective algorithm baselines (experimental details in Appendix F). We note that EHVI, although superior in some metrics, is extremely computationally demanding for higher dimensions. Overall, we observe that our algorithm generally matches or surpasses other baselines in all four metrics, with sampling a much higher density near the soft region. The figures for the other configurations display a similar pattern.

Configurations (normalized to $[0,1]$) : $\{\alpha_{0,S}, \alpha_{0,H}, \alpha_{1,S}, \alpha_{1,H}\}$

1. *Complete-Mid*: $\{0.988, 0.943, 0.856, 0.618\}$
2. *Complete-Top*: $\{0.969, 0.943, 0.935, 0.618\}$
3. *Complete-Bot*: $\{0.998, 0.943, 0.697, 0.618\}$
4. *Top-Mid*: $\{0.969, 0.940, 0.915, 0.856\}$
5. *Bot-Mid*: $\{0.996, 0.975, 0.737, 0.658\}$

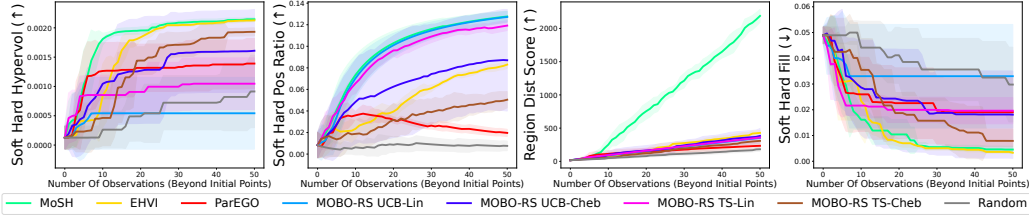


Figure 9: *Complete-Mid* configuration for the Branin-Currin synthetic two-objective function. Results are plotted using the metrics defined in Section 5.3.

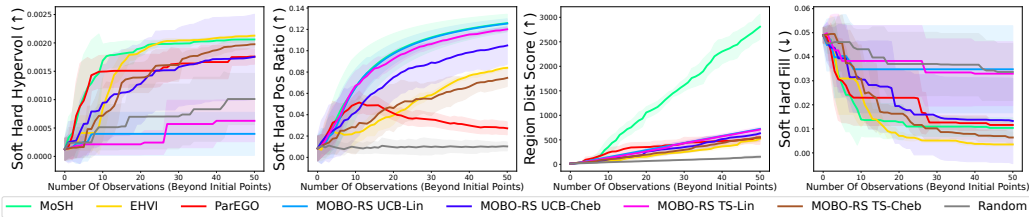


Figure 10: *Complete-Top* configuration for the Branin-Currin synthetic two-objective function. Results are plotted using the metrics defined in Section 5.3.

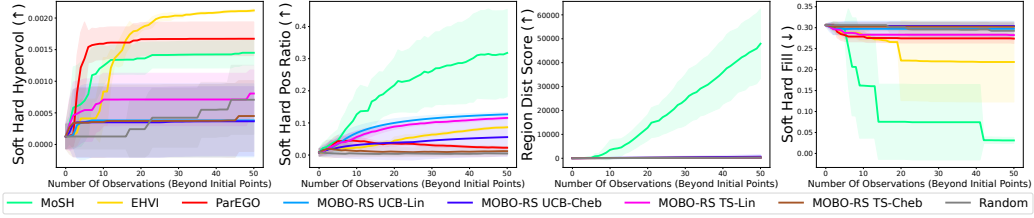


Figure 11: *Complete-Bot* configuration for the Branin-Currin synthetic two-objective function. Results are plotted using the metrics defined in Section 5.3.

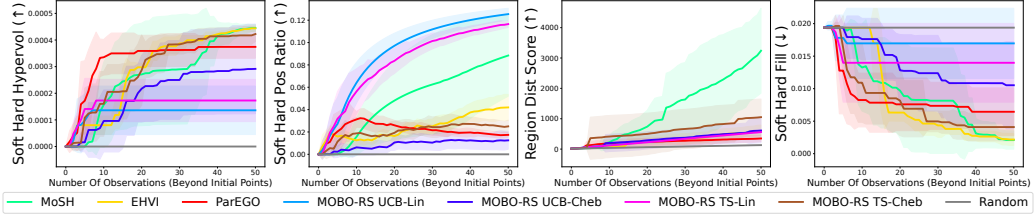


Figure 12: *Top-Mid* configuration for the Branin-Currin synthetic two-objective function. Results are plotted using the metrics defined in Section 5.3.

E.7 Engineering Design Problem: Four Bar Truss

We also evaluate on a MOO engineering design problem, Four Bar Truss, from REPROBLEM (Tanabe and Ishibuchi, 2020), which consists of two objectives and four continuous decision variables (CHENG and LI, 1999). The objectives of the problem are to minimize the structural volume and the joint displacement of the four bar truss. The four decision variables determine the length of the four bars, respectively. To demonstrate our method’s flexibility, we sample from the following variations: (1) *narrow-mid*, (2) *narrow-bot*, (3) *narrow-top*, (4) *bot-mid*, and (5) *top-mid* (over 6 independent trials). Step 1 results are shown in Figures 15, 16, 17, 18, and 19. Figure 20 additionally compares against preference-based evolutionary multi-objective algorithm baselines (experimental details in Appendix F). Based on our experiments, our algorithm consistently matches or surpasses other baselines in all metrics, while sampling at a clearly higher density near the soft region.

Configurations (normalized to [0,1]) : $\{\alpha_{0,S}, \alpha_{0,H}, \alpha_{1,S}, \alpha_{1,H}\}$

1. *Narrow-Mid*: $\{0.62, 0.45, 0.72, 0.55\}$
2. *Narrow-Bot*: $\{0.70, 0.45, 0.65, 0.55\}$
3. *Narrow-Top*: $\{0.55, 0.45, 0.78, 0.55\}$
4. *Bot-Mid*: $\{0.86, 0.70, 0.48, 0.25\}$
5. *Top-Mid*: $\{0.43, 0.20, 0.85, 0.70\}$

E.8 Real Clinical Case: Cervical Cancer Brachytherapy Treatment Planning

We had consent for and leveraged a real clinical case which had previously been treated in the clinic. As a result, we had access to the fully anonymized patient CT data (in the form of DICOM files) and the final treatment plan which had been administered to the patient. For our experiments, we obtained a single treatment plan via a linear program formulated as an epsilon-constraint method (Deufel et al., 2020). The three input parameters to that linear program, which implicitly control the weights of the different objectives (radiation dosage to the bladder, rectum, bowel, and cancer tumor), were employed as the decision variables.

With MoSH, we are able to obtain an informative and concise subset of the treatment plans with metrics which align with the clinicians’ priors . As shown in Figure 1, we are even able to obtain treatment plans which

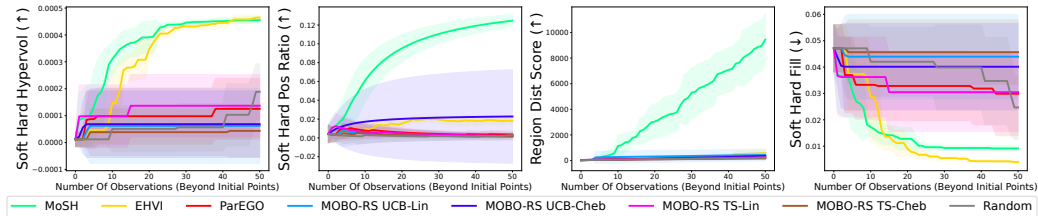


Figure 13: *Bot-Mid* configuration for the Branin-Currin synthetic two-objective function. Results are plotted using the metrics defined in Section 5.3.

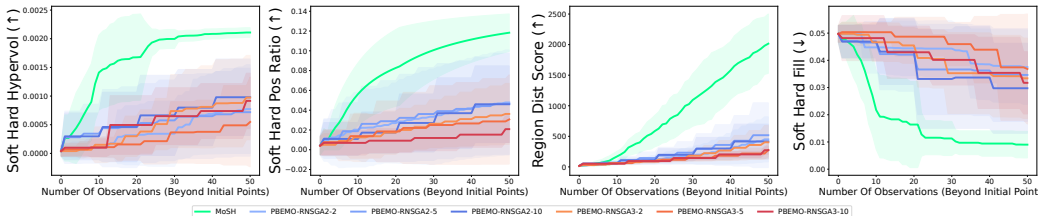


Figure 14: *Complete-Mid* configuration for the Branin-Currin synthetic two-objective function with preference-based evolutionary multi-objective (PBEMO) algorithm baselines. Results are plotted using the metrics defined in Section 5.3.

Pareto dominates those plans found from manual clinician planning, which typically takes significantly more time (~30min) (Morén et al., 2021). Furthermore, the soft and hard bound values we employed for the experiment are exactly the values that the clinician we worked with usually thinks about when performing treatment planning, so our setting closely mimics the real decision-making setting.

The configuration used was: $\{\alpha_{0,S} = 0.95, \alpha_{0,H} = 0.90, \alpha_{1,S} = 513, \alpha_{1,H} = 601, \alpha_{2,S} = 352, \alpha_{2,H} = 464, \alpha_{3,S} = 411, \alpha_{3,H} = 464\}$, where the objectives correspond to PTV_{V700} , $Bladder_{D2cc}$, $Rectum_{D2cc}$, and $Bowel_{D2cc}$, ordered. Before doing the experiment, all of the values were normalized to $[0,1]$ and converted to maximization. Figure 21 additionally compares against preference-based evolutionary multi-objective algorithm baselines (experimental details in Appendix F).

Application to other real-world domains. In this real-world application, the treatment planner typically already has an explicit value which could be used for the soft bound (and hard bound) prior to planning for each patient. Clinical domains, in particular, frequently have specific and standardized hard and soft targets for metrics of interest, such as insulin administration for diabetes, vitals, etc., not solely brachytherapy (Grosman et al., 2010; Cairolì et al., 2019). Our proposed method provides for a principled way of injecting such priors into the optimization process.

However, we acknowledge that there may be cases where the practitioner is unsure of where to set the soft bounds for each of the objectives. In practice, we envision that MoSH can be leveraged in an interactive setting where the soft and hard bounds are iteratively adjusted based on feedback or points observed at each iteration. We leave this for future work.

E.9 Deep Learning Model Selection: Fast and Accurate Neural Network

Similar to Hernández-Lobato et al. (2016), we seek to obtain a neural network which minimizes both prediction error and inference time, two competing objectives. We use the MNIST dataset and consider feedforward neural networks with six decision variables (Deng, 2012). Specifically, we use 50,000 and 10,000 different samples for training and validation, respectively, from the train split. Step 1 results (mean and std. dev computed over 6 independent trials) are shown in Figure 22. Although our algorithm does not surpass the baselines all four metrics, it still performs consistently relatively well in all four (the most consistently well out of the baselines) – in line with the other experimental results.

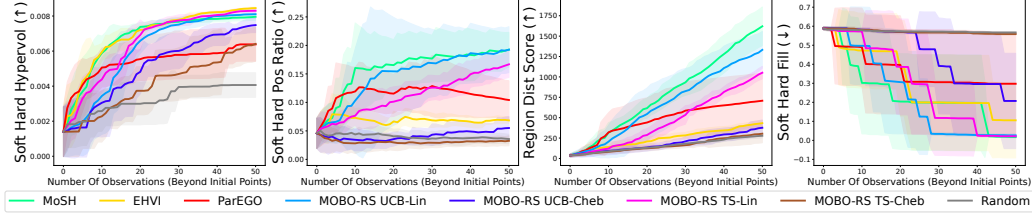


Figure 15: *Narrow-Mid* configuration for the Four Bar Truss engineering design two-objective function. Results are plotted using the metrics defined in Section 5.3.

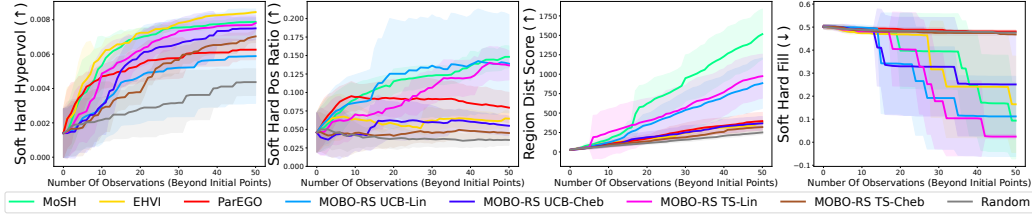


Figure 16: *Narrow-Bot* configuration for the Four Bar Truss engineering design two-objective function. Results are plotted using the metrics defined in Section 5.3.

We used the following decision space: number of hidden units per layer ($[50, 300]$), number of layers ($[1, 3]$), learning rate ($[0, 0.1]$), dropout amount ($[0.4, 0.6]$), l1 regularization ($[0, 0.05]$), and l2 regularization ($[0, 0.05]$). We trained each of the neural networks for 100 epochs and converted both of the objectives such that its a maximization problem. We used the following soft and hard bounds configuration (α_S, α_H) : $(0.87, 0.72)$, $(0.57, 0.37)$ for the dimensions corresponding to accuracy and inference time, respectively, which we assumed was reasonable. In practice, we envision that the soft and hard bounds can be interactively adjusted based on feedback or points observed at each iteration. As mentioned previously, we leave this for future work.

E.10 Large Language Model Personalization: Concise and Informative

We seek to obtain an LLM which generates both concise and informative outputs, two directly competing objectives. We leverage proxy tuning, which steers a large pre-trained model, M , by using the difference between the predicted logits of an expert model (a smaller, tuned model), M^+ and an anti-expert model (the smaller model, un-tuned), M^- (Liu et al., 2024; Mitchell et al., 2023; Shi et al., 2024). We leverage notation from Liu et al. (2024) and obtain the output distribution at time step t , conditioned on prompt $x_{<t}$, from the proxy-tuned model, \tilde{M} , in a two-objective setting as such: $p_{\tilde{M}}(X_t|x_{<t}) = \text{softmax}[s_M(X_t|x_{<t}) + \sum_{i=1}^2 \theta_i (s_{M_i^+}(X_t|x_{<t}) - s_{M_i^-}(X_t|x_{<t}))]$, where $s_M, s_{M_i^+}, s_{M_i^-}$ represent the logit scores for each model and θ_i denotes the input decision variable, the controllable weight applied to the logits difference associated with expert model i . Models M_i^+ , for $i = 0, 1$, are tuned according to the conciseness and informativeness objectives, respectively. By adjusting θ_i at decoding time, we obtain generated output distributions with varying tradeoffs in conciseness and informativeness.

We leveraged models from the TüLU-2 suite of families for our experiments: TüLU-2-13B for the large pre-trained model, M , and TüLU-2-DPO-7B for the expert and anti-expert models, M^+ and M^- (Iverson et al., 2023). M_1^+ and M_2^+ were tuned using direct preference optimization (DPO) with the preference datasets corresponding to the conciseness and informativeness dimensions, respectively, from Jang et al. (2023) and Rafailov et al. (2024). We used the following decision space: θ_1 ($[0.0, 1.5]$) and θ_2 ($[0.0, 1.5]$). We used the following soft and hard bounds configuration (α_S, α_H) : $(0.6, 0.5)$, $(0.6, 0.5)$ for the dimensions corresponding to conciseness and informativeness, respectively. To measure the conciseness dimension, we calculated the number of characters in the output response, which is then normalized to $[0, 1]$. To measure the informativeness dimension, we adapted the prompt from Mitchell et al. (2023) and used GPT-4 to provide a measure between 0 and 100 (which is then normalized):

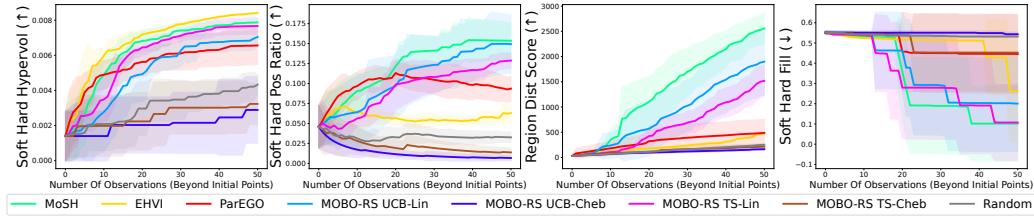


Figure 17: *Narrow-Top* configuration for the Four Bar Truss engineering design two-objective function. Results are plotted using the metrics defined in Section 5.3.

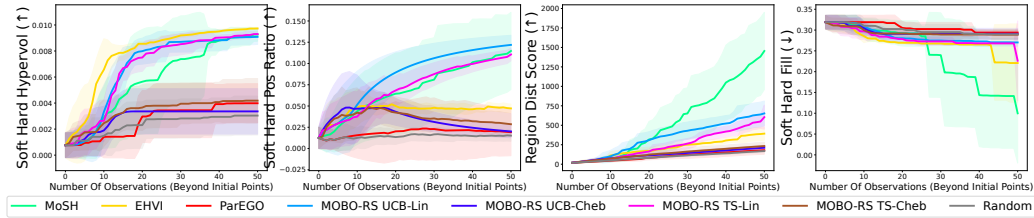


Figure 18: *Bot-Mid* configuration for the Four Bar Truss engineering design two-objective function. Results are plotted using the metrics defined in Section 5.3.

GPT-4 Informativeness Evaluation Prompt

“Score the following chatbot response to the given query on a continual scale from 0 (worst) to 100 (best), where a score of 0 means the response lacks information and provides very little information, and a score of 100 means the response is very informative and provides additional background which will be useful for the user.

QUERY: {{the query}}

RESPONSE: {{the response to evaluate}}

USE THE ANSWER FORMAT:

Reason: {analysis of query and response}

Informativeness score: {a single number 0-100, nothing else}”

Finally, each of the objectives were normalized and converted to maximization and we used the following five prompts for the experiment and averaged the results:

1. “Give three tips for staying healthy.”
2. “What are the three primary colors?”
3. “Describe the structure of an atom.”
4. “Rewrite the following instructions so that they are simpler and easier to understand. Login to the website, input your username and password, then click the “Continue” button.”
5. “How can we reduce air pollution?”

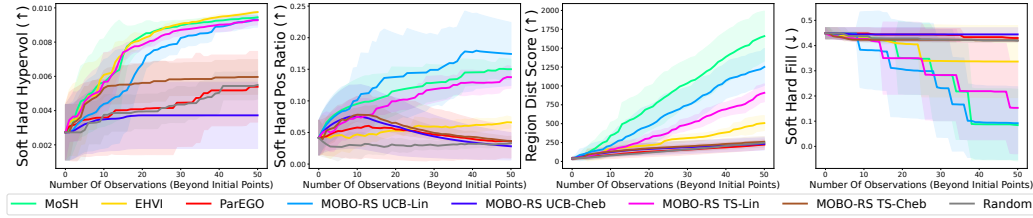


Figure 19: *Top-Mid* configuration for the Four Bar Truss engineering design two-objective function. Results are plotted using the metrics defined in Section 5.3.

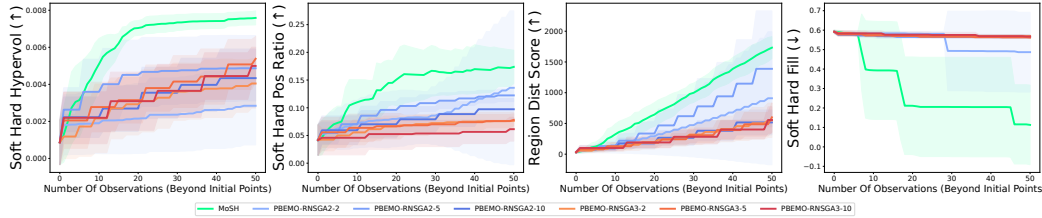


Figure 20: *Narrow-Mid* configuration for the Four Bar Truss engineering design two-objective function with preference-based evolutionary multi-objective (PBEMO) algorithm baselines. Results are plotted using the metrics defined in Section 5.3.

Sample LLM Prompt and Responses

Prompt: Give three tips for staying healthy.

Response 1

▷ θ_1, θ_2 : (1.07, 0.00) which corresponds to high conciseness and low informativeness, respectively
Response:

1. Maintain a healthy diet with plenty of fruits and vegetables.
2. Get regular exercise, at least 30 minutes a day.
3. Practice good hygiene and get enough sleep.

▷ GPT-4 Evaluated Informativeness: 50

Response 2

▷ θ_1, θ_2 : (1.07, 0.58) which corresponds to high conciseness and medium informativeness, respectively
Response:

1. Maintaining a healthy diet with a balanced intake of nutrients
2. Regular physical activity, at least 30 minutes a day.
3. Getting enough sleep, aim for 7-8 hours per night. (4 in case of power nap). bonus: limit stress. (5th tip: stay hydrated) (6th tip: practice good hygiene) (7th tip: limit consumption of processed foods)

▷ GPT-4 Evaluated Informativeness: 70

E.11 Easier Interpretation of Step 1 Results

We acknowledge that it may be more difficult to analyze and determine winners from all four metrics across all the tasks, however, we believe the variety and diversity of metrics and tasks is necessary to highlight the

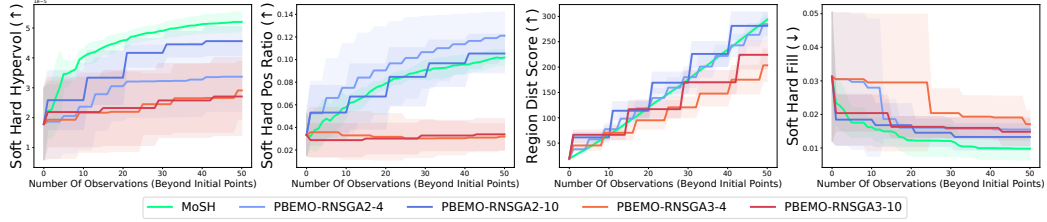


Figure 21: Brachytherapy application with preference-based evolutionary multi-objective (PBEMO) algorithm baselines. Results are plotted using the metrics defined in Section 5.3.

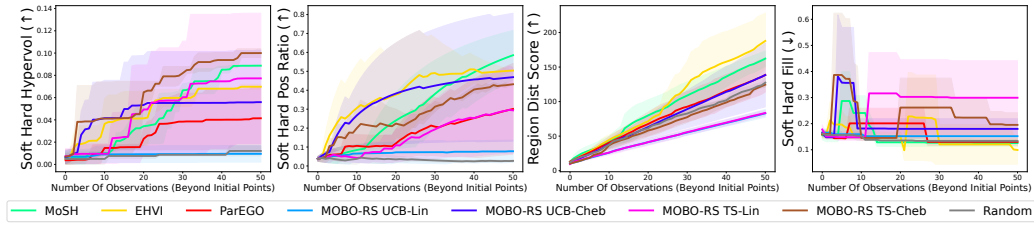


Figure 22: Deep learning model selection problem results, using the metrics defined in Section 5.3.

generalizability and strengths of MoSH. Furthermore, while baselines might occasionally match MoSH on one metric in one experiment, MoSH demonstrates remarkable consistency across all four metrics and all diverse tasks. This broad robustness is a key strength we would like to highlight below.

As a very naive metric, we count the number of times the mean value of each of the methods is within the top 2 highest (for soft-hard hypervolume, positive ratio, and region distance score) and top 2 lowest (for soft-hard fill distance) at the end of the 50 observations. We acknowledge that this is a very simple metric, and there may be other ways to view this, but we propose to exhibit this as a way to facilitate the evaluation process for the reader. Results are shown in Table 2.

Table 2: Aggregated comparison of MoSH against other baselines for step 1 for easier interpretation of results. The table displays the count of the number of times the mean value of each of the methods is within the top 2 highest (for soft-hard hypervolume, positive ratio, and region distance score) and top 2 lowest (for soft-hard fill distance) at the end of the 50 observations, across all applications. Metrics are abbreviated for space, i.e. Soft-Hard Hypervolume \rightarrow SH Hypervolume, Soft-Hard Positive Samples Ratio \rightarrow SH Pos. Ratio, Soft Region Distance-Weighted Score \rightarrow Region Dist. Score.

Method	SH Hypervolume	SH Pos. Ratio	Region Dist. Score	Soft Hard Fill
MoSH	10	13	13	10
EHVI	10	1	1	7
ParEGO	2	1	1	0
MOBO-RS UCB-Lin	0	8	6	1
MOBO-RS UCB-Cheb	0	1	1	0
MOBO-RS TS-Lin	3	1	1	3
MOBO-RS TS-Cheb	2	0	1	3
Random	0	0	0	0

As can be seen in Table 2, MoSH is clearly superior compared to the other baselines, across all metrics and tasks, and Pareto dominates all other baselines. We believe that this clearly illustrates how much MoSH is outperforming the baselines.

E.12 Step 2 Experiments: Synthetic and Real-World Applications

For Figure 4, each successive point that the DM views is determined by the sampling order of the points. Figure 23 shows the remaining evaluations for step (2) on the Branin-Currin, Four Bar Truss, and deep learning model selection problems. All results are displayed with standard deviation bars, with the randomness over a ground-truth λ^* , which was independently sampled 10 different times.

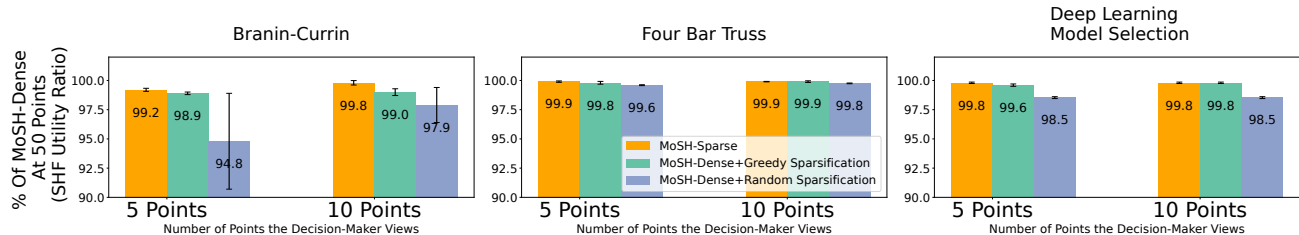


Figure 23: Step (2) evaluation on the Branin-Currin, Four Bar Truss, and deep learning model selection problems, illustrating the SHF utility ratio for the various methods relative to MoSH-Dense’s SHF utility ratio at 50 points, compared when the decision-maker views 5 and 10 points.

E.13 End-to-End Experiments: Synthetic and Real-World Applications

We conducted end-to-end evaluations of our method, MoSH, by varying over the dense set of points, obtained from step 1, provided to SATURATE. Figures 24, 25, 26, and 27 display bar plots representing the maximum SHF utility ratio obtained after the DM views 5 points, for each experiment. Figures 28, 29, and 30 display the results when compared against preference-based evolutionary multi-objective (PBEMO) algorithm baselines (experimental details in Appendix F). For algorithms which select less than 5 points, we use the SHF utility ratio value of the last point.

E.14 Discussion on Preference-Based Evolutionary Multi-Objective (PBEMO) Algorithm Baseline Results

Overall, our method, MoSH, consistently outperforms both of the PBEMO algorithms across our metrics and stages. We observed that the PBEMO algorithms often exhibited slower convergence and seemed to be sensitive to the initial population (illustrated by the higher variance in results). This slower convergence observation also seemed to be observed in other multi-objective Bayesian optimization works (Ji et al., 2024). As a result, the PBEMO-populated Pareto frontiers often were sparser than those of our method, MoSH. This behavior can be noticed in the results.

Furthermore, we found the reference point approach, which most PBEMO algorithms use, to be unsatisfactory in modeling the desired sub-regions of the tradeoff surface. Even when setting multiple reference points, we found it exceedingly difficult to ensure objective values are within certain hard thresholds, and more densely populated in preferred regions. Such behaviors can be seen in the PBEMO baselines’ often lower metrics in step 1. On the contrary, our method, MoSH, explicitly accounts for such user preferences. As a result, such characteristics lead to suboptimal end-to-end results for the PBEMO baselines, compared to those of MoSH. We believe this result further improves our extensive evaluation setting and solidifies the lead our method, MoSH, has in our problem setting.

F EXPERIMENTAL SETUP DETAILS

F.1 Baseline Methodologies

Step 1: Dense Pareto Frontier Sampling. Although none of the baseline algorithms are inherently designed for our setting with SHFs, we aimed to make the comparisons as fair as possible by using a heuristic, similar to what was described in (Paria et al., 2019), to determine the weight distribution $p(\lambda)$. In short, we used the

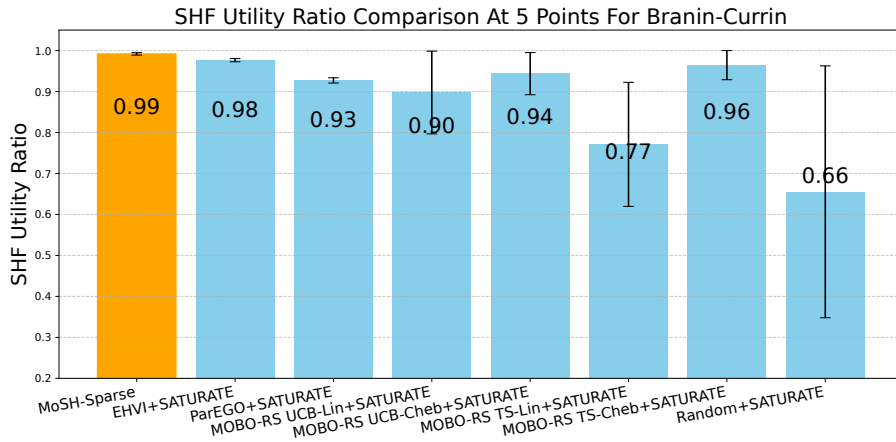


Figure 24: Bar plot illustrating the SHF utility ratio obtained by our method, MoSH-Sparse, compared to the other baselines on the Branin-Currin synthetic function. Only the dense set, from step 1, changes for each bar.

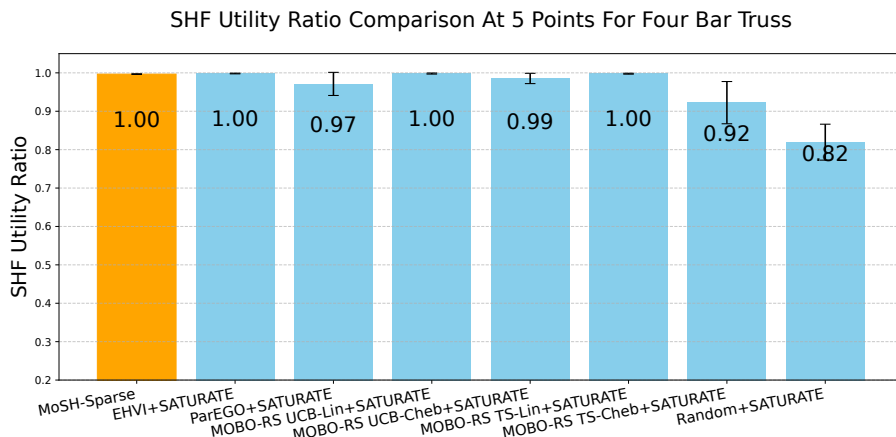


Figure 25: Bar plot illustrating the SHF utility ratio obtained by our method, MoSH-Sparse, compared to the other baselines on the Four Bar Truss application. Only the dense set, from step 1, changes for each bar.

heuristic $\lambda = \mathbf{u} / \|\mathbf{u}\|_1$, where $\mathbf{u}_\ell \sim N(\alpha_{\ell,S}, |\alpha_{\ell,H} - \alpha_{\ell,S}|/3)$, in order to roughly mimic the gradually decreasing weight between the soft bounds $\alpha_{\ell,S}$ and the hard bounds $\alpha_{\ell,H}$, $\forall \ell \in [L]$.

For our proposed method (MoSH), we use the augmented Chebyshev scalarization function (Wierzbicki, 1982), denoted as $s_\lambda(\tilde{f}_\ell(x)) = -\max_\ell[\lambda_\ell|\tilde{f}_\ell(x) - z_\ell^*|] - \gamma \sum_{i=\ell}^L |\tilde{f}_i(x) - z_i^*|$, where z^* is the ideal or utopian vector, L is the number of objective dimensions, and γ is a constant weighing the linear term being added to the traditional Chebyshev scalarization function (Chugh, 2019). We find that this scalarization function performs better since the Chebyshev component allows for non-convex Pareto frontier sampling and the augmented term assists with sampling within the hard bound. All results are displayed with standard deviation bars, computed over 6 independent trials.

Preference-Based Evolutionary Multi-Objective Algorithm Baselines. We ran additional experiments comparing our proposed method, MoSH, against two well-known preference-based EMO (PBEMO) algorithms, R-NSGA II (Deb and Sundar, 2006) and R-NSGA III (Vesikar et al., 2018). Overall, MoSH consistently outperforms both of the PBEMO algorithms across our various metrics and evaluation settings. Our work assumes we are in a setting with expensive function evaluations (Section 3). For fair comparison, we set a terminating condition of 50 function evaluations for both R-NSGA II and R-NSGA III (the same number of function evaluations used for the other multi-objective Bayesian optimization approaches). To closely represent the soft-hard bounds, the reference point for R-NSGA II and R-NSGA III was set to be at the intersection of the pre-defined soft-hard

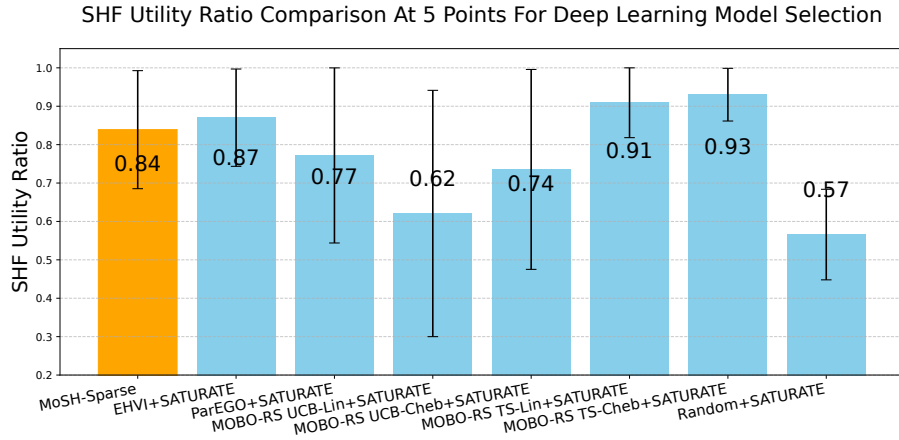


Figure 26: Bar plot illustrating the SHF utility ratio obtained by our method, MoSH-Sparse, compared to the other baselines on the deep learning model selection application. Only the dense set, from step 1, changes for each bar.

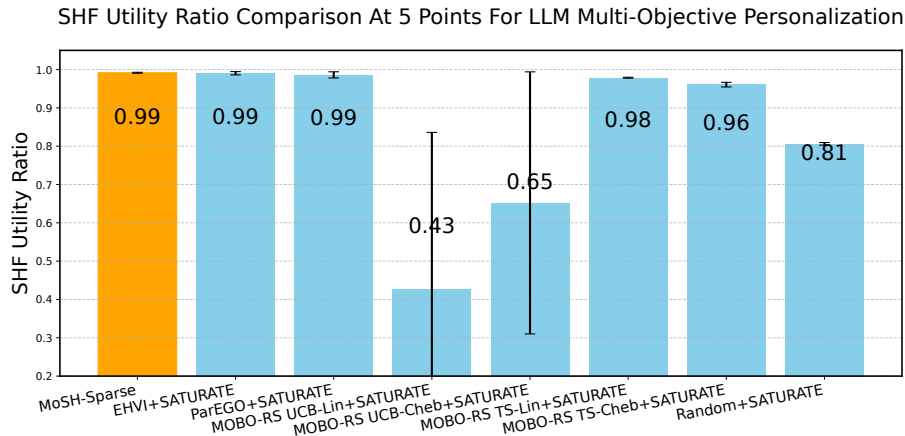


Figure 27: Bar plot illustrating the SHF utility ratio obtained by our method, MoSH-Sparse, compared to the other baselines on the LLM personalization problem. Only the dense set, from step 1, changes for each bar.

bounds. For comprehensiveness, we performed ablations across various parameters for the population size (for R-NSGA II) and the population size used for each reference point (for R-NSGA III). For R-NSGA II, we compared against population size values of 2, 5, and 10. For R-NSGA III, we compared against population size used for each reference point at the values of 2, 5, and 10. For the brachytherapy application, we used the parameters 4 and 10, due to population count requirements. We nicknamed them PBEMO-RNSGA2-x for R-NSGA II and PBEMO-RNSGA3-x for R-NSGA III, where x denotes the parameter value. Extreme points of the Pareto frontier were not set as reference points and the initial population used was set to be the same as the initialization points used for all the multi-objective Bayesian optimization approaches. Only results for the Branin-Currin and Brachytherapy applications are shown, due to character limits in the rebuttal response. Similar results were observed in the Four Bar Truss engineering design application. Finally, the experiments were conducted using the pymoo library (Blank and Deb, 2020).

Step 2: Pareto Frontier Sparsification. All of the baseline algorithms were run immediately after step (1) is completed. In order to ensure all of the algorithms sample an equal number of points, we first run MoSH-Sparse before deciding for greedy and random baseline algorithms to sample the same number of points. All results are displayed with standard deviation bars, with the randomness over a ground-truth λ^* , which was independently sampled 10 different times.

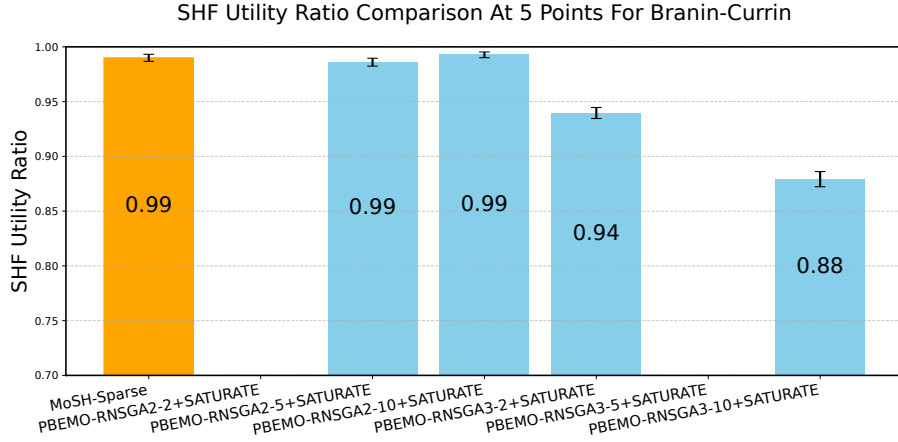


Figure 28: Bar plot illustrating the SHF utility ratio obtained by our method, MoSH-Sparse, compared to the preference-based evolutionary multi-objective (PBEMO) baselines on the Branin-Currin synthetic function. Only the dense set, from step 1, changes for each bar.

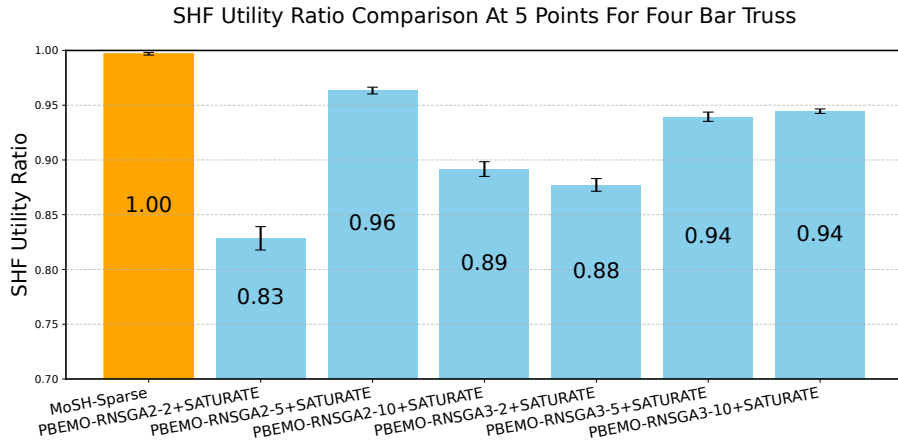


Figure 29: Bar plot illustrating the SHF utility ratio obtained by our method, MoSH-Sparse, compared to the preference-based evolutionary multi-objective (PBEMO) baselines on the Four Bar Truss application. Only the dense set, from step 1, changes for each bar.

F.2 Compute Resources

All experiments were computed on an internal cluster with CPUs, except for the deep learning model selection and LLM personalization experiments which were performed using an NVIDIA A100 GPU with 82GB of VRAM. Final experiments for step 1 generally took within a few hours, except for the deep learning-based experiments, which took longer.

F.3 Performance Evaluation

As mentioned earlier in Section 3, since the DM’s preferences, λ^* are unknown to us, we wish to obtain a set D which is (1) diverse, (2) high-coverage, and is (3) modeled after the DM-defined SHFs. We operationalize those three criteria for soft regions ($A_S = [\alpha_{1,S}, +\infty] \times \dots \times [\alpha_{L,S}, +\infty]$) and hard regions ($A_H = [\alpha_{1,H}, +\infty] \times \dots \times [\alpha_{L,H}, +\infty]$) with the four different metrics described in Section 5.3.

For the soft-hard fill distance, we seek to measure the diversity of sampled points. Malkomes et al. (2021) measures diversity using the notion of fill distance: $\text{FILL}(C, D) = \sup_{x' \in D} \min_{x \in C} \kappa(f(x), f(x'))$ where C is the set of sampled points, D is a full set of precomputed points in the region, and $\kappa(\cdot)$ is the distance metric,

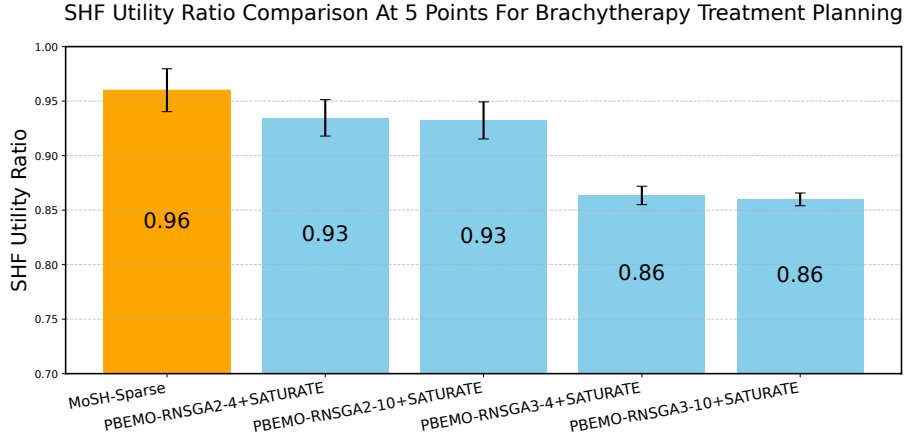


Figure 30: Bar plot illustrating the SHF utility ratio obtained by our method, MoSH-Sparse, compared to the preference-based evolutionary multi-objective (PBEMO) baselines on the Brachytherapy application. Only the dense set, from step 1, changes for each bar.

typically Euclidean distance. We expand this to include the notion of soft and hard regions: $v\text{FILL}_S(C_S, D_S) + (1 - v)\text{FILL}_H(C_H, D_H)$, where $\text{FILL}_S(C_S, D_S)$ and $\text{FILL}_H(C_H, D_H)$ are the fill distances which correspond to the regions defined by the soft bounds and hard bounds, respectively, and C_S, D_S denote the set of points in the soft region, $C \cap A_S, D \cap A_S$, (same for hard region). Intuitively, we wish to obtain a diverse sample set which effectively explores both the soft and hard regions, which a higher weighting towards the soft region. We use the v parameter to control that weighting. For our experiments, we set $v = 0.85$, to indicate a higher weighting of the soft regions.

To calculate the soft-hard fill distance, we first computed a grid search of points on the Pareto frontier, offline, for each experiment. For each experiment, the soft-hard fill distance was then taken with respect to that computed set. As a result, the results for the soft-hard fill distance are somewhat dependent on this offline set of points – however, we ensured that it was kept constant for each set of experiments. In some cases, notably in Figure 22, the soft-hard fill distance does not monotonically decrease. This is due to the heuristic we use in cases where there are no points which have been sampled within either the soft or hard regions, at some iteration. In such a case, we select the worst-case point from set D to represent set C and calculate the metric using that. Once there exists points sampled within the soft or hard regions, we remove that worst-case point and instead use the sampled points. Since the sampled points may result in a soft-hard fill distance value worse than the worst-case point from D , the metric may increase – although the general trend will remain the same. For κ in the soft region distance-weighted score, we used the Euclidean distance.

G ADDITIONAL RELATED WORKS

Multi-Objective Bayesian Optimization. While the literature on Multi-Objective Bayesian Optimization (MOBO) is extensive, we focus our discussion here on several highly relevant frameworks to precisely situate our contributions. For instance, while FERERO also incorporates user guidance, it differs significantly as it assumes differentiable objectives, employs geometric preferences (linear constraints and polyhedral cones) rather than our intuitive soft-hard bounds, and seeks a single optimal solution instead of a robust set (Chen et al., 2024). Other methods are entirely preference-agnostic, aiming to identify the full Pareto front without user guidance. This includes approaches that operate on discrete input sets (Auer et al., 2016) and information-theoretic methods like Predictive Entropy Search (PESMO), which efficiently map the entire front by maximizing uncertainty reduction (Hernández-Lobato et al., 2016). These goals contrast with our approach’s targeted search within a user-defined region. Finally, some works address orthogonal challenges, such as scalability in high-dimensional search spaces (Daulton et al., 2022), a contribution that is complementary to our focus.

Level Set Estimation. In the single-objective setting, Bryan et al. (2005) proposed the straddle heuristic, which was used as part of a GP-based active learning approach for level set estimation (LSE). Although the LSE concept

does not easily extend into a MOO setting, Bryan and Schneider (2008) aims to address that by considering the threshold as part of a composite setting, with scalarized objectives (Iwazaki et al., 2020). In contrast, our proposed MFs are native to the MOO setting and introduce a flexible class of priors DMs may instantiate directly.

H LIMITATIONS

While we offer a novel framework for incorporating general monotonic utility functions directly into multi-objective optimization, we acknowledge certain limitations and areas for future exploration:

Elicitation of Bounds: The effectiveness of MoSH relies on the DM’s ability to specify meaningful soft and hard bounds. We assume in our settings that the DM already possesses knowledge of such priors (e.g. in brachytherapy, diabetes, or vitals monitoring (Deufel et al., 2020; Cairoli et al., 2019; Grosman et al., 2010)). However, in scenarios where DMs are uncertain about these bounds or where such priors do not naturally exist, the initial setup might be challenging. As mentioned in Section E.8, we envision future work where MoSH could be leveraged in a more interactive setting to help iteratively refine these bounds based on feedback.

Form of Monotonic Utility Functions (MFs): We present a specific piecewise-linear form for MFs (Section 5.1) due to its simplicity and interpretability. While we state that other functional classes which are monotonic and bounded would suffice, the impact of choosing different MF forms on the optimization process and the resulting Pareto set has not been extensively studied in this work. The chosen piecewise-linear form for SHFs might not perfectly capture all nuances of a DM’s true underlying utility in all cases.

However, we would like to emphasize what we stated in Section 5.1. While the piecewise-linear form of SHFs is a specific choice, this does not fundamentally limit the framework’s generality. SHFs map directly to standard optimization concepts: hard bounds are equivalent to feasibility constraints on objectives, while soft bounds act as constraints with slack variables. Our experiments consider a fully constrained setting with bounds on every objective for a thorough evaluation. However, this does not imply a loss of flexibility in the SHF model itself. The framework readily accommodates unconstrained objectives, as a user can render any bound non-binding by setting it to a trivial limit, effectively removing the constraint for that objective.

Computational Cost for Very High Dimensions: Although Step 1 of our proposed method aims for efficiency in sampling the PF, there may still be computational cost difficulties when dealing with very high-dimensional problems; *this is inherent to Bayesian optimization*. Similarly, while the worst-case minimax formulation (Equation 1) was converted to an average-case for Step 1 (Equation 2) due to computational feasibility (Appendix E.5), scenarios requiring extremely high-dimensional outputs or fine discretizations of the preference space Λ for the sparsification step (Step 2) could also increase computational demands. *Despite this, we do not view this as a major limitation as our method* has already been evaluated on up to four objectives, which we contend to be appropriate for many real-world MOO problems. For instance, applications of MOO in diabetes typically involve 2-3 objectives (Soares et al., 2024; Mandal and Sutradhar, 2019; El Moutaouakil et al., 2023). Applications of MOO in portfolio optimization typically involve up to 4 objectives (Muteba Mwamba et al., 2025). MOO applications in the field of renewable energy may end up around 4 objectives as well (Shaier et al., 2025). We further assert that all of the real-world applications used throughout this paper (brachytherapy, engineering design, LLM personalization, and deep learning model selection) use the typical number of objectives as well (2-4). Finally, we note that our method is orthogonal to and easily amenable to existing works on scalable Bayesian optimization, if necessary (Daulton et al., 2022).

Addressing these limitations could further enhance the robustness and applicability of our framework.

I BROADER IMPACTS

Potential Positive Impacts: The primary positive impact is enabling more effective and efficient decision support in complex, multi-objective scenarios. In healthcare (e.g., brachytherapy), this could lead to personalized treatment plans better aligned with clinical goals and safety limits. In engineering, it can aid in designing systems that meet performance targets under various constraints. For AI systems like LLMs, it offers a path to better personalization based on user-defined trade-offs (e.g., conciseness vs. informativeness). Overall, our method can enhance human-AI collaboration by translating expert domain knowledge into the optimization process more directly.

Potential Negative Impacts and Considerations: Incorrectly defined utility functions could lead to suboptimal or overlooked solutions. The quality of DM input is crucial. In critical applications, our method should augment, not replace, expert human judgment to avoid potential errors from over-reliance.