


# Preference-Guided Multi-Objective UI Adaptation

Yao Song, Christoph Gebhardt, Yi-Chi Liao, Christian Holz  
Department of Computer Science  
ETH Zürich, Switzerland



**Figure 1:** Our approach enables preference-guided multi-objective UI adaptation. In a coffee shop, an MR user wants to listen to music while watching videos. (a) After the user moves the virtual music player next to their phone, (a-b) our approach determines their preference for the objective of semantic agreement between virtual widgets (i.e., the music player) and physical objects (i.e., the phone) based on changes in objective values and prioritizes it over other terms. (b-c) It then defines a multi-objective optimization problem that conducts a Pareto search based on the set priority rank. Using the music player’s new position as reference, it identifies the closest Pareto-optimal layout candidate by minimizing distance in objective space. (c) The selected layout adjusts all elements to maintain semantic agreement, positioning the video viewer and news page near the iPad, while aligning the messenger, Instagram, and music player closer to the phone.

## ABSTRACT

3D Mixed Reality interfaces have nearly unlimited space for layout placement, making automatic UI adaptation crucial for enhancing the user experience. Such adaptation is often formulated as a multi-objective optimization (MOO) problem, where multiple, potentially conflicting design objectives must be balanced. However, selecting a final layout is challenging since MOO typically yields a set of trade-offs along a Pareto frontier. Prior approaches often required users to manually explore and evaluate these trade-offs, a time-consuming process that disrupts the fluidity of interaction. To eliminate this manual and laborious step, we propose a novel optimization approach that efficiently determines user preferences from a minimal number of UI element adjustments. These determined rankings are translated into priority levels, which then drive our priority-based MOO algorithm. By focusing the search on user-preferred solutions, our method not only identifies UIs that are more aligned with user preferences, but also automatically selects the final design from the Pareto frontier; ultimately, it minimizes user effort while ensuring personalized layouts. Our user study in a Mixed Reality setting demonstrates that our preference-guided approach significantly reduces manual adjustments compared to

traditional methods, including fully manual design and exhaustive Pareto front searches, while maintaining high user satisfaction. We believe this work opens the door for more efficient MOO by seamlessly incorporating user preferences.

## CCS CONCEPTS

- **Mathematics of computing** → **Combinatorial optimization**;
- **Human-centered computing** → **Interactive systems and tools**; **Mixed / augmented reality**.

## KEYWORDS

Multi-objective Optimization, Adaptive User interfaces, Computational Interaction, Mixed Reality

### ACM Reference Format:

Yao Song, Christoph Gebhardt, Yi-Chi Liao, Christian Holz. 2025. Preference-Guided Multi-Objective UI Adaptation. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*, September 28–October 1, 2025, Busan, Republic of Korea. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3746059.3747645>

## 1 INTRODUCTION

User Interfaces (UIs) are the fundamental touchpoint for interaction between users and computing systems, ranging from traditional 2D screens to advanced 3D environments in Mixed Reality (MR) scenarios. The adaptability of UIs in real-time to different user needs and preferences is crucial for enhancing a system’s usability, usefulness, learnability, and user satisfaction, ultimately influencing its success

and acceptance across various applications [28]. To achieve adaptive UIs, previous work has formulated 2D UI layout assignment as a multi-objective optimization (MOO) problem to computationally identify promising layouts. That is, using algorithms to explore various design candidates (i.e., different layouts) to find the ones that lead to optimal objective function values, considering multiple design objective functions (metrics to minimize or maximize).

For example, MOO in 2D GUIs can improve the order of menus by considering the frequency of usage and grouping [2], the sketching interface by optimizing for visual quality and selection efficiency [32], or the layout of control panels governed by ergonomic criterion [10]. Researchers have also leveraged MOO to distribute UIs across multiple devices and users in collaborative environments [29]. For UIs in 3D environments, adaptation is even more critical due to the vast space available for placing UI elements, where the quality of the layout can significantly impact user experience. Moreover, in an open 3D environment, manually adjusting every element is both costly and impractical. In 3D UI layout design, common objective functions include semantic relationships [6], ease of use of interaction methods [7], muscle strain during interaction [15], field of view [11], and more. MOO enables the simultaneous consideration of these objectives, but they inherently involve trade-offs as improving one metric may negatively affect another. The result is typically a Pareto-optimal set of layouts, where each design excels in some objectives while compromising others. While mathematically sound, this approach faces two major limitations that hinder its ability to fully adapt to diverse user needs and preferences.

The first challenge is *the huge effort in manual evaluation on the Pareto-optimal designs*. As discussed, MOO produces a set of trade-off solutions (Pareto-optimal designs), but only one final outcome can be deployed in practice. Determining the final design remains an open question. Several approaches have been proposed. One common approach is to transform a MOO problem into a single-objective problem by weighted sum [6, 7, 11]. This solution leads to information loss since not all Pareto-optimal solutions can be reached when exploring solutions to a single objective. Recent work allows users to *exhaustively test* all the Pareto-optimal designs [15], and identify the most preferred outcome. While this overcomes the limitations of the weighted sum approach, it significantly increases user effort, requiring them to engage with multiple UI layout candidates. This approach is time-consuming and disruptive to a streamlined user experience. The second challenge is *the neglect of user preferences in the MOO process*, which extends in two ways: First, Pareto frontier exploration is typically uniform across objective functions, without prioritizing user preferences. Given limited search time, this can lead to an explored frontier that fails to capture the *true* Pareto frontier, potentially neglecting designs that better align with user needs. Second, when narrowing down to a final design, the process is often disconnected from a concrete estimate of user preferences, leading to subjectively suboptimal outcomes.

To overcome these limitations, we propose a novel optimization approach for UI layout adaptation that *searches the Pareto frontier guided by the user's preference while identifying the final design candidate based on the determined user preference*. Our approach,

combines three features: (a) The system determines the users' preferences for design objectives in a human-in-the-loop manner, directly from their interaction with the UI, (b) our method uses these preferences to guide the dynamic Pareto frontier exploration and (c) automatically derives the most preferred design.

Figure 1 illustrates our approach. Each iteration begins with the user adjusting a UI element in an initial layout. From this input, we determine the user's preference ranking over design objectives by analyzing how the adjustment impacts objective function values. Based on this analysis, objectives are grouped into priority tiers. These rankings are then used by our mixed-Pareto-lexicographic MOO approach [18] to guide the Pareto-front search and identify optimal designs within a fixed search budget. This enables more personalized and preference-aligned results than traditional MOO approaches, which assume fixed Pareto-optimal sets. Finally, the system selects the layout closest to the user's adjustment in objective space. Our framework, illustrated in Figure 2, iterates through this process until the user is satisfied with the layout. It introduces two key distinctions from prior work: (1) it eliminates the need for users to manually explore all Pareto-optimal designs by identifying a candidate based on few user-adjusted elements; (2) it leverages inferred user preferences to guide the Pareto search, enabling personalized, preference-driven UI adaptation.

We demonstrate our approach within a 3D MR context, but our method can equally be applied to 2D UIs. To evaluate its effectiveness and efficiency, we conducted a user study to compare it against two established baselines: manually adjusting the elements of an initial Pareto-optimal layout, and selecting the Pareto-optimal layout closest to the manually adjusted UI element positions. The results showed that our approach significantly reduced the need for manual UI adjustments while maintaining high user satisfaction with the final design. While overall user experience ratings did not show a significant difference from baseline methods, the results indicate promising trends. Notably, the p-values for both metrics were relatively low, with ours consistently trending better than baselines across criteria. These results can be attributed to our approach's ability to conduct a more focused search, as evident by a smaller objective space distance to user-specified references for prioritized objectives. At the same time, the equivalent hypervolume compared to the baseline indicates that this increased focus does not come at the cost of overall design space exploration.

In summary, our contributions are as follows:

- A novel inverse approach that derives the user's preference ranking of objective functions from their UI layout manipulations.
- The first application of mixed Pareto-lexicographic MOO for UI adaptation, leveraging ranked design objectives to model user preferences, effectively guiding a Pareto-frontier search to identify the most preferred UI layouts.
- An empirical study ( $N = 12$ ) demonstrating that, compared to two baseline methods, our approach produced layouts that better matched participants' preferences while requiring less effort to personalize.

## 2 BACKGROUND & RELATED WORK

Our proposed approach builds on two key concepts: Pareto-optimal solutions in MOO problems and the Lexicographic MOO method.

In this section, we provide an overview of the MOO problem, which leads us to Pareto-optimal solutions, and motivates the need for Lexicographic MOO.

## 2.1 Multi-Objective Optimization (MOO)

**2.1.1 MOO problem.** MOO involves simultaneously optimizing multiple objective functions, aiming to identify optimal combinations of decisions across objectives. The general mathematical form of an MOO problem is given by:

$$\begin{aligned} \min \quad & \{f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_k(\mathbf{x})\} \\ \text{subject to} \quad & g_j(\mathbf{x}) \leq 0, \quad j = 1, 2, \dots, m, \\ & h_l(\mathbf{x}) = 0, \quad l = 1, 2, \dots, e \\ & \mathbf{x} \in X \end{aligned}$$

where  $k$  represents the number of objective functions,  $m$  and  $e$  represent the number of inequality and equality constraints respectively. The variable  $\mathbf{x}$  represents the decision variable within the design space  $X$ .

The challenge of a MOO problem is that different objectives may not be fully aligned with each other, and that they could even be in conflict. That is, a design decision,  $\mathbf{x}$ , is good for an objective (e.g.,  $f_1$ ) but may be a poor decision for another design objective (e.g.,  $f_2$ ). Due to such a multi-objective nature, an MOO problem usually can not conclude with one optimal solution. Instead, it yields a set of solutions where improving one objective often comes at the cost of worsening another [24].

These so-called Pareto-optimal solutions are the best trade-offs in this objective space. For more details, we refer the reader to Gunantara's introduction to Pareto optimality and MOO [12].

In the context of UI design and adaptation, the MOO problem is framed within a design space characterized by design factors and a set of objectives that map combinations of these design factors to an objective space [14]. This enables the identification of optimal design factors given the objectives. A common use case is optimizing UI layouts in Virtual Reality. In this context, an optimizer can, for example, determine the positions and dimensions of UI elements, balancing the trade-off between visibility and reachability, as seen in tools like AUIT[11].

**2.1.2 Solvers for MOO problems.** Several types of optimization methods have been proposed to explore and identify the Pareto frontier in the objective space. Evolutionary algorithms are commonly used for the search process because they are gradient-free, allowing them to handle differentiable and non-differentiable objective functions. Within evolutionary algorithms, the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [8] is a common selection to explore the Pareto frontier for its computational efficiency. Below, we briefly provide the working principles of the NSGA-II algorithm:

- (1) **Initialization and Objective Evaluation:** A population of candidate solutions is generated through random sampling, and each solution is evaluated based on multiple objectives.
- (2) **Non-dominated Sorting:** Solutions are ranked into non dominated fronts based on Pareto dominance. Solutions in the first front dominate all others, while subsequent fronts contain increasingly dominated solutions.

- (3) **Crowding Distance and Selection:** Crowding distance is used to maintain diversity within each front. For non-boundary solutions, it is calculated as the normalized difference between the objective values of adjacent solutions. Boundary solutions (those with the smallest and largest values for each objective) are assigned an infinite crowding distance to ensure their preservation [8]. Solutions are selected based on front ranking (prioritizing lower-ranked fronts) and crowding distance (favoring more diverse solutions).
- (4) **Genetic Operations (i.e., Crossover and Mutation):** Selected solutions perform genetic operations (usually, crossover and mutation) to create offspring that explore new areas of the solution space.
- (5) **Elitism and Population Update:** The parent and offspring populations are combined, and the combined set is sorted again into non-dominated fronts. The best solutions are retained for the next generation, ensuring that the elite solutions are carried forward. The process repeats until a stopping criterion is met.

## 2.2 Mixed Pareto-Lexicographic Multi-Objective Optimization (MPL-MOO)

**2.2.1 MPL-MOO problem.** The standard MOO problem assumes equal importance for all objectives, leading to evenly distributed Pareto-frontier searches across the objective space. However, in practice, this may not reflect user preferences. To tackle this limitation, the class of Mixed Pareto-Lexicographic Multi-Objective Optimization problems has been introduced. The generic formulation of the MPL-MOO problem is

$$\begin{aligned} \min \quad & \{f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_k(\mathbf{x})\} \\ \text{subject to} \quad & \wp(f_1, f_2, \dots, f_k) \end{aligned}$$

where  $\wp(\cdot)$  is a general priority distribution among the objectives. In practice, this priority distribution is either realized as priority chains (PC-MPL-MOPs) [17] or as priority levels (PL-MPL-MOPs) [18]. In this paper, we build on the latter. The general mathematical structure of the PL-MPL-MOP is

$$\text{lexmin} \left[ \min \begin{pmatrix} f_1^{(1)}(\mathbf{x}) \\ f_2^{(1)}(\mathbf{x}) \\ \vdots \\ f_{k_1}^{(1)}(\mathbf{x}) \end{pmatrix}, \min \begin{pmatrix} f_1^{(2)}(\mathbf{x}) \\ f_2^{(2)}(\mathbf{x}) \\ \vdots \\ f_{k_2}^{(2)}(\mathbf{x}) \end{pmatrix}, \dots, \min \begin{pmatrix} f_1^{(p)}(\mathbf{x}) \\ f_2^{(p)}(\mathbf{x}) \\ \vdots \\ f_{k_p}^{(p)}(\mathbf{x}) \end{pmatrix} \right]$$

where  $p$  is the number of priority levels and  $f_i^{(j)}$  is the  $i$ th objective in the  $j$ th level of importance. The *lexmin* operator indicates that the minimization of the objectives is ranked lexicographically, meaning that objectives  $f_i^{(1)}$  are infinitely more important than  $f_i^{(2)}$ , which in turn have priority over  $f_i^{(3)}$ , and so on.

The assumption of PL-MPL-MOO is that some groups of objectives have clear precedence over some others in a lexicographic sense. In our work, we leverage this formulation by forming priority groups for objectives based on user preferences determined from their adjustments to an existing MR UI layout. This enables Pareto-frontier searches across the objective space that are guided

by users' preferences leading to designs that are better tailored to individual needs.

**2.2.2 Solvers for PL-MPL-MO problems.** To solve PL-MPL-MO problems, several adaptations of NSGA-II have been proposed for priority-guided Pareto frontier exploration, such as PL-NSGA-II and Lex-NSGA-II [18]. In our work, we employ PL-NSGA-II, which integrates priority levels by performing Pareto dominance checks within each priority level and conducting lexicographic domination checks across priority levels. Here we provide the high-level working principles of PL-NSGA-II:

- (1) **Initialization and Objective Evaluation:** A population of candidate solutions is generated, and each solution is evaluated based on multiple objectives, as in NSGA-II.
- (2) **Non-dominated Sorting:** Solutions are divided into Pareto non-dominated subfronts within each priority level, ensuring a hierarchical ranking based on their priority structure. Within each priority level, the first subfront consists of solutions that dominate all others, while subsequent subfronts contain solutions that are progressively more dominated.
- (3) **Lexicographic Ordering and Crowding Distance:** Solutions are first compared based on higher-priority objectives. If they are equivalent in higher-priority objectives, the algorithm proceeds to evaluate lower-priority objectives using lexicographic ordering (see subsection 2.2.1). To maintain diversity, the crowding distance is computed within each non-dominated subfront at a given priority level.
- (4) **Selection and Genetic Operations:** Solutions are selected for reproduction based on their non-dominated rank, lexicographic preferences, and crowding distance. These selected solutions perform crossover and mutation to generate new offspring.
- (5) **Elitism and Population Update:** The parent and offspring populations are combined, and solutions are re-ranked using non-dominated sorting and lexicographic ordering. The best solutions are retained for the next generation. The process repeats until a termination condition is satisfied.

To summarize, key differences are two-fold: (i) **Lexicographic Ordering:** After non-dominated sorting (i.e., identifying Pareto-optimal solutions) within each priority level, PL-NSGA-II employs a lexicographic hierarchy across priority levels, where higher-priority objectives always take precedence over lower-priority ones. (ii) **Selection Criteria:** PL-NSGA selects solutions not only based on Pareto dominance and crowding distance but also by considering user-defined objective preferences, making it better for cases with objective hierarchies.

## 2.3 Preference-guided MOO in Human-Computer Interaction

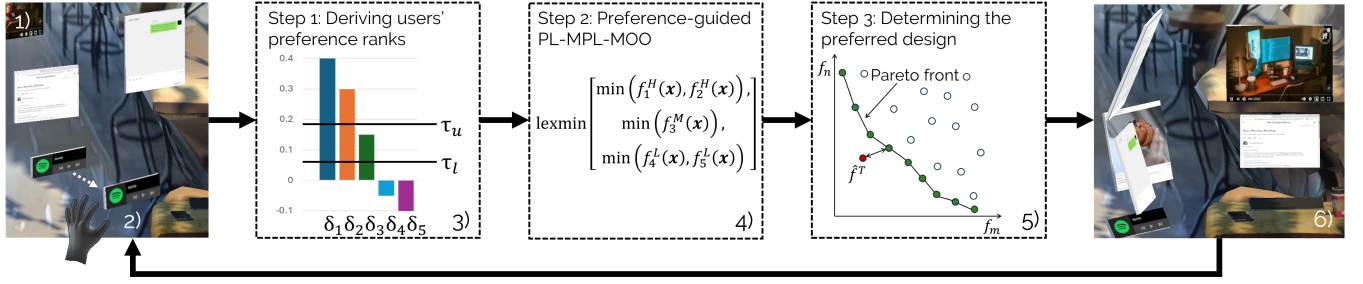
MOO problems typically yield a set of Pareto-optimal solutions, each representing different trade-offs among the objectives. Despite the availability of multiple Pareto-optimal designs, only one solution is usually selected in practice. Selecting a final design from the entire Pareto frontier remains an open research challenge since no single Pareto-optimal design is mathematically superior to the others. Previous research in Human-Computer Interaction has attempted to address this challenge by offloading it onto the user.

One intuitive approach to resolving the challenge of multiple solutions is to transform the MOO problem into a single-objective problem using weighted sums. This method converts the problem into identifying the most suitable weight setting among the design objectives. Numerous studies have explored allowing users and designers to manually adjust these weights to determine the ideal configuration through trial and error. For instance, MenuOptimizer [2] provides an Objective Space Panel where designers can set weights for optimization objectives such as performance, consistency, and similarity. Similarly, OPTIMISM [14] introduces a GUI to help designers and end-users specify objective and heuristic weights for optimization in fabrication tasks. In online 3D UI layout adaptation, tools like SemanticAdapt [6], InteractionAdapt [7], SituationAdapt [20], and AUIT [11] offer control panels within Unity for adjusting weights linked to various objectives. Recently, Liao et al. [21] explored eliciting population-level optimal objective weights based on users' subjective ratings of Pareto-optimal designs. While this method is simple and effective, reducing multiple objectives into a single one inevitably leads to information loss, as certain designs on the Pareto frontier may not be discovered in the weighted-sum solution.

To mitigate this limitation, other approaches have been proposed to explore the Pareto frontier and conclude a final solution based on user preferences. For example, ParetoAdapt [15] presents users with the visualization of the Pareto frontier, allowing them to freely explore and select their preferred Pareto-optimal designs. Instead of focusing on a single Pareto-optimal point, prior research has also explored clustering designs based on similar objective values. For instance, Chan et al. [5] employed multi-objective Bayesian optimization for a 3D input interaction task to find the best trade-offs between accuracy and efficiency. By clustering designs based on objective values, users could choose a design that aligns with their preference for higher accuracy or efficiency. Similar works include clustering haptic device designs based on information transfer and recognition accuracy [22], or clustering other input techniques based on various performance metrics [5], offering users the option to prioritize specific objectives.

While these methods have advanced the selection of a final design in MOO problems based on user preferences, they have three major limitations. First, users must select directly based on objective values, which are often difficult for humans to interpret intuitively. Second, the process involves extensive trial and error, making it time-consuming, tedious, and prone to biases and noise. Finally, the derivation of the Pareto frontier typically occurs independently of user preferences, even though user preferences could provide crucial insights for guiding the Pareto frontier search in more relevant directions. Our approach addresses these limitations by eliminating the need for users to manually select designs based on objective values. Instead, we identify the most preferred design based on users' manipulation of UI elements. This method not only alleviates the challenge of interpreting objective values but also removes the need for repeated trial and error. Moreover, by leveraging the PL-MPL-MOO algorithm, our workflow enables a Pareto frontier search guided by the priority of objectives determined from users' UI interactions, leading to more accurate and customized Pareto frontier exploration.





**Figure 2: Overall pipeline of our approach:** 1) After presenting an initial UI layout, 2) the user can move UI elements according to their preference. 3) Based on these movements, we deduce their preference for optimization objectives based on differences of layouts in the objective space ( $\delta_1, \dots, \delta_5$ , Section 3.1). 4) Via thresholding ( $\tau_u, \tau_l$ ), we then identify priority levels among objectives based on which we formulate a PL-MPL-MOP and perform a preference-guided dynamic Pareto frontier exploration (Section 3.2). 5) Finally, we determine the preferred Pareto-optimal design by using a user's accumulated UI element movements as a reference point  $\hat{f}^T$  and selecting the design on the Pareto frontier with minimal distance (Section 3.3). (6) The adjusted layout is presented to the user, who can then repeat the process until satisfied with the final design.

### 3 APPROACH

Our approach identifies user preferences for objectives and leverages them to conduct a guided exploration of the Pareto frontier, enabling the identification of designs that more closely align with user preferences. Figure 2 illustrates the approach, which consists of three steps: (a) deducing users' preferences for design objectives, (b) preference-guided dynamic exploration of the Pareto frontier, (c) and automatically identifying the most preferred Pareto-optimal design. Below, we provide a detailed explanation of each step.

#### 3.1 Step 1: Deriving users' preference ranks for design objectives

In the first step, we determine user preference ranks over objective terms from their UI element adjustments to an initially generated layout. Specifically, by computing the differences in objective function values before and after user adjustments, we deduce their preferences for design objectives based on the subset of moved elements among the total  $N$  UI elements.

Our core assumption is that a user's preferences for specific design objectives are reflected in the amount of changes made to those objective values. For example, if a user relocates an element in a way that significantly reduces neck strain, we determine that minimizing neck strain has a higher priority than expressed in the current layout, and the optimization process should prioritize this objective. Conversely, if a user moves an element to a position that worsens its reachability, it suggests that reachability has a reduced importance, and should therefore receive less emphasis in the optimization. Building on this, we further assume that the UI layout after user adjustments reflects their preferred state at that moment. With each iteration, user modifications progressively refine the layout, guiding the optimization process toward a final design that best aligns with the user's implicit preferences.

Note that a full UI adaptation may take multiple iterations, with each iteration representing a complete loop in Figure 2. A straightforward solution is seeing each iteration as an independent event; yet, focusing solely on the changes in the current UI movements may introduce noise and miss the information gained from previous

user adjustments. To address this, we adopt a more robust approach by storing all changes to objective values and applying weighted importance to different iterations—recent iterations receive higher weights, while earlier ones receive lower weights. This method enables us to fully utilize prior information while smoothing out noisy human adjustments.

Formally, we define the initial UI positions as  $[x_i, \dots, x_M]$  with the corresponding objective values  $f_j = f([x_i, \dots, x_M])$ . The UI positions of the elements the user has moved are defined as  $[\hat{x}_i, \dots, \hat{x}_M]$  with the corresponding objective values  $\hat{f}_j = f_j([\hat{x}_i, \dots, \hat{x}_M])$ . Here,  $M$  denotes the number of moved UI elements. To get a better estimate of users' preferences, we consider multiple layout adjustments, where  $\hat{f}_j^t$  and  $f_j^t$  are computed for each adjustment  $t \in \{1, 2, \dots, T\}$ , with  $T$  representing the final layout adjustment. Using these objective values, we then compute the difference between objective values  $\delta_j^t$  as

$$\delta_j^t = \frac{\hat{f}_j^t - f_j^t}{f_j^{\max}}, \quad \text{where } j \in \{1, 2, \dots, K\}. \quad (1)$$

Here,  $f_j^{\max}$  represents the maximal value the objective function  $f_j$  can return and  $K$  denotes the total number of objectives. This normalization is a common practice to maintain a consistent objective range, and similar approaches are used in prior human-in-the-loop optimization studies [5, 15]. We then compute the Triangular Moving Average to obtain  $\delta_j$ , representing the difference in values for each objective function in the history of layout adjustments:

$$\delta_j = \frac{\sum_{t=1}^T t \cdot \delta_j^t}{\sum_{t=1}^T t}. \quad (2)$$

The variable  $\delta_j$  reflects whether the user's UI element movements indicate an increase or decrease in preference for the objective function  $f_j$ . By weighting  $\delta_j^t$  based on its position in the adjustment history, we prioritize recent adjustments closer to the final adjustment  $T$ , while gradually reducing the influence of earlier adjustments.

### 3.2 Step 2: Preference-guided dynamic Pareto frontier exploration

In our second step, we incorporate user preferences into the MOO problem by formulating it as PL-MPL-MOP to enable a preference-guided Pareto search. To this end, we use  $\delta_j$  to determine the PL-MPL-MOP's priority levels. Specifically, we define the priority groups as follows:

$$\begin{aligned} H &= \{j \mid \delta_j > \tau_u\} \\ M &= \{j \mid \tau_l \leq \delta_j \leq \tau_u\} \\ L &= \{j \mid \delta_j < \tau_l\} \end{aligned} \quad (3)$$

where  $\tau_l$  and  $\tau_u$  are thresholds for the lower and upper priority level respectively. The final objective function of our problem is then given by

$$\text{lexmin} \left[ \min \left( \begin{pmatrix} f_1^{(H)}(\mathbf{x}) \\ f_2^{(H)}(\mathbf{x}) \\ \vdots \\ f_{k_1}^{(H)}(\mathbf{x}) \end{pmatrix} \right), \min \left( \begin{pmatrix} f_1^{(M)}(\mathbf{x}) \\ f_2^{(M)}(\mathbf{x}) \\ \vdots \\ f_{k_2}^{(M)}(\mathbf{x}) \end{pmatrix} \right), \min \left( \begin{pmatrix} f_1^{(L)}(\mathbf{x}) \\ f_2^{(L)}(\mathbf{x}) \\ \vdots \\ f_{k_p}^{(L)}(\mathbf{x}) \end{pmatrix} \right) \right] \quad (4)$$

where  $H$ ,  $M$ , and  $L$  define the group of objective functions with high, mid, and low priority levels, respectively. The decision to establish three groups allows us to classify objectives into prioritized (important or preferred), neutral (no strong preference), and deprioritized (not preferred) categories. This tri-level classification provides a simple yet robust way to capture user preferences within the objective space. To solve this problem, we use PL-NSGA-II to perform Pareto frontier searches within the objective space, guided by the defined priority groups. This ensures that objectives in  $H$  take precedence over those in  $M$ , and that those in  $M$  take precedence over  $L$ . Since the priority groups are derived from users' UI element adjustments, this approach helps identify Pareto optimal designs that are better aligned with user preferences.

### 3.3 Step 3: Identifying the most preferred Pareto-optimal layout

In the final step of our approach, we leverage the user-adjusted UI elements (element  $i$  to element  $M$ ) to identify the preferred design within the Pareto frontier. Specifically, we map the new positions of the adjusted UI elements  $[\hat{x}_i^T, \dots, \hat{x}_M^T]$  to the objective space to attain a user-defined reference point  $\hat{f}^T = [f_1([\hat{x}_i^T, \dots, \hat{x}_M^T]), \dots, f_K([\hat{x}_i^T, \dots, \hat{x}_M^T])]$ . We then search through the set of Pareto optimal layouts  $X$  to find the point  $f^*$  with minimal distance:

$$\min_{[x_i, \dots, x_M] \in X} d(f, \hat{f}^T), \quad (5)$$

where  $d$  is a function that computes the two-norm distance between the reference point  $\hat{f}^T$  and a Pareto-optimal design  $f = [f_1([x_i, \dots, x_M]), \dots, f_K([x_i, \dots, x_M])]$ , only considering the positions of the  $M$  explicitly adjusted UI elements. Our approach selects the minimal-distance solution  $[x_i^*, \dots, x_M^*]$  and updates the entire layout  $[x_1^*, \dots, x_N^*]$  accordingly.

## 4 EXAMPLE USE CASE: MIXED REALITY UI LAYOUT ADAPTATION

While our approach is applicable to both traditional 2D user interfaces and 3D MR applications, we focus on an MR use case, a common application of adaptive user interfaces [6, 7, 11], to demonstrate its capabilities. Thus, we implemented an MR environment (Figure 1) with a multi-widget UI, similar to those provided by Apple Vision Pro or Meta Quest.

### 4.1 Objective Terms

In the use-case, we aim to utilize our approach to adapt the layout of an MR UI featuring seven virtual widgets. To optimize the layout, we focus on five well-established objectives drawn from previous research: minimizing neck strain [15], minimizing shoulder load [15], balancing visibility and reachability [6], ensuring placement in the field of view and enhancing semantic agreement with environmental objects [6]. All objective functions are defined as cost functions based on the UI elements' positions. Note that our framework is extensible to other virtual widgets and objectives. In the following, we explain the impact of the five objective functions on an MR layout (refer to the respective paper for their equations):

**Minimizing neck strain** [15]: For each virtual widget, determine the angular distance between the user's eye and the widget's position in the x-z plane. Multiply this angular distance by the widget's observation probability. A higher cost signifies that widgets with greater observation probability are positioned above or below the user's eye level, while a lower cost indicates alignment closer to eye level.

**Minimizing shoulder load** [15]: Similar to the neck strain objective, calculate the angular distance between the user's shoulder and each virtual widget's position. Multiply this distance by the widget's interaction probability. A higher cost signifies that frequently interacted widgets are positioned substantially above or below the user's shoulder level, while a lower cost indicates placement closer to the shoulder, facilitating easier interaction.

**Balancing visibility and reachability** [6]: Based on viewing ergonomics [1], this objective is minimized when the UI element is positioned 0.5 meters from the user and maximized at 0 or infinite distance. An exponential interpolation function determines the cost for intermediate distances.

**Ensuring field of view placement** [6]: For each virtual widget, compute the angular difference between its position and the user's gaze direction, then weight it by the widget's observation probability. If this angular difference is less than 5 degrees (i.e., within the foveal region), the cost is zero. A larger angular difference incurs a higher cost (off to the side or behind the user), signifying the widget is farther from the user's gaze, whereas a smaller difference results in a lower cost, indicating proximity to the foveal area (directly in front of the user).

**Semantic agreement** [6]: For each virtual widget, compute its distance to all defined real-world objects around the user and weight it by a semantic association cost derived from WordNet [25]. This approach aligns UI elements with real-world objects based on semantic meaning. A high semantic

association cost and/or a large distance lead to higher overall costs, while weak semantic associations or short distances result in lower costs.

## 4.2 Use-case Scenario

To better illustrate how our approach can support users in a multi-widget MR environment based on the defined objectives, we describe the following scenario: A user is sitting in a coffee shop, wearing an MR headset, and opening various applications on the device (as shown in Figure 1). Physical objects such as a smartphone, an iPad, and a cup of coffee are placed on the table. The visible virtual UI elements include Instagram, a video viewer, a messenger, a paper reader, one news web page, and a music player. Having finished reading the paper, the user now wants to switch to other activities, such as listening to music or watching a video. The user prefers to have entertainment-related applications closer and the paper to the side of the view area. To articulate this preference, the user moves the virtual music player next to the phone (Figure 1a).

Based on the user's adjusted layout (Figure 1b), our approach determines their preference for semantic agreement and adjusts the entire layout accordingly (Figure 1c). The music player is positioned next to the phone, the video viewer is centered, and the paper is placed to the side. Additionally, the video viewer and news pages are moved closer to the iPad, while the messenger, Instagram, and music player are aligned nearer to the phone to achieve higher semantic agreement. The layout also minimizes neck strain while ensuring good reachability and visibility of the UI widgets.

## 4.3 Implementation

We implemented this MR showcase in Unity by building on the implementation of SemanticAdapt [6]. MR environments are scanned real-world 3D rooms where each UI element is depicted as a virtual widget in 3D space (see Figure 3).

Our adaptation approach is implemented in Python 3.9 using the Pymoo package [3]. Specifically, we adapted their NSGA-II implementation to develop a PL-NSGA-II solver. We computed the Pareto frontier with a population size of 100 across 80 generations. We set the thresholds for the priority levels to  $\tau_l = 0.0$  and  $\tau_u = 0.2$ , which were determined via a 3-participant pilot test.

The Unity application communicated with our Python optimizer via websockets.

## 5 EVALUATION

To evaluate if our approach generates MR layouts that better align with users' preferences, we compare our approach with two baseline adaptation methods. Our study aims to address the following two research questions: (1) Can our approach generate layouts that more accurately align with users' preferences? (2) Consequently, can it reduce the number of element movements users need to personalize a layout?

### 5.1 Study Design

We used a within-subject design with two variables: TASK (2 levels: free personalization, targeted personalization) and METHOD TYPE (3 levels: *ParetoSelect*, *Manual*, *Ours*). We also varied the context in which a task was performed by changing the environment (coffee

shop, home office). The order of TASKS was fixed, with free personalization always presented first and targeted personalization second. We applied Latin-square counterbalancing to determine the presentation order of METHOD TYPE and to assign the environment for each TASK. The independent variables are explained in the paragraphs below.

In each condition, we collected usability- and user-experience-related metrics such as the number of elements participants explicitly moved, and participants' perceived task load using the NASA-TLX [13]. After each TASK, participants were further asked to rank the layouts generated by each METHOD TYPE after making their adjustments. Lastly, we asked for participants' overall preference for METHOD TYPE and had them complete the UEQ-S [30].


To compare between MOO-based approaches (i.e., *ParetoSelect* and *Ours*), we measured the average distance between values of participants' moved UI elements in objective space and the nearest solution candidate on the Pareto frontier, using  $d(f, \hat{f}^T)$  as defined in Equation 5. Additionally, we recorded the average objective values. For both metrics, we differentiated between objectives in the highest priority group ( $H$ ) and across all priority groups ( $H, M, L$ ). Lastly, we calculated the hypervolume across all priority groups. Hypervolume is a widely used metric for assessing the results of a Pareto-frontier exploration or MOO. It measures the volume in a multi-dimensional space dominated by the Pareto-optimal points, relative to a given reference point. A higher hypervolume indicates that better overall solutions are identified. For details about the definition and calculation of hypervolume, we refer to Section 2 of Shah and Ghahramani [31].

**Tasks.** The study consisted of two tasks: 1) free personalization and 2) targeted personalization, presented to participants in that order. In free personalization, participants were allowed to freely customize the UI layout according to their preference. In the targeted personalization task, participants were first introduced to the objective functions outlined in Section 4.1 (the explanation of terms used in the experiment is detailed in Appendix A). They were then asked to prioritize the defined terms and make adjustments to the elements, so that movements would align with the objectives of the terms that were more important to them. Both tasks ended when the participants were satisfied with their customized layout.

As our approach models preference in the objective space, it can only estimate preferences related to the objective terms defined in the optimization problem. To evaluate its performance in a controlled setting, we designed the second task to constrain user preferences to these predefined objectives. In contrast, the first task simulates a more realistic scenario where not all user preferences are explicitly modeled, allowing us to assess whether our approach remains effective in such cases. We fixed the task order to always begin with free personalization to avoid priming participants with the predefined objectives of the targeted personalization task.

**Environments.** We used two different 3D real-world scans as MR environments of our study: a coffee shop and a home office. Figure 3 illustrates them, adapted from the SemanticAdapt codebase [6].

In the coffee shop environment, we recreated a leisure scenario where participants were seated at a coffee table with physical items such as a smartphone, an iPad, or a cup of coffee. The virtual UI



**Figure 3: Study environments: (A) coffee shop, (B) coffee shop with initial layout, (C) home office, (D) home office with initial layout.**

elements provided included Instagram, a video viewer, a messenger, a photo viewer, two news web pages, and a music player. Participants could personalize the placement of these UI elements either above the coffee table or in the empty space to their left.

In the home office environment, we designed a literature search scenario where participants were seated at a desk with objects like a phone, headphones, a book, and a tablet. The interface elements included two papers, two web search engines, a messenger, a music player, and Slack. Participants could personalize the placement of interface elements above the desk or to their left or right.

We intentionally selected and designed the scenarios to vary in visual complexity, available space, and social setting. These factors may influence users' mental models and lead to differences in their preferences. The broader range of factors also strengthens our investigation into the approach's generalization. We provided participants with task descriptions tailored to each specific scenario (see Appendix B).

**Methods.** We compared three different METHOD TYPES: *Manual*, *ParetoSelect*, and *Ours*. To ensure a fair comparison, all conditions started with the same initial layout randomly sampled from a Pareto frontier generated by NSGA-II. All conditions used the same set of objective functions (defined in Section 4.1).

- **Manual:** Participants could move all elements of the initial design to create their preferred layout. No further layout optimization was conducted. Participants manually adjusted a layout until they were satisfied with it.
- **ParetoSelect:** It allowed participants to also move elements as they wished, but instead of directly using these positions, we applied Step 3 of our approach to determine the Pareto-optimal layout closest to their adjustments (Section 3.3). Participants could choose to move between one and three elements, after which the next Pareto-optimal layout was selected and displayed. This cycle could be repeated as many times as needed until they were satisfied with the layout. We designed *ParetoSelect* to represent the state of the art in Pareto-optimization-based UI adaptation [15]. This approach suggests selecting a final design candidate from the Pareto frontier of an optimized single UI element by displaying

different positions as "ghosts" for the user to choose from. However, this method does not scale well to full UI layouts with multiple elements. To address this, we implemented *ParetoSelect* as a straightforward approach to integrate candidate selection of full UI layouts into non-prioritized MOO problems. By using the same candidate selection method as in our approach, we ensure that any differences arise solely from the contrast between prioritized and non-prioritized MOO problems.

- **Ours:** It represents our full approach described in Section 3. Based on participants' adjustments, their preference for objectives was determined, which was used to guide the prioritized Pareto frontier exploration. Then, the final optimal layout was identified. Again, participants could move between one and three elements, after which PL-NSGA-II recomputed the Pareto frontier. This process could be repeated as many times as needed until they were satisfied with the layout.

## 5.2 Procedure

Prior to the study, participants completed a consent form and a pre-questionnaire to gather demographic data and assess their previous VR/MR experience. They were then given a tutorial on using hand gestures to interact with UI elements and had the opportunity to explore interactions and elements in a sample scene. Participants proceeded to perform both free and targeted personalization tasks across the literature search and leisure environments, using each of the three adaptation methods. After completing each condition, they filled out a NASA-TLX questionnaire and ranked the layouts generated by each method. At the end of the session, participants completed the UEQ-S and stated their method preference. Each session lasted approximately 30 minutes.

## 5.3 Apparatus

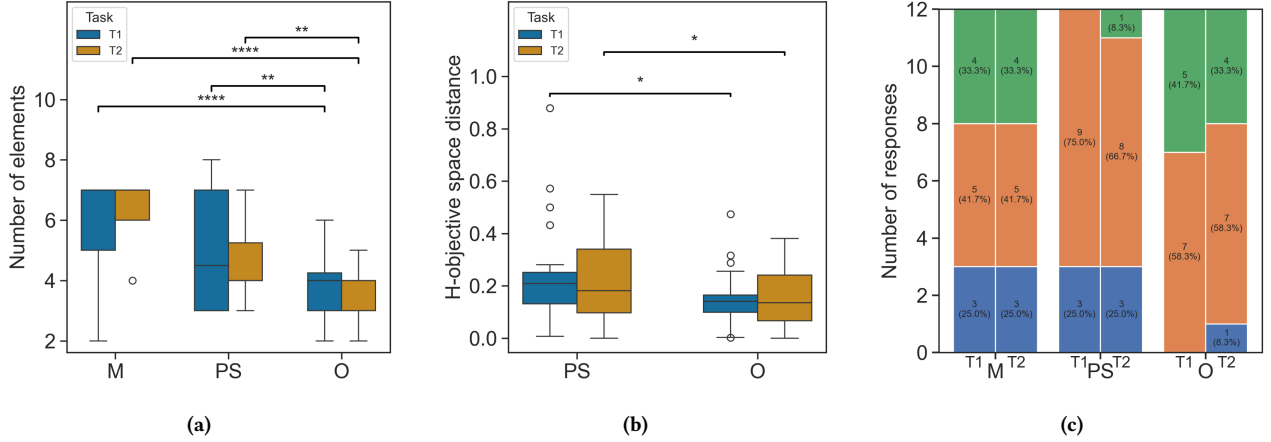
We extended our Unity application to include the baseline conditions and integrated the NSGA-II algorithm into our Python codebase. Participants completed the study while seated, using an Oculus Quest 2 headset tethered to a computer via cable. They used their hands to interact with UI elements via the Quest's hand tracking.

## 5.4 Participants

We recruited 12 participants (2 female, 10 male), aged 22–27 (Mean = 25, SD = 1.7) from a local university. Participants disclosed their familiarity with VR/MR technology. Three participants considered themselves familiar but not experts, while one indicated no prior experience. The remaining participants reported occasional usage, possessing basic operational knowledge of VR devices.

## 6 RESULTS

In the following, we present the results of our study. For significance testing, we performed an ANOVA if data was normally distributed (Shapiro-Wilk  $p > .05$ ) and exhibited equal variances between groups (Levene's  $p > .05$ ). Pairwise comparisons were performed using t-tests with Bonferroni-adjusted p-values. If either assumption was violated, we assessed differences with the non-parametric Aligned Rank Transform (ART) ANOVA. Post-hoc comparisons were then performed using the ART-C algorithm and



**Figure 4: Results from our study of the three methods per task: *Ours* (O), *ParetoSelect* (PS), and *Manual* (M). The metrics displayed are (a) the number of moved elements, (b) H-objective space distance, and (c) layout preference. (a) and (b) are presented as boxplots, with significance indicated at the following levels: \* $p < .05$ , \*\* $p < .01$ , \*\*\* $p < .001$ , \*\*\*\* $p < .0001$ . (c) Illustrates the distribution of preference ranks among the three methods, categorized into Rank 1, Rank 2, Rank 3.**

Bonferroni corrections. If not stated differently, METHOD TYPE and TASK were considered as within-subject factor and participants as a random factor. Since statistical tests showed no significant differences between TASKS across any metrics, we focus solely on effects related to METHOD TYPE in the remainder of this section. Results with significant differences are shown in Figure 4.

### 6.1 Usability metrics

We analyzed the collected metrics related to usability and user experience. A main effect of METHOD TYPE on the *number of moved elements* across TASKS [ $F_{2,66} = 18.97$ ,  $p < .0001$ ,  $\eta_g^2 = .38$ ] was found. Participants manually moved less UI elements with *Ours* compared to both *ParetoSelect* ( $p = .01$ ,  $r = .35$ ) and *Manual* ( $p < .0001$ ,  $r = .61$ ). No other significant differences were identified.

We evaluated the subjective feedback provided by the participants in-between conditions. A significant effect of METHOD TYPE on *layout preference* across TASKS was found [ $F_{2,66} = 3.98$ ,  $p = .02$ ,  $\eta_g^2 = .11$ ]. Post-hoc tests have shown that participants ranked UI layouts generated by *Ours* higher than when generated by *ParetoSelect* ( $p = .02$ ,  $r = .33$ ). Other differences were non-significant. Additionally, no significant differences were identified for NASA-TLX.

To statistically analyze the subjective feedback from the post-questionnaire, conditions were treated as a within-subject factor, while participants were treated as a random factor. No significant differences were found for *method preference* (METHOD TYPE  $p = .09$ ), nor for the dimensions of the UEQ: *pragmatic quality* (METHOD TYPE  $p = .07$ ), *hedonic quality* (METHOD TYPE  $p = .11$ ), and *overall quality* (METHOD TYPE  $p = .07$ ).

### 6.2 Optimization metrics

We evaluated the metrics computed from the optimization problems themselves to further compare *ParetoSelect* and *Ours*. No significant differences were found in the *hypervolume* and the *average costs*

of objective functions within all groups ( $H$ ,  $M$ ,  $L$ ) as well as within priority group  $H$ .

Since the *distance in objective space* of moved UI elements to the nearest candidate on the Pareto front depends on UI element movement, we must ensure these movements are the same for both *ParetoSelect* and *Ours* for a fair comparison. Therefore, for each condition, we recomputed the Pareto front using the movements from participants in the other condition. To track which condition the movements came from and assess a potential effect, we introduced the factor ADJUSTMENT CONDITION. For statistical analysis, METHOD TYPE (the method that generates the Pareto-front designs) was treated as a within-subject factor, and ADJUSTMENT CONDITION (the condition where the user adjustment was originally from) as a between-subject factor.


No differences between conditions were identified for the *distance in objective space* of all objective functions. However, we found a main effect of METHOD TYPE on the *objective space distance* of high priority objectives ( $H$ ) across ADJUSTMENT CONDITIONS [ $F_{1,92} = 4.26$ ,  $p = .04$ ,  $\eta_g^2 = .04$ ]. No other significant differences were identified.

The advantage of our method in achieving a lower distance in the objective space, becomes evident when analyzing the temporal progression of the Pareto front across iterations of our algorithm. Figure 5 illustrates this progression for two participants in our study, demonstrating that the best solution candidates of our PL-NSGA-II algorithm converge closer to the reference point derived from user-moved UI elements. Moreover, these candidates consistently dominate those from the Pareto front generated by NSGA-II.

## 7 DISCUSSION

The goal of our approach is to improve multi-objective UI optimization by more effectively incorporating users' individual preferences in the adaptation of UI layouts. Our generated layouts should require fewer adjustments from users when customizing them to their needs, while increasing overall satisfaction with the layout.





**Figure 5: Pareto fronts of two study participants evolving through successive iterations of our algorithm. The ● represents the initial NSGA-II Pareto Front, the ○ represents the PL-NSGA-II Pareto Front at iteration 1, the × represents the PL-NSGA-II Pareto Front at iteration 2, and the + represents the reference point attained from all user moved UI elements. The larger circles (○, ○, ○) surrounding the Pareto front candidates represent the best candidates from the algorithm iterations with corresponding colors.**

Our results indicate that these objectives were achieved. With *Ours*, participants required fewer UI element movements to customize layouts to their preferences compared to both *ParetoSelect* and *Manual*. Additionally, participants expressed a clear preference for layouts generated by *Ours* over those by *ParetoSelect*. No significant difference between *Ours* and *Manual*, where users could fully control the layout, suggests that *Ours* maintains the flexibility of manual customization while reducing its effort (less adjustments). Note that statistical tests showed no differences between TASKS across any metrics. These results suggest that our method’s advantages hold for cases when objectives are known (Task 2) to more realistic scenarios where users are unfamiliar with them (Task 1).

The higher preference for layouts generated by our method did not yet translate into a clear overall preference or higher user experience compared to the baselines. This is not surprising, as UX metrics such as NASA-TLX and UEQ-S are known to be insensitive to short-term UI adaptation tasks [19, 21]. Nevertheless, the results of our UX-related measures indicate promising trends. Notably, the p-values for the related metrics were relatively low ( $p = .07 - .11$ ), with *Ours* consistently trending better than baselines across criteria. This suggests that with a larger sample size, the results have strong potential to achieve statistical significance.

We further explored whether the performance in these usability-related metrics can be attributed to the underlying algorithm’s functionality. Results showed that *Ours* achieved a lower distance in objective space between moved UI elements and the nearest candidate on the Pareto front for the group of objectives with highest priority, compared to *ParetoSelect*. This shows that *Ours* resulted in Pareto frontiers that are more focused and aligned with the users’ prioritized objectives than other baselines. Furthermore, we highlight that while *Ours* generates a more focused, user-aligned

search (reflected by the shorter distances for high-priority objectives), it maintains overall design space exploration, as indicated by the equivalent hypervolumes across METHOD TYPE. These findings indicate that our method directs the search more effectively toward regions of the Pareto frontier that better reflect user intent, all within the same search budget.

A known limitation of UI optimization is that some user preferences may lie outside the defined objectives. As with state-of-the-art approaches, we assume that the selected objectives capture the majority of user needs. The validity of this assumption is supported by the positive results in our free personalization task, where our method still reduced widget adjustments and led to a preference for the generated layout, highlighting the relevance of the chosen objective set. To mitigate potential mismatches between modeled objectives and user preferences, our approach (a) downranks irrelevant objectives via lexicographic inference and (b) allows manual adjustments beyond the modeled space.

## 8 LIMITATIONS & FUTURE WORK

*Hard-constraint modeling of moved UI elements.* In our current approach, user-adjusted UI elements are leveraged to estimate objective preferences and subsequently identify the optimal UI layout on the Pareto front. As a result, elements manually moved by the user may still be adjusted by the algorithm in later iterations. This is an intentional design choice, as it allows the algorithm to optimize all elements in an unconstrained space with increasing information on user preference. However, future work could explore the integration of user-adjusted UI elements as hard constraints within the optimization process and investigate whether this enhances the perceived usefulness of our approach.



*Handling uncertainty in user preferences.* Our approach infers user preferences in a deterministic manner, assuming that preferences can be directly derived from changes in objective values. However, in real-world scenarios, user adjustments may be inconsistent, exploratory, or even accidental, introducing noise into the preference estimation process. While we have applied filtering techniques to mitigate this issue, more advanced approaches should be explored in future work. A promising direction is to move beyond direct value-based inference and instead adopt probabilistic methods, such as Bayesian preference modeling or probabilistic ranking, which can incorporate confidence levels in inferred preferences. Additionally, stochastic optimization techniques could be leveraged to dynamically refine preference estimations while accounting for uncertainty and inconsistency in user interactions. These approaches would enhance robustness, making UI adaptation more resilient to noisy or ambiguous user feedback.

*Potential of dynamic Pareto frontier exploration.* In this paper, we employ PL-NSGA-II, which ranks design objectives to guide the dynamic exploration of the Pareto frontier. Our approach outperformed traditional NSGA, whose Pareto front exploration lacks guidance from additional information. This demonstrates the potential of dynamically informed Pareto-optimal exploration. Future work should investigate other relevant approaches for such dynamic exploration. One straightforward yet under-explored approach in HCI is the adaptive weighted-sum method [16], which actively learns user preferences for different objective weightings and converges toward a preference-guided optimum. A more complex direction is Trust-Region Optimization [4, 26], where the search process is focused within a “trust region” or around a “trust point” in the design space or objective space, prioritizing exploration in this local area. Trust-region methods have been adapted for evolutionary algorithms, e.g., Reference Point-Based NSGA-II [9], which could be naturally extended to our case. The user adjustment could provide the trust region, enabling a more focused search around nearby designs. Another promising direction is incorporating surrogate models (e.g., Gaussian Processes) to learn user preferences based on their interaction, and guide exploration dynamically [27]. Future research is encouraged to explore these dynamic Pareto-frontier exploration techniques and assess their potential benefits and challenges in various design contexts.

*Beyond the preferences in objective space.* Our approach assumes that users have implicit preferences over objectives, yet, in practice, their preferences may also exist in the design space. For example, a user might prefer certain UI elements to remain fixed in specific locations, even if these preferences do not align with the predefined design objectives. Preference optimization is a parallel topic focused on learning and optimizing user preferences within the design space [33]. Exploring a hybrid approach that infers whether the user’s preference lies within the objective space (e.g., enhancing overall reachability) or the design space (e.g., positioning certain elements in specific locations) would be a valuable direction for future research. By understanding these distinctions, the optimizer can adapt its exploration strategy accordingly, leading to more personalized and effective design solutions.

*Open challenges in re-optimization.* Consistent with prior UI optimization work [2, 6, 7], we assume a one-time adaptation based on stable user preferences where users can restart the process if their preferences change. In practice, nonetheless, user preferences may evolve over time, which calls for the capability to dynamically re-adapt. This introduces several open research questions. For instance, how can we accurately detect when a user’s preferences have significantly shifted? How frequently should the system re-adapt? How can the system re-adapt while balancing the preservation of the current design to avoid drastic changes? Future research should consider addressing these challenges to create systems capable of long-term, dynamic interaction and evolving user satisfaction. Further, future works should extend the optimization to a continual setting [23]; that is, the system keeps accumulating past experience and becomes more informed to make better decisions over time.

*Transferring learned preferences across UIs and users.* Currently, our approach determines user preferences from scratch for each UI adaptation, requiring users to manually adjust layouts even when their preferences remain consistent across similar interfaces. This redundancy increases user effort and slows down optimization. Future work should explore meta-learning or transfer learning to carry over learned preferences from one UI to another, particularly when interfaces share structural similarities. By leveraging past interactions, the system could provide more informed initial suggestions, reducing the need for repeated manual refinements.

Beyond individual users, preference learning could also be extended across users interacting with the same or similar UIs. If multiple users have optimized an interface, new users could benefit from these collective insights, receiving pre-optimized layouts that align with commonly preferred patterns. Clustering techniques could help group users with similar preferences, enabling a personalized yet scalable adaptation strategy. Future research should explore how to balance shared knowledge transfer with individual customization, ensuring UI layouts remain flexible while benefiting from prior optimizations.

*Objective function parameters.* Parameters like arm’s length, ideal viewing distance from an interface, and levels of semantic association and avoidance impact the position of minimum and maximum points in the objective functions. Currently, population averages are used to set these values, but fixed parameters cannot account for individual variations and needs. To further tailor MR layouts to individual users’ preferences and needs, future work should explore the feasibility of automatically inferring these parameters from users’ interaction behavior or other sources.

*Scalability.* In this work, we demonstrated that our approach performs effectively with five objectives and seven widgets. However, real-world use cases may involve a greater number of objectives and widgets, potentially leading to longer optimization runtimes. To address this, we employ NSGA-II [8, 15], a widely used multi-objective optimizer known for its linear runtime scaling with respect to the number of input elements and objectives. Combined with its support for parallelization and a fixed search budget, it enables scalable and predictable real-time performance.

## 9 CONCLUSION

In this paper, we introduced a novel preference-guided multi objective optimization approach for UI layout adaptation, addressing key limitations of traditional MOO methods. By utilizing mixed Pareto-lexicographic optimization, our approach dynamically explores the Pareto frontier based on derived user preference rankings over objective terms, ensuring that the resulting designs are better aligned with the user's needs. This enables more personalized UI adaptation in both 2D interfaces and 3D Mixed Reality environments.

Through a user study in an MR context, we demonstrated that our method significantly reduces the effort required for manual UI adjustments while maintaining high user satisfaction with the final layout. These findings highlight the potential of preference-guided optimization in improving the effectiveness of adaptive UIs.

Overall, our contributions provide a strong foundation towards the true personalization of adaptive UIs. The application of mixed Pareto-lexicographic MOO presents a promising direction for future UI adaptation frameworks, for instance, enabling hybrid approaches that combine it with other methods to directly optimize for user preferences within the design space.

## ACKNOWLEDGMENTS

Yi-Chi Liao was supported by the ETH Zurich Postdoctoral Fellowship Programme.

## REFERENCES

- [1] Dennis R Ankrum and Kristie J Nemeth. 1995. Posture, comfort, and monitor placement. *Ergonomics in design* 3, 2 (1995), 7–9. <https://doi.org/10.1177/106480469500300203>
- [2] Gilles Bailly, Antti Oulasvirta, Timo Kötzing, and Sabrina Hoppe. 2013. Menuoptimizer: Interactive optimization of menu systems. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. 331–342. <https://doi.org/10.1145/2501988.2502024>
- [3] Julian Blank and Kalyanmoy Deb. 2020. Pymoo: Multi-objective optimization in python. *Ieee access* 8 (2020), 89497–89509.
- [4] Gabriel A Carrizo, Pablo A Lotito, and Maria C Maciel. 2016. Trust region globalization strategy for the nonconvex unconstrained multiobjective optimization problem. *Mathematical Programming* 159 (2016), 339–369.
- [5] Liwei Chan, Yi-Chi Liao, George B Mo, John J Dudley, Chun-Lien Cheng, Per Ola Kristensson, and Antti Oulasvirta. 2022. Investigating positive and negative qualities of human-in-the-loop optimization for designing interaction techniques. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–14. <https://doi.org/10.1145/3491102.3501850>
- [6] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. 2021. Semanticadapt: Optimization-based adaptation of mixed reality layouts leveraging virtual-physical semantic connections. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 282–297. <https://doi.org/10.1145/3472749.3474750>
- [7] Yi Fei Cheng, Christoph Gebhardt, and Christian Holz. 2023. InteractionAdapt: Interaction-driven Workspace Adaptation for Situated Virtual Reality Environments. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–14. <https://doi.org/10.1145/3586183.3606717>
- [8] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.
- [9] Kalyanmoy Deb and Jayavelmurugan Sundar. 2006. Reference point based multi-objective optimization using evolutionary algorithms. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*. 635–642.
- [10] Jose Antonio Diego-Mas, Diana Garzon-Leal, Rocio Poveda-Bautista, and Jorge Alcaide-Marzal. 2019. User-interfaces layout optimization using eye-tracking, mouse movements and genetic algorithms. *Applied Ergonomics* 78 (2019), 197–209. <https://doi.org/10.1016/j.apergo.2019.03.004>
- [11] João Marcelo Evangelista Belo, Mathias N Lystbæk, Anna Maria Feit, Ken Pfeuffer, Peter Kán, Antti Oulasvirta, and Kaj Grønbaek. 2022. Auit—the adaptive user interfaces toolkit for designing xr applications. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–16. <https://doi.org/10.1145/3526113.3545651>
- [12] Nyoman Gunantara. 2018. A review of multi-objective optimization: Methods and its applications. *Cogent Engineering* 5, 1 (2018), 1502242.
- [13] Sandra G Hart. 1986. NASA task load index (TLX). *NASA Ames Research Center* (1986). <https://ntrs.nasa.gov/citations/20000021487>
- [14] Megan Hofmann, Nayha Auradkar, Jessica Birchfield, Jerry Cao, Autumn G Hughes, Gene SH Kim, Shriya Kurpad, Kathryn J Lum, Kelly Mack, Anisha Nilakantan, et al. 2023. OPTIMISM: Enabling Collaborative Implementation of Domain Specific Metaheuristic Optimization. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19. <https://doi.org/10.1145/3544548.3580904>
- [15] Christoph Albert Johns, João Marcelo Evangelista Belo, Anna Maria Feit, Clemens Nylandsted Klokmoose, and Ken Pfeuffer. 2023. Towards Flexible and Robust User Interface Adaptations With Multiple Objectives. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–17. <https://doi.org/10.1145/3586183.3606799>
- [16] Il Yong Kim and Olivier L de Weck. 2006. Adaptive weighted sum method for multiobjective optimization: a new method for Pareto front generation. *Structural and multidisciplinary optimization* 31, 2 (2006), 105–116.
- [17] Leonardo Lai, Lorenzo Fiaschi, and Marco Cococcioni. 2020. Solving mixed Pareto-Lexicographic multi-objective optimization problems: The case of priority chains. *Swarm and Evolutionary Computation* 55 (2020), 100687.
- [18] Leonardo Lai, Lorenzo Fiaschi, Marco Cococcioni, and Kalyanmoy Deb. 2021. Solving mixed pareto-lexicographic multiobjective optimization problems: the case of priority levels. *IEEE Transactions on Evolutionary Computation* 25, 5 (2021), 971–985.
- [19] Byungjoo Lee, Mathieu Nancel, Sunjun Kim, and Antti Oulasvirta. 2020. AutoGain: Gain function adaptation with submovement efficiency optimization. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [20] Zhipeng Li, Christoph Gebhardt, Yves Inglin, Nicolas Steck, Paul Streli, and Christian Holz. 2024. Situationadapt: Contextual ui optimization in mixed reality with situation awareness via llm reasoning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–13.
- [21] Yi-Chi Liao, Ruta Desai, Alec M Pierce, Krista E. Taylor, Hrvoje Benko, Tanya R. Jonker, and Aakar Gupta. 2024. A Meta-Bayesian Approach for Rapid Online Parametric Optimization for Wrist-based Interactions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 410, 38 pages. <https://doi.org/10.1145/3613904.3642071>
- [22] Yi-Chi Liao, John J Dudley, George B Mo, Chun-Lien Cheng, Liwei Chan, Antti Oulasvirta, and Per Ola Kristensson. 2023. Interaction Design With Multi-objective Bayesian Optimization. *IEEE Pervasive Computing* 22, 1 (2023), 29–38. <https://doi.org/10.1109/MPRV.2022.3230597>
- [23] Yi-Chi Liao, Paul Streli, Zhipeng Li, Christoph Gebhardt, and Christian Holz. 2025. Continual Human-in-the-Loop Optimization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 795, 26 pages. <https://doi.org/10.1145/3706598.3713603>
- [24] R Timothy Marler and Jasbir S Arora. 2004. Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization* 26 (2004), 369–395. <https://doi.org/10.1007/s00158-003-0368-6>
- [25] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41. <https://doi.org/10.1145/219717.219748>
- [26] A Mohammadi and Ana Luisa Custódio. 2024. A trust-region approach for computing Pareto fronts in multiobjective optimization. *Computational Optimization and Applications* 87, 1 (2024), 149–179.
- [27] Hee-Seung Moon, Antti Oulasvirta, and Byungjoo Lee. 2023. Amortized inference with user simulations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [28] Antti Oulasvirta, Niraj Ramesh Dayama, Morteza Shiripour, Maximilian John, and Andreas Karrenbauer. 2020. Combinatorial optimization of graphical user interface designs. *Proc. IEEE* 108, 3 (2020), 434–464. <https://doi.org/10.1109/JPROC.2020.2969687>
- [29] Seonwook Park, Christoph Gebhardt, Roman Rädle, Anna Maria Feit, Hana Vrzakova, Niraj Ramesh Dayama, Hui-Shyong Yeo, Clemens N Klokmoose, Aaron Quigley, Antti Oulasvirta, et al. 2018. Adam: Adapting multi-user interfaces for collaborative environments in real-time. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–14.
- [30] Martin Schrepp, Jessica Kollmorgen, and Jörg Thomaschewski. 2023. A Comparison of SUS, UMUX-LITE, and UEQ-S. *Journal of User Experience* 18, 2 (2023).
- [31] Amar Shah and Zoubin Ghahramani. 2016. Pareto frontier learning with expensive correlated objectives. In *International conference on machine learning*. PMLR, 1919–1927.
- [32] Kashyap Todi, Daryl Weir, and Antti Oulasvirta. 2016. Sketchplore: Sketch and explore with a layout optimiser. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. 543–555.
- [33] Kenta Yamamoto, Yuki Koyama, and Yoichi Ochiai. 2022. Photographic lighting design with photographer-in-the-loop bayesian optimization. In *Proceedings of the 35th annual ACM symposium on user interface software and technology*. 1–11.

## A TASK EXPLANATIONS

We introduced participants of our study to the two tasks, free personalization and targeted personalization. In addition, we gave them instructions on how to adjust UI elements manually. Below, we provide the instructions used in our study:

### Free personalization:

You are free to adjust UI widgets' position until you are satisfied with the overall layout.

### Targeted personalization:

When adjusting the position of virtual UI elements and evaluating the updated layout, you should focus only on the following factors:

- **Shoulder Load** – UI should be easily reachable without lifting your arms too high.
- **Neck Strain** – UI should be at a comfortable eye level to avoid neck strain. (Neither too high nor too low)
- **Distance** – UI should be within arm's reach—neither too far nor too close.
- **Central Vision** – UI should be positioned near the center of your field of view.
- **Semantic Agreement** – UI elements and environmental objects can possess a semantic association, meaning that certain UI components are more semantically connected to specific environmental (physical) objects. For example, the UI of a video viewer is more semantically aligned with a laptop than with a book. This indicates that positioning the video viewer UI closer to the laptop rather than the book enhances semantic agreement.

### Manual UI element adjustments:

*Manual*: Review all UI widgets and adjust any that you're not satisfied with until you are satisfied with the layout.

*ParetoSelect, Ours*: Adjust 1-3 UI elements as an example. The tool will automatically refine the layout accordingly. If you're satisfied with the updated layout, you can stop; otherwise, continue refining elements.

## B SCENARIO EXPLANATIONS

We introduced participants of our study to the two scenarios, coffee shop and home office, with the following text:

You will complete two tasks that involve customizing the positions of virtual UI elements. These tasks take place in two different scenarios:

**Relaxing in a coffee shop** – Imagine you're seated at a table in a coffee shop. UI elements can be arranged in the space directly in front of you and to your left side. You arrange UI widgets for an enjoyable browsing experience using apps like Instagram, a video viewer, a messenger, news websites, and a music player. Your goal is to adjust the position of these widgets to create the most comfortable and convenient UI layout for an enjoyable browsing experience.

**Working from home** – Imagine you're working from home, sitting at a desktop. UI elements can be positioned in the area in front of you, as well as to your left and right sides. You'll organize UI elements for efficient literature research using research papers, search engines, Slack, a messaging app, and a music player. Your goal is to arrange these widgets to create the most comfortable and convenient UI layout for your literature search.