

Computational Adaptation of Extended Reality Interfaces Through Interaction Simulation

Kashyap Todi
kashyap.todi@gmail.com
Meta Reality Labs
USA

Ben Lafreniere
benlafreniere@fb.com
Meta Reality Labs
Canada

Tanya Jonker
tanya.jonker@fb.com
Meta Reality Labs
USA

ABSTRACT

Adaptive and intelligent user interfaces have been proposed as a critical component of a successful extended reality (XR) system [15]. In particular, a predictive system can make inferences about a user and provide them with task-relevant recommendations or adaptations. However, we believe such adaptive interfaces should carefully consider the overall *cost* of interactions to better address uncertainty of predictions. In this position paper, we discuss a computational approach to adapt XR interfaces, with the goal of improving user experience and performance. Our novel model, applied to menu selection tasks, simulates user interactions by considering both cognitive and motor costs. In contrast to greedy algorithms that adapt based on predictions alone, our model holistically accounts for costs and benefits of adaptations towards adapting the interface and providing optimal recommendations to the user.

1 INTRODUCTION

Extended reality (XR) is a growing area of post-desktop computing, spanning virtual reality (VR), augmented reality (AR), and mixed reality (MR). In VR, users are immersed in complex 360° environments where digital content can be freely accessed. In AR and MR, users can seamlessly interact with both the real world and virtual content, blurring the boundaries between them. This is radically different from current interaction paradigms where digital interfaces occupy their own dedicated physical spaces (screens) and offer users optimized and performant input methods to interact with them. It is still unclear what new interaction capabilities and interfaces can best support the user and enable them to better interact with content here. Given recent advancements in computational interaction [11, 14], we see great potential for applying or developing computational approaches to generate or adapt XR user interfaces (UIs). In contrast to static UIs, these approaches could dynamically adapt the UI to a user’s context and activity, thereby improving interactions.

In this position paper, we introduce a model that can simulate a user’s interactions with an XR interface, with the goal of selecting optimal UI adaptations. We assume that the system has a predictive model that provides a probability distribution over all possible actions in the interface. In a menu interface, for instance, the system could predict the selection probability for each of the menu items.

Prior work has developed methods and models for predicting users’ intentions given some prior information such as the user’s interaction history and their context (e.g. [2, 6, 10, 19]). Given a probability distribution over available actions, a predictive system can then generate the UI. For example, it could generate a simplified menu that prioritizes items with high probability (e.g. [7]), preferentially reveal items [4], gradually reorganize items [18], or augment the standard menu interface with recommendations (e.g. [17]).

A key aspect of our proposed model is that it consider various *costs* during interaction. Recent work has studied the cognitive costs and recovery time associated with recognition errors [9], and planning-based approaches have been developed to estimate long-term costs and benefits of adaptations [18]. Here, we study cognitive and interaction costs in the context of a predictive system that makes adaptive recommendations based on contextual predictions about the user’s intents¹. We argue that a ‘greedy’ recommender approach – one that exclusively prioritizes high probability actions – might be sub-optimal when used directly because it does not take into account the interface design and potential adverse interaction costs of inaccurate predictions. We develop a simulation model of a user’s interactions with a hierarchical menu. The model takes into account both *benefits* of a suitable adaptation and *costs* of inaccurate predictions. We consider three key time-based costs of interactions – inspection cost, selection cost, and correction cost – to compute the *utility* of any given adaptation. In contrast to a greedy baseline that always selects the highest-probability action, our model attempts to minimize the overall selection time. The model is instantiated in a menu-based XR system, enabling it to adaptively select a utility-maximizing starting point in the menu hierarchy. We demonstrate the model’s predictions in a system with illustrative scenarios.

2 OVERVIEW

2.1 Background: Adaptive User Interfaces

In this paper, we discuss XR systems that can adapt their interface to a user. By observing the user’s context, such systems can make probabilistic predictions about their future intentions [15]. An *adaptation policy* determines how the interface should adapt given a contextual probability distribution. Prior works have explored several policies, such as frequency, recency, and access rank, among others (e.g. [5, 7, 13]). However, these typically assume that the system’s predictions are accurate without fully considering the cost of inaccurate predictions. As such, they can be ‘greedy’ in that they aim to optimize adaptations solely based on these predictions by recommending highest probability actions. When making adaptations, they do not fully take into account the UI design or

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI 2022, April 30–May 06, New Orleans, Louisiana, USA.

© 2022 Copyright held by the owner/author(s).

¹In this paper, we assume that contextual predictions are available to the adaptive system; we do not develop this predictive model here.

the fact that the user might have learned parts of it. To address these gaps, in the present work, we introduce a model-based approach for adapting XR interfaces. The system considers the utility of adaptations given contextual observations of the user to select an optimal adaptation that can both maximize the benefits of an accurate prediction and minimize the costs incurred otherwise.

2.2 General Approach: Model-Based Simulation

In our model-based approach, the goal is to estimate the *utility* [8, 16] of an adaptation given a contextual probability distribution over actions supported by the interface. This utility is a combination of the *benefits* provided by an adaptation and the *costs* incurred due to inaccurate predictions. As such, it is tightly coupled to both the system’s predictions and the user interface design itself. The key to computing utility is simulation of the user’s interaction sequence when completing a given task. We consider target acquisition tasks, where a user interacts with the UI to find and select a *target item*, which is common when interacting with typical graphical UIs such as menus, toolbars, and interactive applications. Each interaction sequence – navigating from a starting point to the target item – has a *total cost* associated with it. This could be quantified by, for example, the time taken to complete the interaction sequence, cognitive load, or physical exertion. Conversely, an adaptation has a *benefit* associated with it; this is the cost that is circumvented through adaptation of the starting point.

For any general adaptive interface, the utility of an adaptation is a combination of the benefit provided by the adaptation and the costs incurred due to inaccurate adaptations where the user’s target item differs from the one favored by the adaptation. By modeling the sequence of actions taken by a user to complete an interaction task and estimating the costs incurred, an adaptive system can effectively use simulations to estimate the utility of each possible adaptation. It can then select the utility-maximizing adaptation to optimize the interface given the user’s current context. We develop one such adaptive system and elaborate upon the simulation model that drives interface adaptations here.

3 MODEL FORMULATION

3.1 Application: Adaptive XR System

Menus have received extensive attention in HCI research as they are widely used and adaptation has potential to improve usability [1]. Our simulation-based interaction model studies the case of *hierarchical menus*. In our design, the user can traverse the menu hierarchy, where items are grouped semantically into sub-menus, and select items using *selection* actions. A dedicated *correction* action (e.g. back button or gesture) enables the user to undo a selection, or go one step up the hierarchy to a parent menu. While static menus always provide the top-level item (root node) as the starting point for interaction, we develop an *adaptive menu* that can dynamically select any item as a starting point. This could either be the root node, any intermediate sub-menu, or a leaf item (final action). Figure 1 illustrates the adaptive system using an example menu. Our simulation model adapts such a menu with the objective of minimizing *selection time*.

3.2 Interaction Cost Parameters

Previous literature has developed models to explain a user’s search process within linear menus [18]. We build upon prior work to formulate a model for simulating user interactions with our hierarchical menu-based XR interface. The model takes into account cognitive and motor components of the interaction, given by the three *time-based cost* parameters, to estimate the total cost of a specific interaction event.

3.2.1 Inspection Cost ($T_{inspect}$): When traversing a menu, the user needs to first visually inspect the interface to find relevant items. Each visual inspection has a cost associated with it. In this work, we consider that inspection cost is item-independent, and use a constant time-based cost of inspection $T_{inspect}$.

3.2.2 Selection Cost (T_{select}): This parameter defines the cost associated with taking a selection action. As the user inspects the interface and finds relevant sub-menus or the target item, they provide input by selecting the item. This selection could be using a point-and-click technique, a dedicated gesture, or other input modalities. The cost associated with this action is directly dependent on the input technique, and is influenced by factors such as input precision, encumbrance, and noise.

3.2.3 Correction Cost ($T_{correct}$): During interaction, the user might need to take corrective actions either when they make an incorrect selection or when the system presents an adaptation that is not suitable towards the user’s actual intention. Such corrective actions also have costs associated with them. Similar to the selection cost, the input technique used to enable corrective actions influences the correction cost.

3.3 Simulating User Interactions

We formulate a model for simulating user interactions with the described adaptive XR interface. Given a starting menu item as stimulus, it produces the set of actions required to acquire a target item. We consider that a user adopts a serial search strategy, where they inspect each menu item until they find their target item or they reach the end of the menu.

3.3.1 Search-and-Select: During a target acquisition task, the user inspects (reads) the menu interface until they either find a sub-menu item that contains their target item or the target item itself; upon finding this item, they make a selection. The total time cost for this sequence of actions ($T_{search}(i_l)$) is given by:

$$T_{search}(i_l) = \sum_{j=1}^l T_{inspect} + T_{select} \quad (1)$$

where i_l is either the target sub-menu or target item and l is the location of i_l . To recursively search and select within sub-menus, the above search process is repeated and costs are accumulated.

3.3.2 Backtracking: If the user encounters a sub-menu that does not contain the target item, they need to take a sequence of corrective actions, called *backtracking*, to traverse the menu hierarchy and find the appropriate parent menu before beginning the above *search-and-select* process. Each corrective action incurs a *correction cost*, following which, items in the parent sub-menu are displayed. The

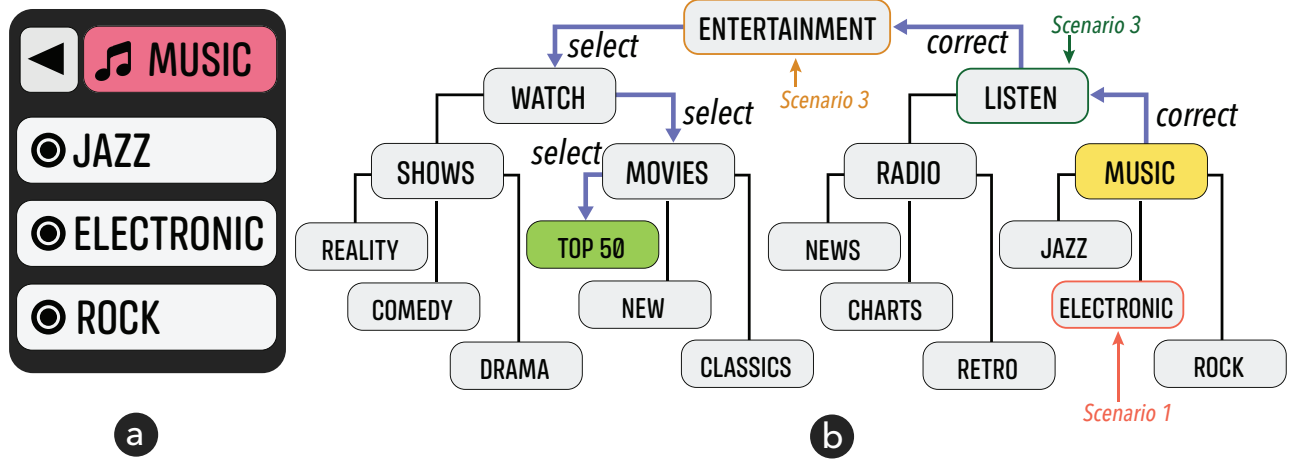


Figure 1: (a) An example application UI, with ‘Music’ as the currently selected item. (b) Hierarchical structure of the application menu, with ‘Entertainment’ as the top-level item (root). The structure shows the path from the current item (‘Music’, highlighted in yellow) to the target item (‘Top 50’, highlighted in green). To complete this target acquisition task, the user needs to make two corrections followed by three selections. The tree diagram is also annotated to show the selected adaptations for the three scenarios.

user now serially searches for a relevant item within this sub-menu; if they fail to find such an item, they repeat these backtracking steps until they encounter a relevant item. The total backtracking cost to recursively find an item related to the target item is given by:

$$T_{backtrack} = \sum_{k=1}^n (T_{correct} \cdot (l_k \cdot T_{inspect})) \quad (2)$$

where n is the number of sub-menus and l_k is either the position of the relevant item within a sub-menu k or the number of items in sub-menu k if the relevant item is not present.

3.3.3 Total Interaction Cost: Using the two components above, search-and-select and backtracking, we can compute the total cost of any target acquisition task. When the user is presented with an initial menu item k and has a target l , this cost is given by:

$$T(k, l) = T_{backtrack} + T_{search} \quad (3)$$

where $T_{backtrack}$ is the total time cost required to backtrack to the closest common parent item of k and l , and T_{search} is the time cost for inspecting the menu and making selections until the target item l is acquired.

3.4 Selecting Adaptations

With the above formulation, we can predict the total cost for any interaction sequence. As described in subsection 3.1, the system adapts the interface by selecting a suitable starting point k , given a probability distribution p over all leaf menu items. Each adaptation has a *utility* associated with it, which is the combination of the *benefit* and *costs* of an adaptation k :

$$Utility_k = \left(\sum_{i=0}^n p_i \cdot T(k, i) \right) - (p_k \cdot Benefit(k)) \quad (4)$$

where n is the number of leaf items, $T(k, i)$ is the cost for starting from item k when the target is i , p_i is the predicted probability of leaf item i , and $Benefit(k)$ is the benefit of starting at item k . For our menu-based interface, given a target item l , benefit of adaptively selecting a starting item k is given by $Benefit(k) = p_k \cdot T(0, k)$.

With this formulation, we can predict the utility of every possible adaptation in the interface. As this utility estimates the average selection time, we can then select the adaptation a that minimizes this value:

$$A = \underset{i \in n}{\operatorname{argmin}} U_i \quad (5)$$

4 WALKTHROUGH

4.1 Setup

Menu Interface: We implement an application that allows users to select from a range of actions: watching TV shows or movies and listening to radio stations or specific music genres. These actions are presented in a hierarchical menu, where a (sub-)menu and its contents are displayed in a list. Additionally, a ◀ “back” button allows for corrective actions (Figure 1-a). Figure 1-b illustrates the hierarchical structure of the menu interface and the sequence of actions taken by a user to navigate from a starting item to select a target item.

Probability Distribution: The XR device observes the users context to generate a probability distribution over the given action space. It can, for instance, observe the user’s current location, activity, and time-of-day to make inferences about what their intentions would be. We consider the following exemplary probability distribution for leaf items to demonstrate various scenarios:

$p = \{\text{Reality: } 0.073, \text{ Comedy: } 0.024, \text{ Drama: } 0.098, \text{ Top 50: } 0.024, \text{ New: } 0.024, \text{ Classics: } 0.122, \text{ News: } 0.11, \text{ Charts: } 0.085, \text{ Retro: } 0.122, \text{ Jazz: } 0.073, \text{ Electronic: } 0.22, \text{ Rock: } 0.025\}$.

Item	Probability	Scenario 1: Utility (in ms)	Scenario 2: Utility (in ms)	Scenario 3: Utility (in ms)
Entertainment	1.000	8121.0	2121.0	7605.0
Watch	0.365	6731.0	3461.0	8050.0
Listen	0.635	4984.0	2254.0	5525.0
Shows	0.195	6883.5	5663.5	10427.5
Movies	0.170	7067.0	5747.0	10495.0
Radio	0.317	4808.9	4076.9	7442.5
Music	0.318	4736.5	4008.5	7114.5
Reality	0.073	7687.2	8125.2	12441.0
Comedy	0.024	8226.3	8370.3	12931.5
Drama	0.098	7370.5	7958.5	11982.5
Top 50	0.024	8299.2	8443.2	12996.0
New	0.024	8294.4	8438.4	12972.0
Classics	0.122	7152.8	7884.8	11674.0
News	0.110	5857.0	6517.0	9659.0
Charts	0.085	6120.0	6630.0	9849.0
Retro	0.122	5673.8	6405.8	9283.0
Jazz	0.073	6228.8	6666.8	9865.0
Electronic	0.220	4523.7	5843.7	7954.5
Rock	0.025	6761.2	6911.2	10367.0

Table 1: Model results showing the list of menu items, probabilities, and utility values for the three scenarios. The value associated with the optimal adaptation for each scenario are shown in bold (Electronic = 4523.7 ms for scenario 1, Entertainment = 2121.0 ms for scenario 2, Listen = 5525.0 ms for scenario 3.)

4.2 Example Scenarios

Here, we present a set of illustrative scenarios that imitate different conditions or properties of the XR system, to study how they might result in varying model predictions and adaptations. For each scenario, Table 1 presents the model outputs (utility) for all menu items.

4.2.1 Scenario 1: low inspection cost, high selection cost, low correction cost. Here, the XR system is designed such that inspections have a low cost, corrections have a low cost, and selections have high cost. This might be the case when the UI displays texts in large fonts, simple gestures are used for correction (e.g. swipe), and selection is tedious (e.g. point-and-click). We use the following values: $T_{inspect} = 100$ ms, $T_{select} = 2500$ ms, $T_{correct} = 500$ ms.

The system selects Electronic as the optimal adaptation. Since this item has high probability and selection is more time-consuming than correction or inspection, it is advantageous for the system to make this *optimistic* adaptation.

4.2.2 Scenario 2: low inspection cost, low selection cost, high correction cost. In this scenario, while inspection is similar to scenario 1, selection now has a low cost (e.g. pinch), while correction has a higher cost (e.g. point-and-click). We use the following values: $T_{inspect} = 100$ ms, $T_{select} = 500$ ms, $T_{correct} = 2500$ ms.

The system selects Entertainment as the adaptation to optimize the utility. Given the current context, with uncertainty in the contextual predictions, the system chooses the root item for a *conservative* adaptation to avoid the high correction costs.

4.2.3 Scenario 3: high inspection cost, average selection cost, average correction cost. Finally, we consider a case where inspection costs are relatively high; this might be the case when text is rendered in small fonts, or items are represented using abstract icons. Additionally, we select average selection and correction costs, where similar gestures are used for both actions (e.g. left-swipe and right-swipe).

Here, the system selects Listen as the adaptation. By selecting a high-probability sub-menu, which is neither the root nor a leaf item, it considers the trade-offs to optimize the average selection time using a *balanced* adaptation.

4.3 Summary

As the above scenarios illustrate, our simulation model successfully takes various interaction costs into account to compute utility. By doing so, it can select optimal adaptations, which vary between scenarios, that minimize the average selection time. In contrast, to the above results, a greedy approach would always select the highest probability item (Bonobo) as the adaptation.

5 DISCUSSION

In this position paper, we have argued for a computational approach to adapt XR interfaces that considers both costs and benefits of adaptations. We presented a novel model that considered three cost components – inspection, selection, and correction – towards simulating user interactions with a hierarchical menu. Given a contextual probability distribution, the model can estimate utility and select optimal adaptations that minimize interaction time. Our illustrative scenarios demonstrate the benefits of such a model:

by considering varying cost parameters, it can enable a system to select varying adaptations that are suitable for particular scenarios.

We see several exciting topics for future research on model-based adaptive interfaces. While we consider a serial search strategy, an extended model could account for varying visual search strategies such as foraging and recall. Additionally, model predictions can be improved by identifying user expertise from prior interactions and including it in the model. Our model considers total interaction time as the objective for optimization; exploration of other aspects, such as ergonomics [3] and learnability [12] and their associated cost parameters will be important for optimal adaptive XR systems. Furthermore, extending the model to consider different adaptation styles will increase its applicability in various systems. For instance, our model considers a hierarchical menu where a single item is recommended as the starting point. Future adaptive systems can benefit by similarly developing simulation-based models for scenarios where multiple recommendations, or custom shortcuts, are presented.

To conclude, computational methods for generating or adapting their interfaces are a promising approach for the future of highly usable XR systems. By adapting based on both contextual inference and interaction utility, they will be able to overcome new challenges imposed by the device and the environment. We believe future applications can benefit from applying simulation-based interaction models, and look forward to developing and discussing novel models and approaches for such adaptive interfaces.

REFERENCES

- [1] 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* (St. Andrews, Scotland, United Kingdom) (UIST '13). Association for Computing Machinery, New York, NY, USA, 331–342. <https://doi.org/10.1145/2501988.2502024>
- [2] Brendan David-John, Candace Peacock, Ting Zhang, T. Scott Murdison, Hrvoje Benko, and Tanya R. Jonker. 2021. Towards Gaze-Based Prediction of the Intent to Interact in Virtual Reality. In *ACM Symposium on Eye Tracking Research and Applications* (Virtual Event, Germany) (ETRA '21 Short Papers). Association for Computing Machinery, New York, NY, USA, Article 2, 7 pages. <https://doi.org/10.1145/3448018.3458008>
- [3] João Marcelo Evangelista Belo, Anna Maria Feit, Tiare Feuchtnner, and Kaj Grønbaek. 2021. XRgonomics: Facilitating the Creation of Ergonomic 3D Interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 290, 11 pages. <https://doi.org/10.1145/3411764.3445349>
- [4] Leah Findlater, Karyn Moffatt, Joanna McGrenere, and Jessica Dawson. 2009. Ephemeral Adaptation: The Use of Gradual Onset to Improve Menu Selection Performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 1655–1664. <https://doi.org/10.1145/1518701.1518956>
- [5] Krzysztof Z. Gajos, Mary Czerwinski, Desney S. Tan, and Daniel S. Weld. 2006. Exploring the design space for adaptive graphical user interfaces. In *Proceedings of the working conference on Advanced visual interfaces*. ACM, 201–208.
- [6] Christoph Gebhardt, Brian Hecox, Bas van Opheusden, Daniel Wigdor, James Hillis, Otmar Hilliges, and Hrvoje Benko. 2019. Learning Cooperative Personalized Policies from Gaze Data. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 197–208. <https://doi.org/10.1145/3332165.3347933>
- [7] Camille Gobert, Kashyap Todi, Gilles Bailly, and Antti Oulasvirta. 2019. SAM: Self-Adapting Menus on the Web. In *Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion* (Marina del Rey, California) (IUI '19). Association for Computing Machinery, New York, NY, USA, 75–76. <https://doi.org/10.1145/3308557.3308705>
- [8] Eric Horvitz. 1999. Principles of Mixed-Initiative User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania, USA) (CHI '99). Association for Computing Machinery, New York, NY, USA, 159–166. <https://doi.org/10.1145/302979.303030>
- [9] Ben Lafreniere, Tanya R. Jonker, Stephanie Santosa, Mark Parent, Michael Glueck, Tovi Grossman, Hrvoje Benko, and Daniel Wigdor. 2021. False Positives vs. False Negatives: The Effects of Recovery Time and Cognitive Costs on Input Error Preference. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, 54–68. <https://doi.org/10.1145/3472749.3474735>
- [10] Luis A. Leiva. 2011. Restyling Website Design via Touch-Based Interactions. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services* (Stockholm, Sweden) (MobileHCI '11). Association for Computing Machinery, New York, NY, USA, 599–604. <https://doi.org/10.1145/2037373.2037467>
- [11] Yang Li and Otmar Hilliges. [n. d.]. *Artificial Intelligence for Human Computer Interaction: A Modern Approach*. Springer.
- [12] Linda S. Liebelt, James E. McDonald, Jim D. Stone, and John Karat. 1982. The Effect of Organization on Learning Menu Access. *Proceedings of the Human Factors Society Annual Meeting* 26, 6 (1982), 546–550. <https://doi.org/10.1177/154193128202600613> arXiv:<https://doi.org/10.1177/154193128202600613>
- [13] Jeffrey Mitchell and Ben Shneiderman. 1989. Dynamic versus static menus: an exploratory comparison. *ACM SigCHI Bulletin* 20, 4 (1989), 33–37.
- [14] Antti Oulasvirta, Per Ola Kristensson, Xiaojun Bi, and Andrew Howes. 2018. *Computational interaction*. Oxford University Press.
- [15] Tanya R. Jonker, Ruta Desai, Kevin Carlberg, James Hillis, Sean Keller, and Hrvoje Benko. 2020. The Role of AI in Mixed and Augmented Reality Interactions. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA). ACM Press, New York, NY. <https://doi.org/10.1145/3334480.XXXXXXX>
- [16] Leonard J. Savage. 1972. *The foundations of statistics*. Courier Corporation.
- [17] Andrew Sears and Ben Shneiderman. 1994. Split menus: effectively using selection frequency to organize menus. *ACM Transactions on Computer-Human Interaction (TOCHI)* 1, 1 (1994), 27–51.
- [18] Kashyap Todi, Gilles Bailly, Luis Leiva, and Antti Oulasvirta. 2021. Adapting User Interfaces with Model-Based Reinforcement Learning. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 573, 13 pages. <https://doi.org/10.1145/3411764.3445497>
- [19] Kashyap Todi, Jussi Jokinen, Kris Luyten, and Antti Oulasvirta. 2018. Familialisation: Restructuring Layouts with Visual Learning Models. In *23rd International Conference on Intelligent User Interfaces* (Tokyo, Japan) (IUI '18). Association for Computing Machinery, New York, NY, USA, 547–558. <https://doi.org/10.1145/3172944.3172949>