MRxD: A Usage-Pattern Framework for Mixed-Reality Cross-Device Interaction

Yang Zhan
North Carolina State University
Raleigh, North Carolina, USA
yzhan3@ncsu.edu

Yuxuan Huang
North Carolina State University
Raleigh, North Carolina, USA
yhuang94@ncsu.edu

Yuchen Cai
McMaster University
Hamilton, Ontario, Canada
caiy72@mcmaster.ca

Yichen Andy Yu
North Carolina State University
Raleigh, North Carolina, USA
yyu55@ncsu.edu

Xiaoran Yang
North Carolina State University
Raleigh, North Carolina, USA
xyang49@ncsu.edu

Ye Yuan
McMaster University
Hamilton, Ontario, Canada
irene.yuan@mcmaster.ca

Qiao Jin
North Carolina State University
Raleigh, North Carolina, USA
qjin4@ncsu.edu

arXiv:2611.16171v1 [cs.HC] 12 Nov 2026

Abstract

Mixed reality (MR) cross-device experiences often combine MR headsets with non-immersive devices such as laptops, tablets, and phones. Yet such cross-device design remains challenging because users must coordinate workflows across heterogeneous devices with different roles. This motivates our focus on usage patterns: recurrent ways in which people distribute and coordinate work across devices. Through a systematic analysis of 238 MR cross-device papers, we developed MRxD, a usage-pattern framework organized around three dimensions: user attention, task relatedness, and cognitive symmetry. Using these dimensions, we identified six usage patterns and synthesized the design strategies commonly used within each pattern. We further investigated the framework's application in design practice through a workshop (n=6) and the results suggest that MRxD can guide pattern-based design exploration and identify blind spots. We further discuss how usage patterns may shift or combine, and outline research gaps and design implications for future MR cross-device systems.

CCS Concepts

• **Human-centered computing** → **User centered design**; **Mixed / augmented reality**; **Virtual reality**.

Keywords

mixed reality, cross-device interaction, usage patterns

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference acronym 'XX, Woodstock, NY
© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

ACM Reference Format:
Yang Zhan, Yuxuan Huang, Yuchen Cai, Yichen Andy Yu, Xiaoran Yang, Ye Yuan, and Qiao Jin. 2018. MRxD: A Usage-Pattern Framework for Mixed-Reality Cross-Device Interaction. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 22 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Mixed Reality (MR)[1] cross-device interaction combines head-mounted displays (HMDs) with non-immersive devices (e.g., phones, desktops) to support activities that are difficult to handle through the headset alone. While HMDs provide immersive and spatial presentation, non-immersive devices often remain better suited for dense reading [8], long text entry [125], and routine productivity tasks [12, 142]. These complementary strengths have motivated prior systems to explore this combination in varied ways, including augmented displays [101, 116], control offloading [43, 150, 161, 162], and desktop-to-MR authoring [54, 72]. MR cross-device interaction is more than adding peripheral devices to an HMD, but a broader design space for coordinating various display and interaction resources. Prior cross-device research describes this design space through the distribution of interfaces, information, and interaction across device ecologies [18]. A recent work on MR Hybrid User Interfaces (MR-HUI) revisits the earlier notion of HUI for contemporary MR systems [55]. Building on Feiner's view [37] of HUIs as combinations of heterogeneous display and interaction technologies, MR-HUI characterizes systems that combine conventional 2D devices with MR environments in complementary roles [55]. Their synthesis describes these systems mainly through device and component configurations, dependency, anchoring, interaction dynamics, and system-level challenges.

However, device and component configuration alone does not fully explain how users interact with the cross-device system. The same device configuration can support different forms of cross-device activities. For example, an HMD-phone setup may use the

[1] MR refers to the combination of Virtual Reality (VR) and Augmented Reality (AR) in this paper, consistent with broad definitions of MR in prior work [132, 133].

handheld device as a standard smartphone [162], a menu surface [124], a spatial controller [81], or a text-entry device [22]. Despite the identical hardware setup, the underlying patterns of attention allocation, action flow, and cross-device interpretation beyond the phone's functional role may differ substantially, leading to distinct design requirements [114]. In cross-device research, usage-oriented analyses have been widely used to describe and form how cross-device systems fit into concrete activities in addition to how devices are connected [79, 155], which helps explain why the same device ecology can require different design choices.

This motivates a usage-centric account of MR cross-device interaction. Existing system-oriented frameworks identify user-facing issues such as visual attention switching and perceptual synchronization [55] as vital challenges, but the usage patterns behind them remain under-specified. We use the term **usage pattern** to describe recurring ways in which users coordinate immersive and non-immersive devices during a usage activity [155]. A usage-pattern framework can support both analysis and design in MR cross-device interaction. It helps researchers compare systems by how devices work together in use beyond device setup. It also helps designers begin with intended user activities before deciding which devices and interaction techniques to use. Thus, this paper develops a usage-centered synthesis of MR cross-device interaction. We ask:

RQ1. What usage patterns are reported in prior MR cross-device research?

RQ2. What design strategies are associated with these usage patterns, and how can these design strategies be applied to generate new MR cross-device designs?

To answer these questions, we conducted a systematic review of MR cross-device research published between January 2015 and December 2025. The review followed a PRISMA-informed [98] collection and screening process and resulted in a corpus of 238 papers. We then developed a usage-pattern framework using a hybrid deductive–inductive method, applied it across the corpus, and inductively synthesized recurring design strategies within each pattern. Based on the findings, we proposed **MRxD**, a framework characterizing MR cross-device usage through three dimensions: *User Attention*, or how users distribute attention across devices; *Task Relatedness*, or whether devices support the same activity or separate activities; and *Cognitive Symmetry*, or whether users can carry the same interpretation across devices or must construct a new interpretive mapping between them. We use these dimensions to identify recurring usage patterns and to organize design strategies that prior systems use to support them. A design workshop ($n = 6$) was further conducted to examine how the framework can support design exploration. In summary, this paper contributes:

(1) **MRxD, a usage-pattern framework derived from a systematic analysis of 238 MR cross-device papers**, which characterizes systems through User Attention, Task Relatedness, and Cognitive Symmetry.
(2) **Six recurring usage patterns and associated design strategies** that explain how immersive and non-immersive devices are coordinated across different forms of MR cross-device activity.
(3) **An exploratory design workshop ($n = 6$) and resulting design guidance**, providing initial evidence of how MRxD can support pattern-based design exploration and help designers identify overlooked design opportunities.

2 Background

2.1 Cross-Device Interaction

Cross-device interaction broadly concerns interactions that extend beyond the boundary of an individual device by coordinating multiple devices, displays, or interaction resources within a shared activity or device ecology [18]. Research in this area has developed over several decades, from multi-monitor workspaces [126] and augmented display environments [121] to increasingly mobile and ad-hoc device ecologies [46, 76, 112]. This long development has also produced several overlapping research traditions. For example, *distributed user interfaces* consider how interface components and interaction resources can be distributed across multiple computational and physical resources [30], while multi-display research has examined how visual attention shifts across displays [114]. Related work has also been framed using overlapping terms such as *cross-surface interaction* [52], *multi-device interaction* [5], and *cross-display interaction* [25]. Although these traditions overlap and findings from one can often inform another, they developed with different scopes and terminology, leaving the broader research landscape conceptually fragmented. Brudy et al. [18] synthesized these strands under the umbrella of *cross-device computing*, providing a common vocabulary for interactions spanning diverse device configurations and ecologies. Their taxonomy organizes the cross-device design space along six dimensions, including *temporal*, *configuration*, *relationship*, *scale*, *dynamics*, and *space*, which collectively characterize how cross-device ecologies are organized, used, and changed over time. Because these dimensions are not tied to a particular device class, the taxonomy encompasses conventional displays, mobile and wearable devices, tangibles, and head-worn AR/VR devices.

Brudy et al.'s broad framing of cross-device interaction encompasses both homogeneous and heterogeneous combinations of devices. Within this broader landscape, HUIs, originally introduced by Feiner and Shamash [37] as combinations of heterogeneous display and interaction device technologies, emphasize the complementary use of heterogeneous components. Focusing on the prevalent combination of conventional 2D components with MR environments, Hubenschmid et al. [55] characterize this subset as MR-HUIs. Their taxonomy builds on Brudy et al.'s cross-device framework, adapting its six dimensions and introducing *interaction dynamics* and *anchoring* to account for interaction and spatial relationships across MR and 2D components. Collectively, these taxonomies provide rich vocabularies for describing how cross-device systems are composed and organized across device ecologies.

2.2 Usage-Oriented Perspectives in Cross-Device Interaction

Beyond characterizing device configurations and interaction techniques, cross-device research has also examined how activities unfold across devices in use. Cross-device taxonomy has captured usage-facing characteristics such as temporal organization and user–device relationships [18], while empirical studies have examined task organization and attention across devices. For example,

Jokela et al. distinguished sequential and parallel use as well as related and unrelated tasks [65], while Rashid et al. [114] showed that display arrangement and content coordination can shape attention switching. In cross-reality settings, a prior work found that transitions among desktop, AR, and VR can introduce friction when they disrupt users' spatial mental models, highlighting the importance of maintaining spatial continuity across realities [146].

To capture such behavior more systematically, prior work has also organized recurring cross-device activities into usage patterns. Yuan et al. [155] studied everyday multi-device work practices and identified five recurring patterns: *integrating*, *cloning*, *expanding*, *partitioning*, and *migrating*. Their analysis positions usage patterns as a frame of reference for grounding cross-device technologies in workflows observed in practice. They further related these patterns to physical configurations, user motivations, and opportunities for cross-device design. Li et al. [79] applied a similar usage-pattern perspective to in-the-wild cross-device learning while additionally examining attention allocation and shifting across devices. Their findings show how recurring workflows and attention issues can reveal design opportunities for cross-device learning platforms.

Together, these studies demonstrate the value of examining cross-device systems through how users coordinate tasks, attention, and transitions in use, complementing descriptions of how devices and interfaces are configured. However, existing usage-pattern studies have primarily focused on conventional non-immersive device ecologies. This motivates examining how such a usage-oriented perspective extends to interactions that combine immersive and non-immersive devices.

2.3 Scope and Research Positioning

For the purposes of this review, we use *MR cross-device system* to refer to an HMD-based MR system used together with at least one physically separate or peripheral non-immersive device during a concrete user activity. We use *MR cross-device* as a descriptive review term grounded in two lines of prior work: cross-device taxonomy, which studies interaction across device ecologies [18], and MR-HUI review, which focuses on complementary combinations of conventional 2D devices and MR environments [55]. Our review focuses on widely available, general-purpose non-immersive devices and peripherals, including phones, tablets, laptops, desktops, keyboards, mice, pens, and smartwatches. These devices recur across MR systems and often take clear roles in display, input, control, text entry, pointing, and sensing. This boundary helps us identify usage patterns and design strategies around general forms of cross-device coordination, while excluding highly specialized or one-off hardware whose interaction is mainly shaped by device-specific mechanisms.

Our scope differs from the MR-HUI review [55] in its organizing perspective and inclusion criteria. While that review primarily characterizes MR-HUI systems through device and component configurations, our review is organized around *usage patterns*: how users coordinate immersive and non-immersive devices during concrete activities. Because our framework is intended to characterize recurring forms of cross-device coordination rather than particular MR-HUI component configurations, we use the broader term *MR cross-device*. MR-HUI nevertheless provides an important basis for defining the systems examined in this review. We also retain physically distinct non-immersive devices when they shape the user's cross-device activity, even if they provide only input, output, or sensing support, as well as asymmetric collaborative setups in which users participate through different device arrangements within the same shared activity.

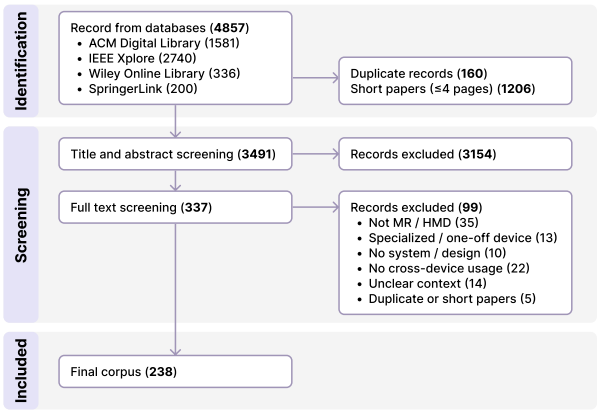


Figure 1: PRISMA diagram of paper collection and screening

3 Methods

We conducted a PRISMA-informed systematic review [98] to answer our RQs. Based on the identified papers, we developed a usage-centric framework through combined deductive and inductive process, and analyzed the design strategies associated with each usage pattern.

3.1 Paper Collection

We searched four digital libraries that publish substantial portions of HCI and MR research: ACM Digital Library, IEEE Xplore, Wiley Online Library, and SpringerLink. The search covered work published between January 2015 and December 2025 to capture recent MR cross-device systems and contemporary devices. We searched titles, abstracts, and keywords where supported by each digital library. For SpringerLink, whose advanced search does not support abstract-specific searches, we searched titles and keywords only. The search query comprised three keyword sets: a main set, a supplementary set, and a reality set. In abstract form, the query combined the main and supplementary sets with the reality set, while a small subset of terms in the main set that inherently describe hybrid or cross-reality interfaces were searched independently. The main set was informed by prior reviews of MR-HUI and cross-device taxonomies [18, 55]. It included terms directly related to MR cross-device interaction, such as "cross-device," "hybrid user interfaces," and "multi device." Considering the fragmented terminology in this area [18, 55], the supplementary set included device and interaction terms that could imply cross-device usage, such as "phone," "tablet," "laptop," "smartwatch," "keyboard," "mouse," "pen," and "asymmetric interaction." The reality set included "augmented reality," "virtual reality," "mixed reality," and "extended reality." The full search query

is provided in Appendix A. The search returned 4,857 records[2]. We removed 160 duplicates and excluded papers of four pages or fewer. Because our analysis required sufficient descriptions of user activities and system design to analyze cross-device usage behavior, we used paper length as a pragmatic minimum-reporting threshold, following similar page-length criteria adopted in prior systematic reviews [116, 153]. This process resulted in 3491 records for title and abstract screening.

3.2 Screening and Eligibility Check

Screening proceeded in two stages. First, we screened titles and abstracts to remove papers that were not written in English or were clearly outside the scope of the review. Specifically, we excluded papers that were (1) content-level duplicates (e.g., poster [45] of a full paper [44]), (2) unrelated to MR or MR head-mounted displays (e.g., handheld AR [71]), (3) centered on highly specialized or one-off devices whose roles were not transferable to common cross-device settings (e.g., haptic gloves [131]), or (4) did not propose a system or design for MR cross-device interaction (e.g., surveys [92]). This stage left 337 papers for the second stage: full-text eligibility assessment. In addition to the criteria above, we excluded papers when the full text did not describe actual cross-device usage or interaction, such as papers focused only on backend algorithms [134, 149], mock devices [160], or comparisons between devices [137]. We also excluded papers when the usage context was unclear (i.e. without a concrete application scenario [75, 148]). This process resulted in a final corpus of 238 papers[3]. Figure 1 provides a detailed breakdown of the numbers of papers at different review stages.

Before screening, the first four authors discussed the exclusion criteria to establish a shared interpretation of the review scope. In the first stage, records were divided among the four authors and screened individually. In the second stage, three authors reviewed the assigned papers. Across both stages, uncertain cases were marked and discussed until the authors reached an agreement.

3.3 Framework Development and Corpus Analysis

Our framework development and strategy synthesis were qualitative and iterative, supplemented by descriptive frequency and co-occurrence analyses. We first developed usage dimensions and used their configurations to identify recurring usage patterns (**RQ1**); we then analyzed how prior systems realized each pattern through design strategies (**RQ2**). Rather than coding usage patterns directly, we derived them from a small set of *usage dimensions* (observable aspects along which MR cross-device use differs), where a *usage pattern* is a recurring configuration of values across these dimensions. Coding dimensions first avoided grouping systems by surface similarities such as shared devices or applications, and instead made the criteria for distinguishing patterns explicit and consistent across a heterogeneous corpus. We analyzed patterns before strategies because a design strategy is a way of realizing a particular usage pattern; interpreting strategies therefore depends on first establishing the patterns they serve. As summarized in Figure 2, this proceeded in three stages: developing usage dimensions, coding usage patterns, and analyzing design strategies.

The unit of analysis was the *usage episode*: a continuous segment of activity in which the user's goal, involved devices, and interaction workflow remained stable. Before identifying or applying the dimensions, we segmented the reported workflows into candidate episodes based on observable changes in these features. For example, [57] contributed two episodes: a complementary sequential usage and an in-situ interaction. This episode-level unit allowed us to distinguish different forms of cross-device use within the same paper or system. We retained episodes central to a paper's design motivation or system contribution: those explicitly motivated, implemented, evaluated as a study condition, or discussed as a contribution. We excluded episodes that appeared only as comparative baselines [125] or as conventional interactions retained for system operation, such as standard typing beneath an augmented keyboard mapping [93]. When a paper gave equal attention to multiple techniques or scenarios, as in comparative studies [35] and design-space explorations [161], we included all corresponding episodes. These episodes served as the cases for developing the MRxD framework.

3.3.1 *Identifying Usage Patterns (RQ1).*

Developing usage dimensions. We developed the usage dimensions by alternating between conceptual (top-down) and empirical (bottom-up) analysis, informed by the iterative logic of taxonomy-development methods [97]. Following this approach, we anchored the process in a *meta-characteristic*, which provides the overarching basis for selecting relevant dimensions [97]. Given our usage-centric goal, we defined this meta-characteristic as how users coordinate immersive and non-immersive devices within a usage episode. A candidate dimension was retained only if it (1) was a logical consequence of this meta-characteristic, that is, captured a user-facing difference in cross-device coordination rather than an implementation detail, (2) could be coded from the interaction workflows reported for an episode, and (3) distinguished meaningful differences across episodes that were not already captured by another dimension. Notably, this scoping is deliberate: the dimensions are not intended to describe every way MR cross-device systems differ, but to provide a parsimonious and explanatory lens tied to our analytical purpose. Consistent with the philosophy of iterative taxonomy development, our goal was a useful rather than a uniquely optimal set of dimensions [97]; aspects that fell outside the meta-characteristic were excluded or kept as lower-level descriptors rather than elevated to dimensions.

We first assembled a broad set of candidate dimensions. Deductively, we drew usage-related dimensions from prior reviews, taxonomies, and empirical studies of cross-device use [4, 18, 55, 79, 116, 155], such as attention, task properties, usage flow, interaction and information distribution, spatial and temporal relations, and device roles. Inductively, three authors independently open-coded the usage episodes identified in a random 20% subset of the corpus to surface aspects not covered by prior work. We then compared, grouped, and abstracted these candidates into a small set of *initial* working dimensions. For example, data similarity and interaction mapping first appeared as separate aspects, but both reflected how users must interpret the relationship between devices and were subsumed under a single provisional dimension.

[2]The search date was April 28th, 2026.
[3]The full list of papers is available at https://mrxd-paper.pages.dev/.

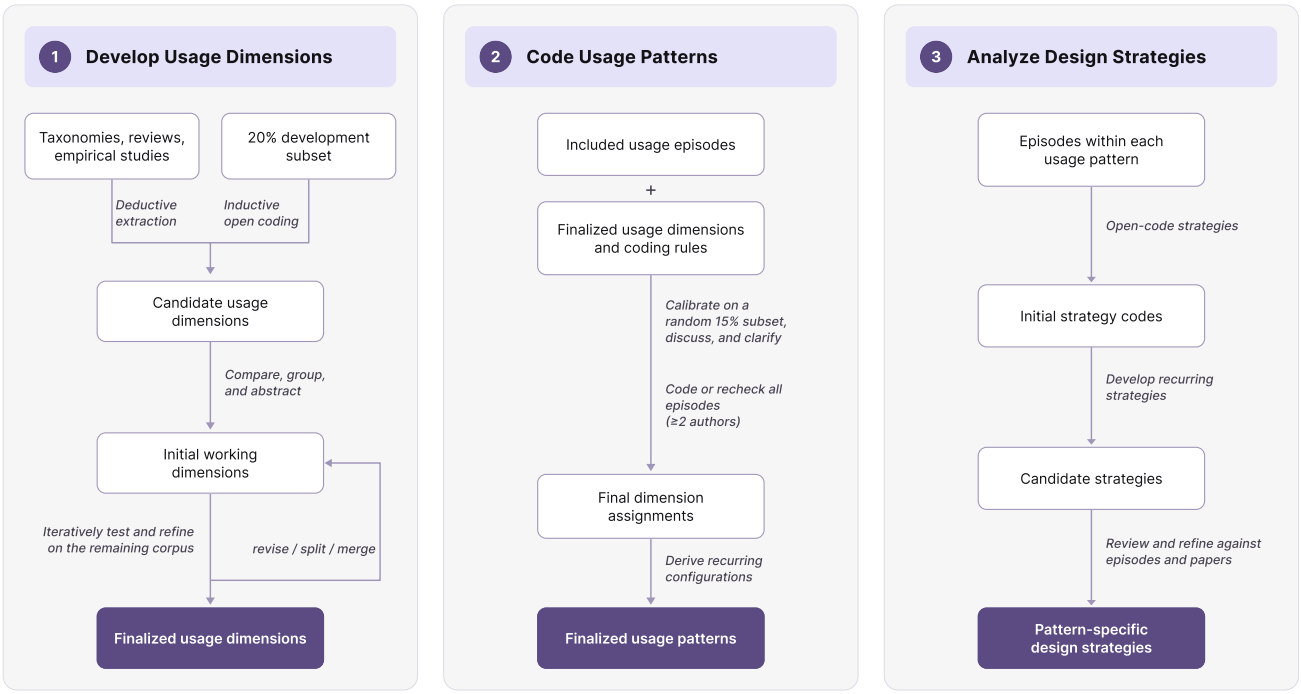


Figure 2: Overview of the MRxD framework development process. The analysis proceeded in three stages: (1) developing usage dimensions through deductive extraction from prior work and inductive open coding of a 20% development subset, followed by iterative corpus-wide refinement; (2) coding included usage episodes with the finalized dimensions and coding rules, calibrating and clarifying before deriving recurring configurations as usage patterns; and (3) inductively analyzing episodes within each usage pattern, from initial strategy codes through recurring and refined strategies, to derive pattern-specific design strategies.

One author then led an exploratory refinement pass over the remaining corpus to test whether the provisional dimensions and their distinctions could characterize usage episodes clearly and consistently. Dimension assignments made during this pass were functioned as diagnostic probes and used only to guide refinement. The other two authors reviewed the proposed revisions and flagged episodes, and all three authors discussed whether these cases required revised definitions, split or merged dimensions, or lower-level descriptors. Whenever a dimension changed, the affected episodes were revisited to test the revised formulation. For instance, we narrowed User Attention to overt visual attention, which could be inferred more consistently from the interaction behavior reported in papers. Drawing on Nickerson et al.'s use of explicit ending conditions [97], we finalized the dimensions when their definitions and distinctions remained stable through the final pass, all included episodes could be characterized, and no recurring residual cases required an additional independent dimension.

This produced three final usage dimensions, which we define and justify in Section 4.1: *User Attention*, *Task Relatedness*, and *Cognitive Symmetry*. We operationalized each dimension as binary with two categories to support consistent coding across the corpus. A usage pattern is then defined by a recurring configuration of values across these dimensions. This binary treatment involved a trade-off: while it made large-corpus coding feasible, it simplified experiences that may vary along a spectrum. We therefore treat the dimensions as design-oriented analytical categories for characterizing recurring MR cross-device use rather than as a closed taxonomy, and return to this limitation in Section 6.

Coding usage patterns. After finalizing the dimensions and their operational definitions, we conducted formal pattern coding across the corpus. Unlike the exploratory assignments used during dimension refinement, this stage aimed to produce the final and consistent pattern assignments using a fixed set of definitions and decision rules. Earlier assignments made during dimension refinement were not treated as final results; each episode was coded or rechecked against the finalized codebook. For each included episode, the three author assigned the applicable dimension values and derived the corresponding pattern. To calibrate the finalized coding scheme, the three authors independently coded the usage episodes from a random 15% subset papers. Pairwise Cohen's κ exceeded .70 across all three author pairs, indicating substantial agreement [94]. After discussing disagreements and clarifying the rules for edge cases, all other episodes were coded or rechecked against the final codebook by at least two authors, with disagreements resolved through discussion. For corpus-level prevalence and co-occurrence analyses, repeated instances of the same pattern within a paper were counted

once. Distinct implementations were retained as separate cases for design-strategy analysis.

3.3.2 Analyzing Design Strategies (RQ2). To understand how each pattern was instantiated in prior systems, we analyzed the design approaches within each pattern using reflexive thematic analysis [17], following prior review practices for synthesizing design strategies [2, 86]. The unit of analysis remained the usage episode, and an episode could contribute multiple strategies. One author open-coded the strategies across all usage episode and merged similar designs into higher-level categories; the other authors reviewed these categories against the associated papers to ensure that each reflected recurring design choices.

Alongside coding, we extracted structured metadata describing system context, device and user configurations, cross-device interaction and coordination, and physical arrangements, partly adapted from prior cross-device taxonomies [18, 55] and from the lower-level descriptors identified during dimension development. We treated these as descriptive context for the usage patterns rather than as the primary basis of the framework. Descriptive distributions of these metadata are provided in the Supplementary Material. For co-occurrence analyses, we used lift to normalize observed frequencies by category prevalence. This distinguished associations driven by the prevalence of common patterns or devices from those occurring more or less often than expected under independence. We applied lift to both pairwise pattern combinations and associations between device type and usage pattern.

4 MRxD Framework

Through the process described in Section 3, we propose **MRxD**, a usage-pattern framework for analyzing MR cross-device systems from a user-centered perspective. As shown in Figure 3, MRxD is built around three usage dimensions—*User Attention*, *Task Relatedness*, and *Cognitive Symmetry*. Different configurations of these dimensions define recurring usage patterns. For task-related episodes, crossing User Attention and Cognitive Symmetry yields four configurations: Clone, Proxy, Extend, and Complement. For task-unrelated episodes, Cognitive Symmetry is not applicable because the devices do not support a shared task, leaving two configurations distinguished by User Attention: Standby and Partition. For each pattern, we further identify design strategies that explain how prior systems supported that pattern in practice.

Figure 4 provides an overview of how these dimensions and usage patterns were distributed across the reviewed corpus. Most usage episodes involved task-related device use, while attention and cognitive symmetry showed more balanced distributions. At the pattern level, Clone and Proxy were the most prevalent, followed by Extend and Complement, whereas Standby and Partition occurred only occasionally.

4.1 Usage Dimensions

4.1.1 User Attention. User Attention describes whether users distribute visual attention across immersive and non-immersive devices during a usage episode. Prior work treats visual attention as a key design trade-off: additional displays can expand available resources, but may also introduce attention-switching costs or distraction [79, 114]. Attention distribution also helps reveal which device acts as the activity foreground [43] and where designers may need to reduce or guide context switching [79]. In MRxD, we focus on *overt* visual attention, observable through gaze or head orientation, rather than *covert* attention through peripheral perception [107], because overt attention is more observable in system descriptions and easier to operationalize for analyzing cross-device coordination. [4] We distinguish two forms: a usage episode is **Focused** when visual attention is directed to one device while another plays a supporting role, and **Fragmented** when users must visually attend to multiple devices to complete the task. The criterion is not switching frequency, but whether multiple spaces are visually necessary for task progress. For example, a sequential and exclusive workflow [56, 72] is considered as fragmented since users must move visual attention between device spaces. Mirrored output [10] to MR is also treated as fragmented when the content remains tied to the non-immersive device and requires coordination with the MR view.

Across the reviewed systems, focused overt attention was enabled in two broad ways. In many cases, the secondary device remained part of the interaction, but users could operate it through **tactile or peripheral attention** instead of direct visual inspection. This was common when MR systems were paired with keyboards, mice, or pens, or touchscreens [53, 90, 109, 135, 143], where users relied on touch [161], motor memory [81], or simple repeated gestures [29, 77]. In other cases, the cross-device configuration required **attention to only one device**. This occurred when relevant resources were duplicated, mirrored, or placed directly in the user's current view, removing the need to inspect the source device [109, 150]; when another device operated in the background by supplying sensed information without requiring user interaction [111]; or in collaborative settings where responsibilities were distributed across users, allowing each person to remain focused on their own device [44, 50, 74].

Two main reasons appeared for fragmented usage. For **resource access**, attention shifted because information needed for the ongoing task was distributed across devices. A common case involved moving between 2D screen-based content and 3D spatial content to understand correspondences [12, 40, 57], either across different representations of the same data [40, 129] or across different data entities within one activity [118]. Fragmented attention also occurred in multi-screen settings, where users shifted between one physical and multiple virtual 2D displays to access distributed resources [101]. Resource access could involve peeking behaviors, where users briefly returned to another device to compare information or confirm its status [128, 146]. Data transfer could likewise fragment attention when users accessed and moved a digital object across devices [115, 151]. For **control access**, attention shifted in two main situations. First, some control interfaces themselves required visual engagement, and thus users needed to look at the device to guide actions such as typing on touchscreen keyboards [120], sketching precisely [119], adjusting parameters [57, 146], or manipulating on-screen controls [24, 110]. Second, some interactions created a cross-device dependency between input and output, requiring users to act on one device and then monitor or assess the corresponding result on another. In some systems, such switching

[4] Unless otherwise stated, visual attention in this paper refers to overt visual attention.

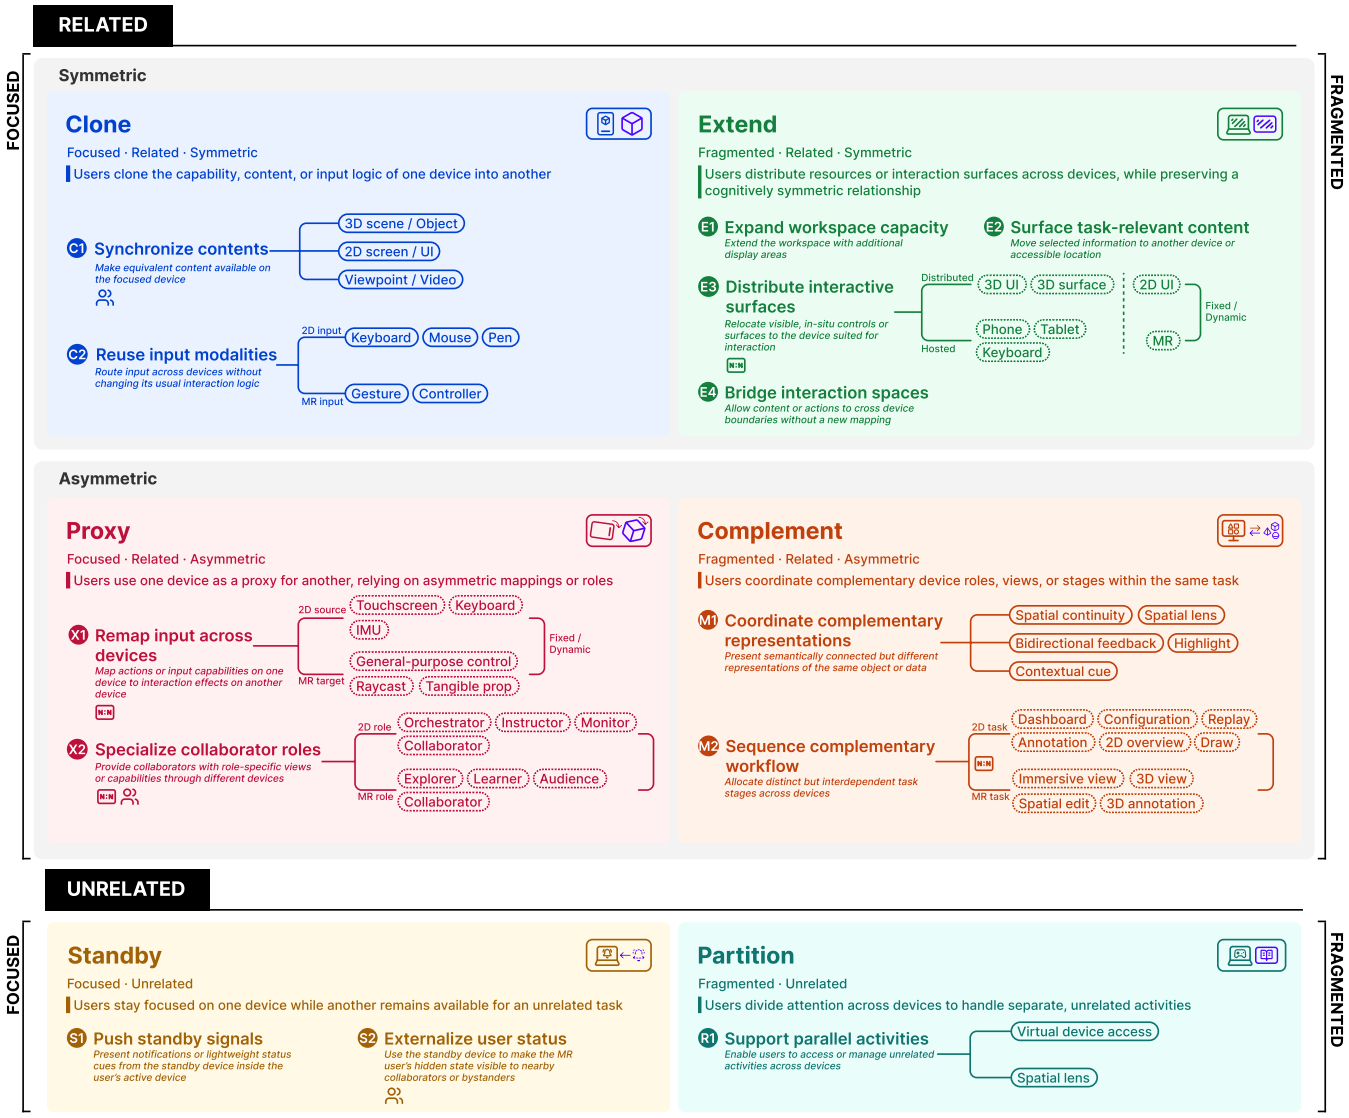


Figure 3: MRxD framework. The horizontal axis represents User Attention (Focused to Fragmented). The framework first separates Task Relatedness into Related and Unrelated use. Within Related use, patterns are further distinguished by Cognitive Symmetry (Symmetric vs. Asymmetric). Because Cognitive Symmetry is not applicable to Unrelated tasks, the framework yields six usage patterns rather than eight combinations. indicates collaborative cases, and indicates many-to-many mapping from two branches. The purple icons represent the virtual objects in MR.

was deliberately built into the cross-device interaction [24, 40, 120], and could also become embedded in users' own workflows as they coordinated actions and inspection across devices [128].

4.1.2 Task Relatedness. Task Relatedness describes whether tasks on immersive and non-immersive devices are connected. Tasks are **Related** when the devices contribute to a single overarching goal, even if each device supports a different subtask, representation, or interaction role. Tasks are **Unrelated** when the devices support different goals or parallel activities. This dimension extends prior non-MR cross-device work [155] and helps explain when devices are used together for one activity and when MR cross-device creates multitasking or coordination challenges across separate activities.

When tasks were related, we identified several recurring device relationships. In an **input medium** relationship, one device supplied input for the same task on another, through either user actions or sensed information. User-operated input flowed from non-immersive to immersive devices [20, 53], or in the reverse direction [16, 161], through modalities such as touch [20], typing [46], drawing [33], mouse movement [159], and gesture [16]. Devices also supplied sensed information, including physiological signals [111, 158], voice [48], and video streams [68, 130]. In

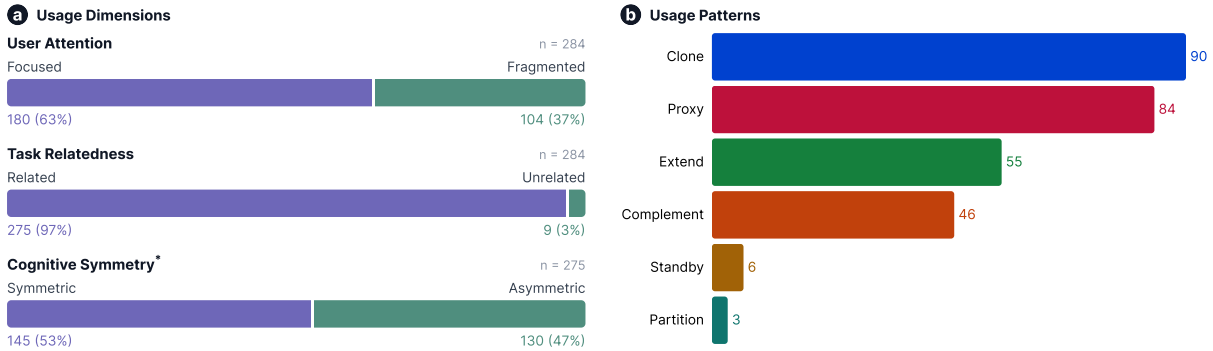


Figure 4: Prevalence of usage dimensions and usage patterns in the corpus. (a) Distribution of usage episodes across the three usage dimensions. Attention and Task Relatedness are reported for all 284 episodes, while Cognitive Symmetry is defined only for the 275 Related episodes. (b) Number of usage episodes exhibiting each usage pattern.

collaboration relationships, multiple users coordinated a shared task through different devices, often with asymmetric views [67], differentiated roles [61], or complementary interaction capabilities [44]. In **auxiliary display** relationships, one device presented information that users consulted to support task activities on another, such as references [8, 151], previews [72, 120], or supplementary views [83, 162]. **Workflow chain** relationships captured cases where users advanced the same task across stages, moving materials, decisions, or actions from one device to another sequentially [56, 84, 139, 146].

Unrelated tasks appeared less frequently, with two main forms. In **multitasking**, users maintained separate activities across immersive and non-immersive devices, such as messaging or watching video independently of the ongoing immersive task [3, 89]. In **standby** relationships, one device remained available as a background channel for notifications or an outward-facing channel for communication and social awareness, while the primary user stayed engaged in an immersive activity [14, 43, 88, 104].

4.1.3 Cognitive Symmetry. Cognitive Symmetry describes whether the user can carry the same interpretation across the devices, or must construct a new interpretive mapping to relate a representation or action on one device to its counterpart on another. A relationship is **Symmetric** when the correspondence is semantically direct, so users reuse an existing interpretation across devices without an additional interpretive mapping. It is **Asymmetric** when they must reinterpret how a representation or action on one device corresponds to information or effects on another. This distinction of Cognitive Symmetry is informed by the notion of semantic distance, which characterizes how directly interface expressions correspond to users' intentions and the interpretive transformation required to bridge them [59]. We extend this underlying intuition from user–interface mappings to cross-device correspondences involving representations and actions. Cognitive Symmetry matters because immersive systems often combine 3D spatial representations with conventional 2D interfaces and devices, which can create gaps in how information is represented, aligned, and manipulated across devices. Prior work has suggested that spatial mental models can help users predict outcomes and reduce friction across reality [146]. Cognitive Symmetry helps explain when designers should preserve direct cross-device correspondence and when clearer metaphors, alignment cues, or device-specific interaction designs may be needed.

To operationalize this distinction in our coding, we considered both representational and action–effect relationships. For representations, a change in device or display location does not by itself introduce asymmetry when the underlying representational structure is preserved; for example, a conventional 2D interface extended from a desktop into an MR virtual screen can still be interpreted in the same way. In contrast, when information is recast into a different representational structure, such as from a 2D visualization to a spatial 3D representation, users must interpret how elements and relations in one form correspond to those in the other. The same principle applies to action–effect relationships: asymmetry arises when an action on one device must be reinterpreted to understand its effect on another, rather than carrying over through an already familiar correspondence.

In the reviewed systems, symmetric relationships were often relatively straightforward. For representations, devices preserved similar representational forms, such as physical and virtual screens [42, 101], or presented the same or semantically equivalent content, including shared documents [83, 109], snapshots [58, 151], and views [117, 157]. For actions, familiar inputs retained their conventional effects, such as typing text with a keyboard [13], drawing strokes with a pen [33], or selecting a button by tapping it [124]. Asymmetric relationships instead introduced a new mapping between structures across devices. For representations, information could be organized or expressed in different forms. This commonly occurred when 3D spatial content was paired with 2D parameters, lists, charts, or abstract views [56, 57, 128, 129, 146]. For actions, users needed to interpret how an action on one device was mapped to an effect elsewhere, including what it controlled and where the resulting change appeared. For example, a smartphone could act as an MR controller [161, 162], a trackpad for a virtual cursor [81], or a spatial proxy for an object [89, 144].

Across these asymmetric cases, we observed two broad ways in which users could bridge the resulting correspondence, including **explicit bridging** and **implicit bridging**. These mechanisms are not mutually exclusive; users may combine several of them during

actual use. Explicit bridging relied on external support that made the cross-device relationship more observable or interpretable. One form is **system alignment**, where the interface provides perceptual cues that expose the correspondence between representations or action effects (e.g. bidirectional highlighting) [60, 76, 85]. When such cues are absent or insufficient, users may need **instructional scaffolding**, such as tutorials, onboarding, system prompts, researcher guidance [23, 57, 77, 147]. Under multi-user settings, users may also rely on **collaborative mediation**, where verbal communication [12, 44], monitoring [61], or shared annotations [102, 123] help reconcile differences between device views and roles.

Implicit bridging occurs when users infer an otherwise asymmetric cross-device relationship without relying on explicit external support, instead drawing on prior knowledge or the immediate task context. One form relies on **established mappings**, where users apply a correspondence already familiar from prior experience. Although familiar, such mappings still require users to relate representations across different domains and are therefore coded as asymmetric. Examples include interpreting a 2D projection of a 3D object [119], or connecting a 2D floor plan to a 3D space [102]. In other cases, users transfer a relational schema from a different familiar context to the cross-device setting. Following Structure-Mapping Theory (SMT) [41], this **analogical transfer** relies on aligning relational roles across domains, such as treating a phone as a handheld controller [162] or a pen as a pointer [78]. Users may also rely on **contextual grounding** to infer the relationship. For example, when users revise parameters or code on a desktop and immediately inspect the resulting change in MR, the temporal and interaction context helps them infer how the desktop action relates to MR [56, 118].

4.2 Usage Patterns and Design Strategies

4.2.1 ***Clone****: Focused, Related, and Symmetric.* In the Clone pattern ($n = 90$), users clone the capability, content, or input logic of one device into another while remaining focused on a primary device. The devices support the same task, and the cross-device relationship is cognitively symmetric because users can understand the relationship directly through the duplication. Among the usage patterns, Clone represents a foundational motivation for cross-device usage: it serves both as a basic synchronization mechanism and as a fallback or alternative interaction mechanism that retains existing device functions alongside more elaborate cross-device capabilities. However, this copy-and-paste logic can sometimes limit adaptation to the affordances of the receiving device. Two design strategies were used to realize the Clone pattern:

C1 Synchronize content or views. The system clones the same or equivalent content across devices so that users can build common ground, access resources without returning to the original device, or allow another user to observe the same activity. This includes shared 3D scenes for collaboration [50, 66, 127], shared first-person views [36, 74] or video streams [26, 122] for embodied understanding, and shared 2D screens or UIs that are reproduced in MR [22, 109, 150]. In single user setting, designers can use this strategy when they want users to peek, adjust, or continue a 2D activity while keeping their visual attention within the MR workspace, such as by cloning a 2D screen or interface into MR [22, 109]. In collaborative settings, cloned first-person views provide external users with a straightforward way to follow the MR user's view or egocentric perspective [26, 36, 74]. However, such views may provide limited surrounding context, which can be complemented by a shared 3D scene [36] and asymmetric information (see 4.2.2).



Figure 5: Illustrations of design strategies in the Clone pattern. (C1) Synchronize contents: content displayed on a laptop is replicated in an MR window. (C2) Reuse input modalities: pen input performed on a tablet is routed to a virtual canvas in MR.

C2 Reuse input modalities. This strategy retains the input function of one device while routing it to another device or environment without changing its interaction logic. Because these inputs can often be used without shifting visual attention, they help users remain focused on the primary display. For example, non-immersive input can be reused in immersive applications through keyboards [63, 90], mice [105, 113], or tablet-and-pen input [31, 50]. This direction is useful for addressing limitations of current MR input, such as the lower efficiency or precision of controllers and hand gestures for certain tasks. Conversely, immersive input can be reused to operate non-immersive interfaces, such as using gestures [16, 161] or controllers [150] to interact with a phone or desktop. In both directions, designers can further leverage this strategy to avoid unnecessary switches to a different input modality. This can be useful for brief, lightweight cross-device adjustments where switching input methods would add unnecessary overhead.

4.2.2 ***Proxy****: Focused, Related, and Asymmetric.* In the Proxy pattern ($n = 84$), users remain visually focused on one device while another device supports the same task through an asymmetric input mapping or role relation. Unlike Clone, which preserves a shared representation or interaction logic across devices, Proxy assigns devices different but corresponding functions. For example, a handheld device may provide mapped control over spatial content in MR, while collaborative systems may present role-specific views of the same activity across devices. This distinction allows designers to tailor interaction, information, or user roles to the characteristics of each device. However, these differentiated functions require users to understand how their actions and representations correspond across devices. Such relationships can create cognitive gaps, making the bridging mechanisms described in Section 4.1.3 particularly relevant.

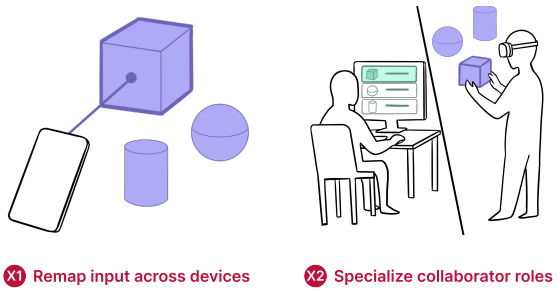


Figure 6: Illustrations of design strategies in the Proxy pattern. (X1) Remap input across devices: input from a smartphone is mapped to the manipulation of virtual objects in MR. (X2) Specialize collaborator roles: a desktop user guides an MR user through a virtual assembly task, with role-specific interfaces supporting task guidance on the desktop and direct manipulation in MR.

X1 Remap input across devices. The physical affordances or input capabilities of one device are mapped to interaction effects on another. Similar to C2, interaction with the secondary device relies on proprioceptive, or peripheral cues, allowing users to keep their gaze on the immersive view. These mappings draw on IMU data [29], touch input [57], or keyboard actions [23], and support interactions including 3D object manipulation [20, 24], raycasting [138, 162], tangible props [89, 144], and context-aware commands [53, 77, 108]. Within this broad design space, input types and interaction effects could be combined flexibly. We further observed several recurring pairings. Continuous inputs based on IMU or touchscreen movement were often paired with continuous spatial operations such as object manipulation [20, 29, 43, 53]. In contrast, discrete inputs, such as touchscreen gestures or touches at specific device regions, were more commonly mapped to invoking specific MR functions and shortcuts [57, 77]. When leveraging this strategy, one design consideration is how users understand the asymmetric mapping between input and effect. Designers can reduce this cognitive gap by reusing mappings established in everyday device interaction or gaming, allowing users to infer the new action–effect relationship through analogical transfer as described by SMT. In our corpus, we found no clear cases in which MR input served as a proxy for interaction on a non-immersive device.

X2 Specialize collaborator roles. In collaborative settings, different devices support distinct roles while contributing to the same shared task. These systems align each role with the strengths of its device. Non-immersive devices can support information-dense or overview-oriented work through text, dashboards, timelines, and other 2D interfaces, making them suitable for roles such as orchestrators or instructors [44, 61, 67, 123, 136]. Immersive devices, in contrast, provide spatial, embodied, and situated access to the task environment, supporting roles such as learners or explorers [44, 61, 67, 123, 136]. This specialization allows collaborative systems to combine the output and interaction strengths of heterogeneous devices. However, users must understand not only their own role, but also what information and actions are available to others and their effects on task state. Designers should therefore expose the collaborator actions and task states needed for coordination while minimizing unnecessary monitoring and attention shifts [44, 123].

4.2.3 ***Extend****: Fragmented, Related, and Symmetric.* In the Extend pattern ($n = 55$), users distribute task-related resources, interfaces, or interaction surfaces across devices. The devices support the same activity, but users need to visually attend across them because each device provides task-relevant content or interaction opportunities, which distinguishes Extend from Clone. The cross-device relationship remains cognitively symmetric: users can understand the extension as additional space, equivalent content, direct and continuous controls. Across the reviewed systems, Extend reflects two recurring directional motivations. When task resources are extended into MR, designers leverage the spatial flexibility of immersive space to overcome the fixed display area of physical screens, creating additional room for views, references, and controls. This flexibility, however, introduces a coordination challenge as resources must be arranged so that they remain easy to process and their relationships remain perceptually clear without making attention shifts across a broader workspace unnecessarily burdensome. In the reverse direction, designers extend UIs or interactive surfaces onto non-immersive devices to make use of stable, directly manipulable physical surfaces that can remain readily available alongside the immersive workspace.

E1 Expand workspace capacity. The system provides additional display space through user-anchored, device-anchored, or space-anchored virtual screens. This strategy also echos the design space of augmented displays [116]. The extended content may serve as independent resources [10, 101, 117], or supplements to the main display [46, 83, 103]. When extending into MR, the immersive space serves as a large workspace for placing virtual screens, references, or resources around physical devices, resembling multi-display work [114, 126] but with greater flexibility in spatial placement, scale, and anchoring [10, 42, 101]. One work extended the workspace to the non-immersive device, using the 2D tablet to provide an extended 3D view of the MR virtual object (i.e. a fish tank view) [157]. Such design also benefits from the Proxy to leverage the tablet surface and pen for efficient selection.

E2 Surface task-relevant content. This strategy is motivated by the need to reduce the effort of locating or monitoring information during an ongoing task. Systems support this by bringing frequently consulted resources, such as reference materials [151] or quick-access information [58], into more accessible positions, for example by relocating them into virtual space or placing them near the main task area. Although this strategy facilitates information retrieval, it shifts part of the interaction cost to the selection and organization of relevant content, especially when information needs change frequently. Designers therefore need lightweight mechanisms for moving content across devices and organizing surfaced information without introducing excessive interaction overhead or spatial clutter [58, 151].

E3 Distribute interactive surfaces. The system relocates interactable surfaces to the device best suited for interaction. Since many

MR controls remain familiar 2D UIs, such as menus [124, 144], tool palettes [34, 162], or shortcuts [93, 125], they can be redistributed across the device ecology: to tablets [34, 100], phones [145, 162], or keyboards [93, 125], whereas mobile controls can be reversely reproduced in MR as added-on interfaces [161, 162]. Placing these interfaces on physical devices can exploit direct and precise input surfaces while keeping frequently used controls readily accessible. Such devices are often positioned within easy reach or integrated into the user's working posture, reducing the physical transition required to access them during immersive work. This distribution may also be adaptive, with controls switching across surfaces according to context or user behavior, such as gaze [117], proximity [124], or interaction focus [145, 162]. Beyond UI controls, task-specific surfaces, such as a 3D plane or object surface, may be relocated to a tablet to leverage precise 2D drawing [10, 33, 91]. Compared with Clone, these UIs usually contain visual elements that users need to check before acting, which leads to fragmented attention. Compared with Proxy, the interaction remains symmetric when the UI elements preserve direct action-effect relationships, such as tapping a visible button or drawing on a tablet surface.

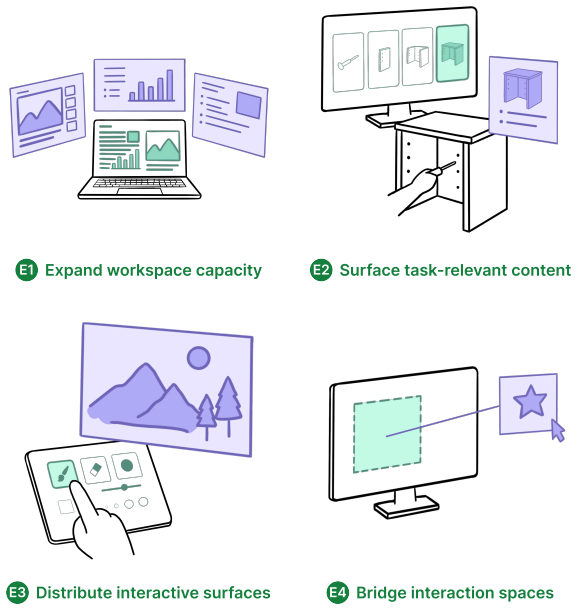


Figure 7: Illustrations of design strategies in the Extend pattern. (E1) Expand workspace capacity: additional virtual displays around a laptop extend the available display space beyond the physical screen. (E2) Surface task-relevant content: frequently consulted assembly instructions are brought closer to the user's ongoing workspace for easier access. (E3) Distribute interactive surfaces: drawing controls are placed on a tablet while the resulting effect is displayed on a virtual canvas in MR. (E4) Bridge interaction spaces: an object is dragged across the boundary between a desktop display and an adjacent MR window.

E4 Bridge interaction spaces. The system connects immersive and non-immersive spaces as adjacent or continuous interaction regions, so users can migrate content or continue an action without switching tools or interaction modes. For example, a cursor may move beyond the desktop screen into AR space [27], an object may be dragged through a screen border into a 3D region [28, 115, 161], or a brush or lasso trajectory may continue across physical and virtual surfaces [157]. The key feature is semantic continuity: although users may attend across devices, they can understand the cross-boundary action as one continuous manipulation rather than a new mapping. This continuity, however, is local to the boundary-crossing action and need not imply identical interaction semantics after the transition. Once content enters the target space, interaction may shift to its native modality or another device-specific mapping. Designers should therefore preserve the ongoing manipulation state while making such post-transition changes easy to anticipate, so users can continue the task without an abrupt interaction-mode switch [28, 115, 157].

4.2.4 ***Complement****: Fragmented, Related, and Asymmetric.* In the Complement pattern ($n = 46$), users coordinate immersive and non-immersive devices as different but interdependent parts of the same task. The devices provide complementary representations, task stages, or interaction roles. Each device contributes task-relevant but different information or action and thus users distribute visual attention across them and need to interpret how the parts relate. Unlike Extend, Complement does not simply expand or redistribute surfaces within the same action context; it divides the task into different but connected roles. Unlike Proxy, the asymmetry here is less about action mapping and more about coordinating attention and interpretation across task views and stages.

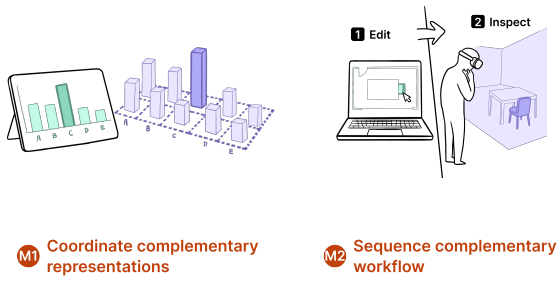


Figure 8: Illustrations of design strategies in the Complement pattern. (M1) Coordinate complementary representations: a tablet presents data as a 2D bar chart while MR presents the corresponding data as spatial 3D bars. (M2) Sequence complementary workflow: a user first edits a room layout on a laptop and then puts on the headset to inspect the resulting spatial arrangement in MR.

M1 Coordinate complementary representations. The system presents different but semantically connected representations of the same underlying object or data across devices. For example, one device may show a 2D chart, table, map, parameter panel, or overview, while the other shows a 3D spatial scene or 3D visualization, which

often appears in immersive analytics, sense-making, and data visualization systems [19, 60, 76, 106, 118, 129]. Some systems further couple these representations through spatial lenses, where a phone or tablet is registered as a movable window into the MR scene and displays 2D information (e.g. cross-section view) based on its position or orientation [85, 140]. The design goal is to help users integrate information that cannot be fully understood from one representation alone. However, the differences between representation can create cognitive gaps. Systems bridge this through spatial continuity [40, 76], bidirectional highlighting [21, 129], synchronized feedback [85, 140], or contextual cues [80].

M2 **Sequence complementary workflow chains.** The system allocates distinct but interdependent task stages across devices, forming a sequential workflow. The output of one stage becomes the input, reference, or object of manipulation for another stage. A common chain is 2D editing followed by 3D validation or editing, where the non-immersive device supports precise editing and the MR device supports in-situ inspection [51, 56, 57, 146] or spatial adjustment [12, 129, 139, 142, 157]. Other chains include moving from 2D drawing to 3D object creation [95, 119], or from coarse spatial selection to fine-grained planar refinement [96]. Rather than transferring an entire task phase, some systems coordinate tightly coupled steps within a single interaction [96, 157]. These workflows are complementary because each device supports a different form of reasoning or action within the same task. This strategy differs from the E3 since the latter usually remains part of the same immediate action context whereas this strategy supports more distinct stages of work, so users must connect representations and carry task meaning from one stage to the next.

4.2.5 ***Standby****: Focused and Unrelated.* In the Standby pattern ($n = 6$), users maintain focus on one device while another device remains available for a separate purpose, such as notification, status display, or awareness support. The tasks are unrelated because the standby device does not directly contribute to the user's current primary goal. One strategy is to S1 **push standby signals** into the active device, where notifications or lightweight status cues are presented within the user's current display space [43, 104]. In multi-user settings, another strategy is to S2 **externalize users' status** through the standby device. Rather than serving the MR user directly, the standby device allows nearby people or bystanders to interpret the MR user's otherwise hidden state, such as facial expressions or activity-related warnings [88, 152]. This supports social awareness around the MR activity while allowing the MR user to maintain focus on the primary task.

4.2.6 ***Partition****: Fragmented and Unrelated.* In the Partition pattern ($n = 3$), users divide attention across devices that support separate activities, making this the primary pattern for MR cross-device multitasking. Systems can facilitate this pattern by R1 **supporting parallel activities**. However, few of the reviewed systems were explicitly designed around this pattern. Instead, they typically enabled this through mechanisms also seen in other patterns. Some systems clone a phone screen into MR in the same way as in the Clone, but to enable potentially unrelated tasks, such as texting and photo-taking in the MR environment [3, 89]. In addition, a phone may serve as a spatial lens for inspecting the surrounding physical environment [156]. This limited design focus suggests a possible gap in current MR cross-device research, which we return to in Section 6.3.



Figure 9: Illustrations of design strategies in the Standby pattern. (S1) Push standby signals: a notification from a smartphone is surfaced in MR while the user is engaged in an unrelated MR task. (S2) Externalize user status: an MR user's busy status is displayed on a nearby tablet that is not involved in the current MR task, signaling their unavailability to people nearby.



Figure 10: Illustration of the design strategy in the Partition pattern. (R1) Support parallel activities: a VR user follows a virtual presentation while simultaneously playing an unrelated game on a smartphone.

4.3 Combinations Across Usage Patterns

We identified several systems that integrate multiple usage patterns. We further examined these cases to characterize how patterns were combined. 34 papers (14.3% of the corpus) introduced systems that meaningfully combined multiple usage patterns, as shown in Fig. 11. No particular pair showed strong systematic affinity, with lift values generally around or below 1. The three most frequent combinations were Clone–Proxy, Extend–Complement, and Proxy–Extend. Combinations involving Unrelated patterns were rare and did not form recurring pairwise structures; we therefore discuss them only in 4.3.1.

A recurring approach was to leverage Proxy as a flexible input layer within mixed-pattern designs. When used in this role, we identified two recurring design motivations: to supplement the input already supported by another pattern, or to provide an input

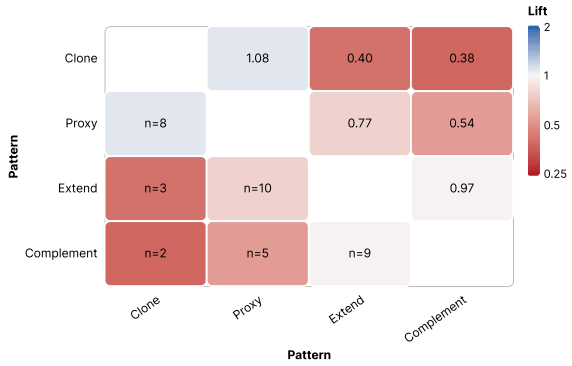


Figure 11: Pairwise combinations among the four Related usage patterns. The lower triangle shows the number of papers containing each pattern pair (n), and the upper triangle shows lift, representing observed relative to expected co-occurrence, with values above (below) 1 indicating more (less) co-occurrence than expected. Because a paper may contain more than two patterns, it can contribute to multiple pairwise counts; these counts are therefore not mutually exclusive.

method for the output introduced by another pattern. To supplement existing input, Proxy is often combined with Extend, where portable non-immersive devices such as phones and tablets can simultaneously host redistributed interface components (E3) and serve as proxy controllers (X1). For example, a phone can retain familiar touch input for interfaces redistributed into XR, while dynamically transitioning into a spatial controller as application contexts change [162]. This combination retains the intuitiveness of symmetric interaction with familiar 2D interfaces while introducing specialized asymmetric controls better suited to manipulating objects in 3D MR space. Beyond altering input behaviors at the task level [162], several systems exploit physical orientation and spatial proximity to dynamically coordinate the patterns. For example, a tablet may display redistributed UI elements when positioned vertically or flipped, while also being used as a tangible prop for selecting nearby 3D objects [57, 138]. Together, these adaptive mechanism illustrate a broader coordination strategy in which task and spatial context can mediate transitions between usage patterns.

For the second motivation, Proxy is frequently paired with output-oriented strategies from Extend, Complement, and Clone. In these combinations, the other pattern determines how information is presented, such as through complementary spatial presentation (M1) or workspace expansion (E1), while Proxy provides controllers or interaction shortcuts for frequent inspection and adjustment (X1). This arrangement appears particularly useful in applications involving dense information or repeated manipulation, such as immersive sense-making, extended workspaces, and data visualization [24, 42, 57, 99], where relying exclusively on menus or spatial input can impose substantial interaction overhead. For example, pen-based interaction such as rotation and movement along the z-axis can be remapped to rotate and scale virtual slides in MR, while gaze is used for selection, reducing the need to repeatedly access explicit menus [10]. A similar division occurs when Proxy is combined with Clone: the cloned view (C1) decouples visual output from the physical non-immersive device, whereas proxy interactions preserve convenient selection and predefined shortcuts, such as taking screenshots [109].

In Extend–Complement combinations, we identified two recurring design motivations. One is to support hybrid workflows across devices. These systems often combine a PC for dense or precise editing (M2) with a portable device that brings lightweight controls into the in-situ MR workflow (E3) [72, 95]. Thus, this combination enhances users' fine-tuning ability within the complementary workflow. Extend can also bridge the asymmetric stages of the complementary workflow, supporting a smoother transition from desktop to in-situ MR interaction by preserving the active task context [72]. For example, users can construct the main program structure on a PC, then use a phone for mobile parameter adjustment and convenient text or number entry while working in MR [72]. The second motivation is to combine an extended 2D workspace with complementary 3D representations, a configuration commonly found in 3D design and sensemaking systems. The non-immersive surface provides familiar and precise planar interaction, while MR expands the available display space for arranging data and interfaces (E1). Complementary representations (M1) then allow the same underlying data to be examined through both symmetric and asymmetric views [76, 119, 142]. Their spatial continuity helps users relate the two representations and reduces the cognitive gap between them.

A similar motivation appears in Clone–Proxy combination. In multi-user settings, many systems combine a shared first-person perspective (C1) with role-specific views available only to an external collaborator (X2) to enhance common ground [9, 32, 44, 123, 136]. For example, desktop or VR analysts can access a synchronized window into one another's current view while retaining platform-specific 2D or 3D representations [123], while tablet presenters can monitor a mirrored audience view alongside a top-down room view showing audience position and orientation [44]. Here, the same symmetric–asymmetric logic shifts from coordinating representations across devices to coordinating information across user roles, which Clone establishes a shared reference while Proxy preserves role-specific information or controls needed by individual collaborators.

When viewed through the lens of usage dimensions, these recurring combinations reveal a more systematic structure than the pattern pairs alone suggest. Clone–Proxy and Extend–Complement both combine symmetric and asymmetric patterns while preserving the same attention regime, allowing systems to retain a stable shared representation or familiar interaction while selectively introducing device- or role-specific capabilities where they are useful for current task. Proxy–Extend additionally crosses the attention dimension. In these cases, systems often coordinate the transition between focused and distributed device use through task state or interaction need, suggesting that combinations across attention regimes require an explicit mechanism for determining when each pattern becomes appropriate. Overall, the observed combinations therefore reflect different ways of coordinating stability and specialization, as well as focused and fragmented interaction.

4.3.1 Less Common and Incidental Combinations. Beyond the recurring combinations discussed above, we observed several less common or less tightly integrated forms of pattern combinations. First, several combinations appeared in only one or two systems but nevertheless exhibited a clear functional rationale. A few systems are designed for both single-user and multi-user settings, which leads to a part of Complement–Proxy [12, 62], Complement–Clone [139], and Extend–Clone [50] combination. In single-user workflow (M2,E1), the system allows an individual user to switch between devices. By further enabling the real-time data synchronize, they can also support two users to control the different devices to collaborate around a role-specific view (X2) or shared scene (C1). These systems provide a potential direction to expand the future MR system from single usage to multi-users. Second, some systems combined multiple patterns toward the same overall task goal, but with relatively loose coupling between them. In these cases, the patterns addressed separate local needs within the task, without one pattern directly enabling or coordinating the other. Such combinations often took the form of pairing multiple input-related functions (e.g., E1+M2 [157], C2+E3 [34]), combining independently useful input and output functions (e.g., C2+M1 [105]), or mixing on-demand unrelated usage (R1 [156], S1 [43]). Finally, additional combinations could arise if patterns outside a paper's primary design motivation were also counted. As mentioned in 3.3.1, we excluded incidental uses during pattern coding to avoid inflation. A common pattern for these cases to combine is Clone as C2 retains the basic interaction mechanism of input devices. Examples includes cloning the laptop input at extended workspace (E1) [101], supporting basic keyboard typing or pen drawing as alternative when proposing X1 [34, 109] and E3 [93, 125]. These cases represent forms of pattern coexistence that are present in the systems but intentionally omitted from the quantitative combination analysis.

4.4 Boundary Cases

We encountered several boundary cases when applying either our coding framework or eligibility criteria, concerning how included systems should be classified or whether closely related systems should be included.

Some included systems used additional devices mainly as contextual sensors, such as camera-based recognition [6, 68, 130], microphone or phone-call input [48], and smartwatch heart-rate sensing [1, 111, 158]. Although users may not explicitly operate or attend to these devices, we retained them because their outputs directly participated in the cross-device usage episode and affected the interaction on the other device. For cognitive symmetry, we coded whether the sensed information preserved a direct semantic correspondence with its resulting effect or was transformed into a different interpretation. These cases therefore represent a more passive form of cross-device coordination rather than a separate class outside our framework.

We also encountered devices that closely resembled common consumer devices but incorporated specialized affordances, such as a tablet with a transparent display [73], as well as simulated devices such as virtual tablets [160]. Although their interaction concepts may transfer to more common cross-device settings, we excluded them under our existing criteria to avoid case-by-case judgments about how much specialization remained sufficiently transferable. Such systems may nevertheless help extend the framework as emerging device capabilities become more commonplace.

5 Usage of MRxD Framework

This section shows how MRxD can be used in practice through example system analysis, concept generation, and a formative workshop on early-stage MR cross-device design.

5.1 Example Applications

Analyzing existing systems. MaRginalia supports classroom note-taking by projecting a tablet interface into MR to reduce visual attention shifts between the tablet and the lecture [109]. Its cloned tablet screen and familiar pen-based note-taking fit the **Clone** pattern, while pen interactions that map the pen to an MR cursor can be analyzed as **Proxy**. This pattern analysis clarifies the system's rationale: its interaction techniques consistently support focused attention. By mapping Proxy designs into MRxD, we further identify a latent trade-off, where some Proxy mappings remain inconsistent, such as using pen-down for button activation but pen squeezing for manual snapshot capture. Although the paper does not frame this as a mapping issue, a participant's workaround of capturing transcripts by tapping a dot suggests that inconsistent asymmetric mappings may introduce local friction.

Generating new design concepts. Consider a desktop-to-MR generative design workflow where a designer edits prompts and parameters on a desktop and inspects generated 3D objects in MR. MRxD characterizes this workflow as **Complement**: both devices support the same authoring task, but attention is fragmented between desktop editing and MR inspection, and the relationship is cognitively asymmetric because text prompts and parameters differ from 3D outputs. This diagnosis suggests bridging strategies that make the asymmetry easier to interpret. For example, hovering over a prompt phrase or parameter on the desktop could highlight the corresponding object component or material in MR. The workflow could also introduce **Proxy** for quick inspection, such as using a phone's IMU to manipulate the MR object through lightweight motion input.

5.2 Formative Design Workshop

To understand how designers apply MRxD in practice and collect feedback, we conducted a formative one-on-one design workshop with six HCI designers (2 female, 4 male; ages 23–29), all with AR/VR-related design or research experience, including four academic researchers and two industry engineers. The workshop adapted three stages from the Google design sprint [69]—map, diverge, and decide—to support problem framing, design exploration, and reflective selection. In the **map** stage, participants chose whichever of the two design cases they felt would better support their ideation and spent about five minutes identifying user needs and cross-device design challenges: VR classroom learning with a headset, tablet, pen, and laptop ($n = 2$), or AR home decoration with an AR headset, smartwatch, smartphone, and desktop ($n = 4$). The cases provided concrete design contexts and were not intended for comparison. In the **diverge** stage, participants first spent ten minutes designing from their own knowledge while thinking aloud

and sketching. We then introduced MRxD, including its three dimensions, six usage patterns, and related design strategies, after which participants spent another ten minutes revisiting and extending their ideas with the framework when appropriate. In the **decide** stage, participants spent about ten minutes refining their ideas and reflecting on coordination trade-offs using MRxD. The workshop ended with a 15-minute semi-structured interview on MRxD's usefulness and other feedback. We collected participants' initial and revised design artifacts, session notes, and interview transcripts. We analyzed these materials using reflexive thematic analysis [17]. The study was approved by the local IRB.

5.2.1 Results. Our results show that participants used MRxD in three main ways. **Revealing overlooked needs and design blind spots.** After MRxD was introduced, participants added elements that were absent from their initial designs, including multi-user support (P1, P2, P5, P6) and missing input or output components (P1, P3, P4, P6). Participants described the design strategies as directly useful (P2–P4), and P1 described usage patterns as a "connector between usage behavior and design strategies" that "helped avoid redundant elements and find supplemental designs."

Encouraging capability-level decomposition of devices. Participants initially tended to assign whole devices to broad roles, such as computer for editing, AR for display, or phone for note-taking. After using MRxD, several decomposed devices into capabilities such as input, output, sensing, and transfer. For example, P1 and P3 moved toward reusing keyboard input inside VR, while P6 shifted from treating the phone as a note-taking tool to using it as a proxy controller for placing furniture in MR.

Prompting dimension-guided design and reasoning. The three dimensions helped participants reason about coordination choices. User Attention led participants to consider whether users should remain focused in MR or move across devices, as when P6 proposed cloning phone notes onto an AR virtual screen to reduce repeated glances at the phone. Cognitive Symmetry prompted participants to move beyond direct or mirrored device relationships and explore asymmetric mappings, such as using a pen (P1) or smartphone (P4) as a laser pointer for selection. P4 and P5 further suggested that this dimension-guided reasoning could be better supported through a more interactive format, such as a website that allows designers to browse usage patterns, compare dimensions, and access relevant examples during ideation.

6 Discussion

6.1 MRxD as a Design Scaffold

MRxD supports design reasoning by connecting intended user activities with concrete choices about cross-device coordination. The formative workshop illustrates how designers used this connection in early-stage design. Participants used the framework to identify support that was absent from their initial concepts, reconsider broad device-level roles in terms of specific capabilities such as input and output, and reflect on coordination choices through the usage dimensions. These observations suggest two design roles of MRxD. First, the framework provides an analytic structure for examining whether a proposed design aligns with its intended design goal. For example, a design intended to maintain an immersive and focused experience may still introduce interactions that require users to visually attend to a secondary device. Mapping the design back onto the usage dimensions makes such tensions explicit and can prompt designers to reconsider the design. This highlights the value of the framework in making design reasoning explicit, allowing underlying assumptions, trade-offs, and inconsistencies to be more readily examined and revised [87]. In addition, patterns and strategies not represented in the current design can similarly suggest alternative design possibilities. When multiple patterns appear within a workflow, MRxD can further help designers reason about how they interact and what coordination trade-offs their combination may introduce (Section 4.3). Second, MRxD can serve as a generative scaffold that connects intended usage behavior to concrete design strategies. Designers can begin with a usage behavior, identify the corresponding pattern, and use its associated strategies as starting points for design exploration. Such structured prompts may be particularly useful for less-experienced designers by helping them consider alternatives beyond their initial concepts [11, 154]. In the future, because these patterns are defined by usage rather than specific device configurations, MRxD may also accommodate emerging devices and interaction techniques according to the cross-device relationships they create, with recurring new realizations further extending the strategy set. This allows MRxD to remain useful as a design scaffold even as MR devices and interaction techniques evolve.

6.2 Designing Along MRxD Dimensions

For systematic coding across a large corpus, we operationalized the three MRxD dimensions as two-sided distinctions. In design use, however, these dimensions can also be treated as spectra, allowing designers to steer a system toward more focused or fragmented attention, or toward more symmetric or asymmetric cognition, by implementing specific designs. This also allows the analyze of user behavior along the spectra, leading to more detailed design for the middle state.

6.2.1 Shifting User Attention. As discussed in Section 4.1.1, fragmented attention often resulted from distributed resources or controls that required visual access across devices. To shift user attention toward the focused end, strategies in Clone and Proxy suggest a common principle: *reduce the need to visually confirm the secondary device*. Continuous input and stable, learnable, or tactile controls can support this by allowing users to act while monitoring feedback in MR, whereas visually locating controls or confirming actions (e.g. E3, M2) draws attention back to the device. This is consistent with prior work showing lower visual demands for gestural than soft-button interaction [15]. Our multi-pattern results further reveal input–output decoupling as another way to support focused attention. Designers can separate visual output from the physical device while retaining that device for proxy input, combining Clone and Proxy strategies to keep visual attention in MR. For example, a tablet can serve as a trackpad or stable input surface while the corresponding cursor or guidance appears near the MR task [109]. When visual access across devices is unavoidable, as often occurs in Extend, Complement, and Partition, designers can instead manage fragmentation by reducing cost of attention switches. AR resumption cues can be used as contextual guidance

that facilitates returning to an interrupted task [7]. The workflow also can be optimized to shift the attention closer to the focused end by reducing the switch frequency.

Attention can also shift along the spectrum as user state or momentary task demands change, even within the same design. One source is familiarity: users may initially inspect a device to locate controls, but require less visual guidance as location memory and motor control develop with experience [64]. Systems could therefore provide stronger cross-device cues during early use and gradually reduce this scaffolding. Attention may also shift temporarily when a task requires greater visual precision, such as locating an infrequently used control or performing a precision-critical action. These foreseeable shifts suggest that systems can adapt their support accordingly, providing visual guidance when confirmation is needed while otherwise allowing attention to remain focused on MR.

6.2.2 Managing Cognitive Symmetry. At the system level, designers can build on the bridging mechanisms synthesized from our corpus to manage Cognitive Symmetry. In particular, cues used for system alignment between representations can also make asymmetric interaction mappings more interpretable. For example, when different touch regions on a Proxy device are mapped to distinct MR commands [57], labels or icons placed at these regions can make the mapping directly observable. However, when such cues require users to inspect another device or representation, reducing the cognitive gap may introduce fragmented attention, revealing a potential trade-off between usage dimensions.

Designers can also reduce cognitive gaps by preserving continuity across states and grounding new mappings in familiar schemas. SeamlessVR showed that continuously transforming an immersive visualization into its corresponding screen view reduced the cognitive effort required to connect the two representations [82]. A similar principle appears in our observations, where Extend bridges asymmetric stages of a Complement workflow by preserving the active task context as users transition from desktop to in-situ MR interaction [72], echoing broader work on transitional interfaces. Familiarity provides another form of bridging: consistent with Structure-Mapping Theory [41], designers can favor mappings that preserve familiar relational schemas, such as treating a phone as a controller [162] or a pen as a pointer [78]. More specialized mappings that depart from familiar device roles [57, 138] may instead require stronger feedback or guidance while users form the corresponding mental model.

Our E4 findings further suggest that Cognitive Symmetry can be local to an interaction state rather than stable throughout a workflow. A cross-boundary action may remain directly interpretable during transfer, yet a subsequent interaction may become asymmetric once this interaction ends and adopts the target space's device-specific mapping [28, 115, 157]. Designers should therefore consider how easily users can carry their interpretation across such state transitions, not only the symmetry of each state in isolation. Similar shifts can arise in adaptive systems that redistribute interfaces based on contextual signals such as gaze or distance [117, 124]. Such adaptation may improve contextual fit, but users must understand both the resulting state and why the correspondence changed; prior work similarly shows that unpredictable interface adaptations can cause confusion and relearning costs [141]. Designers should therefore consider not only the Cognitive Symmetry of each state, but also whether users can carry their interpretation across transitions between them.

6.3 Research Opportunities

Unrelated patterns remain underexplored in current MR cross-device research, partly due to our search strategy and the field's focus on primary MR tasks. Nevertheless, unrelated usage remains important because users may still need to handle secondary activities while working in MR. Future work could explore bridging mechanisms that help users remain aware of activities across devices and smoothly switch between related and unrelated tasks. For example, prior systems have used visual focus to infer whether users are attending to an unrelated phone task or using the phone as an MR controller [43]. Future systems could extend this direction through activity awareness, visual cues, and adaptive reminders that support smoother suspension and resumption.

Another opportunity is to further examine interaction flows from MR to non-immersive devices. Echoing the MR-HUI review [55], most MR cross-device systems focus on integrating non-immersive devices into MR, while fewer systems explore the reverse direction. Existing examples include gesture or controller-as-mouse interactions for non-immersive devices [16, 150, 161], data migration across device boundaries [115], and tablet-based fish-tank views for inspecting 3D content [157]. Future work could more systematically explore this direction, for example, to consider how MR input modalities can serve as proxy inputs for non-immersive devices. Existing MR interaction vocabularies, such as gaze/gesture targeting [38], palm-based interfaces [49], and mid-air gesture sets [70], may provide a basis for such mappings, but they have rarely been examined as MR-to-non-immersive Proxy strategies.

MRxD also suggests the need for configurable design toolkits. Many strategies contain recombinable variants, such as input sources, output targets, mapping rules, and bridging cues, yet current systems often implement them as one-off designs. Inspired by prior cross-reality toolkits [47], an MRxD-based toolkit could help designers configure these components and examine how a design behaves across the framework dimensions.

6.4 Limitations

We identify several limitations of this study. First, our search strategy and review scope may not capture the full landscape of MR cross-device research. Despite accounting for fragmented terminology and including device-specific keywords, relevant work may still be missed when its cross-device characteristics are not reflected in searchable metadata (e.g., Friedl-Knirsch et al. [39]). In addition, our focus on HMD-based MR with common non-immersive devices excludes adjacent settings such as handheld MR, spatial AR, and systems using specialized devices. Our exclusion of papers of four pages or fewer may omit relevant short-form contributions. MRxD should therefore not be interpreted as an account of all MR cross-device systems or design strategies, and future work could examine how the framework extends to these settings.

Second, developing the MRxD dimensions involved interpretive analytical decisions and is therefore not fully reproducible in the

sense that another research team, or a different initial development sample, might derive a different but plausible set of dimensions from the same corpus. The initial 20% development corpus may also have made some dimensions more salient and thereby shaped the subsequent analytical trajectory. We sought to mitigate this dependence by treating the initial framework as provisional rather than fixed and repeatedly testing and refining the dimensions against the remaining corpus. MRxD therefore should not be interpreted as the uniquely optimal decomposition of the design space, but as an empirically grounded and parsimonious abstraction developed to support the analytical goals of this study.

Third, MRxD operationalizes each usage dimension as a binary distinction, although the underlying phenomena may lie on a spectrum or change over time. This abstraction supports parsimonious comparison across the corpus but may compress intermediate cases and corresponding design strategies.

Fourth, our analysis retained usage episodes central to each paper's design motivation or contribution rather than exhaustively coding all supported interactions. This focuses the framework on patterns actively explored in prior work, but makes episode selection partly dependent on authors' framing and our interpretation. Accordingly, pattern distributions reflect research emphasis rather than behavioral prevalence, and secondary patterns or combinations may be underrepresented. Future work could systematically examine all instantiated episodes within systems or conduct in-the-wild data collection to provide a more comprehensive account of pattern co-occurrence and prevalence.

7 Conclusion

This paper contributes MRxD, a usage-pattern framework for MR cross-device interaction. Based on a systematic review of 238 papers, MRxD characterizes usage episodes through User Attention, Task Relatedness, and Cognitive Symmetry, and identifies six recurring patterns with associated design strategies. A formative design workshop further suggests how the framework can support design reflection and concept generation. This paper offers MR community valuable insights for analyzing and designing cross-device systems that coordinate HMDs with non-immersive devices.

References

[1] Sarfraz Ahmed, Joannes Sam Mertens, and Laura Galluccio. 2025. A Digital Twin Framework for Human Organs modeling and forecasting in the Metaverse. In *2025 3rd International Conference on Intelligent Metaverse Technologies & Applications (iMETA)* (Durbovnic, Croatia). IEEE, 39–44. doi:10.1109/imeta66706.2025.11306455

[2] Oscar Alvarado, Nyi Nyi Htun, Yucheng Jin, and Katrien Verbert. 2022. A Systematic Review of Interaction Design Strategies for Group Recommendation Systems. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2, Article 271 (Nov. 2022), 51 pages. doi:10.1145/3555161

[3] Tatsuya Amano, Shugo Kajita, Hirozumi Yamaguchi, Teruo Higashino, and Mineo Takai. 2018. Smartphone Applications Testbed Using Virtual Reality. In *Proceedings of the 15th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services* (New York, NY, USA) *(MobiQuitous '18)*. Association for Computing Machinery, New York, NY, USA, 422–431. doi:10.1145/3286978.3287028

[4] Jonas Auda, Uwe Gruenefeld, Sarah Faltaous, Sven Mayer, and Stefan Schneegass. 2023. A Scoping Survey on Cross-reality Systems. *Comput. Surveys* 56, 4, Article 83 (2023), 38 pages. doi:10.1145/3616536

[5] Md Tanvir Islam Aumi, Sidhant Gupta, Mayank Goel, Eric Larson, and Shwetak Patel. 2013. DopLink: Using the Doppler Effect for Multi-device Interaction. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Zurich, Switzerland) *(UbiComp '13)*. Association for Computing Machinery, New York, NY, USA, 583–586. doi:10.1145/2493432.2493515

[6] Mihai Bâce, Gábor Sörös, Sander Staal, and Giorgio Corbellini. 2017. HandshakAR: Wearable Augmented Reality System for Effortless Information Sharing. In *Proceedings of the 8th Augmented Human International Conference* (Silicon Valley, California, USA) *(AH '17)*. Association for Computing Machinery, New York, NY, USA, 1–5. doi:10.1145/3041164.3041203

[7] Kilian L. Bahnsen, Lucas Tiemann, Lucas Plabst, and Tobias Grundgeiger. 2024. Augmented Reality Cues Facilitate Task Resumption after Interruptions in Computer-Based and Physical Tasks. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3613904.3642666

[8] Sunyoung Bang, Hyunjin Lee, Seo Young Oh, and Woontack Woo. 2025. AReading with Smartphones: Understanding the Trade-offs between Enhanced Legibility and Display Switching Costs in Hybrid AR Interfaces. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3706598.3713879

[9] Jalal Safari Bazargani, Phil Joong Kim, and Soo-Mi Choi. 2025. Beyond the Headset: Enabling Group Collaboration in Cross-Reality Environments. In *2025 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Daejeon, Korea, Republic of). IEEE, 6–10. doi:10.1109/ismar-adjunct68609.2025.00008

[10] Verena Biener, Travis Gesslein, Daniel Schneider, Felix Kawala, Alexander Otte, Per Ola Kristensson, Michel Pahud, Eyal Ofek, Cuauhtli Campos, Matjaz Kljun, Klen Copic Pucihar, and Jens Grubert. 2022. PoVRPoint: Authoring Presentations in Mobile Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 28, 5 (2022), 2069–2079. doi:10.1109/tvcg.2022.3150474

[11] Nis Bornoe, Anders Bruun, and Jan Stage. 2016. Facilitating Redesign with Design Cards: Experiences with Novice Designers. In *Proceedings of the 28th Australian Conference on Computer-Human Interaction*. Association for Computing Machinery, 452–461. doi:10.1145/3010915.3010921

[12] Marcel Borowski, Jens Emil Sloth Grønbæk, Peter W. S. Butcher, Panagiotis D. Ritsos, Clemens Nylandsted Klokmose, and Niklas Elmqvist. 2025. Spatialstrates: Cross-Reality Collaboration through Spatial Hypermedia. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology* (Busan, Republic of Korea) *(UIST '25)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3746059.3747708

[13] Sidney Bovet, Aidan Kehoe, Katie Crowley, Noirin Curran, Mario Gutierrez, Mathieu Meisser, Damien O Sullivan, and Thomas Rouvinez. 2018. Using Traditional Keyboards in VR: SteamVR Developer Kit and Pilot Game User Study. In *2018 IEEE Games, Entertainment, Media Conference (GEM)* (Galway). IEEE, 1–9. doi:10.1109/gem.2018.8516449

[14] Evren Bozgeyikli and Victor Gomes. 2022. Googly Eyes: Displaying User's Eyes on a Head-Mounted Display for Improved Nonverbal Communication. In *Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play* (Bremen, Germany) *(CHI PLAY '22)*. Association for Computing Machinery, New York, NY, USA, 253–260. doi:10.1145/3505270.3558348

[15] Andrew Bragdon, Eugene Nelson, Yang Li, and Ken Hinckley. 2011. Experimental Analysis of Touch-Screen Gesture Designs in Mobile Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 403–412. doi:10.1145/1978942.1979000

[16] Eugenie Brasier, Emmanuel Pietriga, and Caroline Appert. 2021. AR-enhanced Widgets for Smartphone-centric Interaction. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction* (Toulouse & Virtual, France) *(MobileHCI '21)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3447526.3472019

[17] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. doi:10.1191/1478088706qp063oa

[18] Frederik Brudy, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmose, and Nicolai Marquardt. 2019. Cross-Device Taxonomy: Survey, Opportunities and Challenges of Interactions Spanning Across Multiple Devices. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–28. doi:10.1145/3290605.3300792

[19] Simon Butscher, Sebastian Hubenschmid, Jens Müller, Johannes Fuchs, and Harald Reiterer. 2018. Clusters, Trends, and Outliers: How Immersive Technologies Can Facilitate the Collaborative Analysis of Multidimensional Data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal, QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173664

[20] Wolfgang Büschel, Annett Mitschick, Thomas Meyer, and Raimund Dachselt. 2019. Investigating Smartphone-based Pan and Zoom in 3D Data Spaces in Augmented Reality. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services* (Taipei, Taiwan) *(MobileHCI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3338286.3340113

[21] Chen Chen, Matin Yarmand, Varun Singh, Michael V. Sherer, James D. Murphy, Yang Zhang, and Nadir Weibel. 2022. VRContour: Bringing Contour Delineations of Medical Structures Into Virtual Reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (Singapore, Singapore). IEEE, 64–73. doi:10.1109/ismar55827.2022.00020

[22] Sibo Chen, Junce Wang, Santiago Guerra, Neha Mittal, and Soravis Prakkamakul. 2019. Exploring Word-gesture Text Entry Techniques in Virtual Reality. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland, UK) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3290607.3312762

[23] Subramanian Chidambaram, Alex C Williams, Min Bai, Satyugjit Virk, Patrick Haffner, Matthew Lease, and Erran Li. 2024. Annorama: Enabling Immersive At-Desk Annotation Experiences in Virtual Reality with 3D Point Cloud Dioramas. In *ACM Symposium on Spatial User Interaction* (Trier, Germany) *(SUI '24)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3677386.3682081

[24] Neil Chulpongsatorn, Wesley Willett, and Ryo Suzuki. 2023. HoloTouch: Interacting with Mixed Reality Visualizations Through Smartphone Proxies. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3544549.3585738

[25] Haeyong Chung and Chris North. 2018. SAViL: Cross-display Visual Links for Sensemaking in Display Ecologies. *Personal and Ubiquitous Computing* 22, 2 (2018), 409–431. doi:10.1007/s00779-017-1091-4

[26] Marina Cidota, Stephan Lukosch, Dragos Datcu, and Heide Lukosch. 2016. Workspace Awareness in Collaborative AR using HMDs: A User Study Comparing Audio and Visual Notifications. In *Proceedings of the 7th Augmented Human International Conference 2016* (Geneva, Switzerland) *(AH '16)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/2875194.2875204

[27] Robbe Cools, Matt Gottsacker, Adalberto Simeone, Gerd Bruder, Greg Welch, and Steven Feiner. 2022. Towards a Desktop-AR Prototyping Framework: Prototyping Cross-Reality Between Desktops and Augmented Reality. In *2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Singapore, Singapore). IEEE, 175–182. doi:10.1109/ismar-adjunct57072.2022.00040

[28] Robbe Cools, Inne Maerevoet, Matt Gottsacker, and Adalberto L. Simeone. 2025. Comparison of Cross-Reality Transition Techniques Between 3D and 2D Display Spaces in Desktop-AR Systems. *IEEE Transactions on Visualization and Computer Graphics* 31, 5 (2025), 2798–2805. doi:10.1109/tvcg.2025.3549907

[29] Murilo Santos de Castro, Pedro Raphael Inácio Gomes, and Thamer Horbylon Nascimento. 2023. Exploring Interaction in a Virtual Music Studio through Gesture Recognition on Smartwatches and HMD Devices. In *Symposium on Virtual and Augmented Reality* (Rio Grande, Brazil) *(SVR '23)*. Association for Computing Machinery, New York, NY, USA, 284–288. doi:10.1145/3625008.3625049

[30] Alexandre Demeure, Gaëlle Calvary, Jean-Sebastien Sottet, and Jean Vanderdonkt. 2005. A Reference Model for Distributed User Interfaces. In *Proceedings of the 4th International Workshop on Task Models and Diagrams* (Gdansk, Poland) *(TAMODIA '05)*. Association for Computing Machinery, New York, NY, USA, 79–86. doi:10.1145/1122935.1122952

[31] Seyed Amir Ahmad Didehkhorshid, Siju Philip, Elaheh Samimi, and Robert J. Teather. 2020. Text Input in Virtual Reality Using a Tracked Drawing Tablet. In *Proceedings of the 20th International Conference on Human-Computer Interaction, Part II (Lecture Notes in Computer Science, Vol. 12425)*. Springer International Publishing, 314–329. doi:10.1007/978-3-030-59990-4_24

[32] Morgan Diederich, Jina Kang, Taehyun Kim, and Robb Lindgren. 2021. Developing an In-Application Shared View Metric to Capture Collaborative Learning in a Multi-Platform Astronomy Simulation. In *LAK21: 11th International Learning Analytics and Knowledge Conference* (Irvine, CA, USA) *(LAK21)*. Association for Computing Machinery, New York, NY, USA, 173–183. doi:10.1145/3448139.3448156

[33] Tobias Drey, Jan Gugenheimer, Julian Karlbauer, Maximilian Milo, and Enrico Rukzio. 2020. VRSketchIn: Exploring the Design Space of Pen and Tablet Interaction for 3D Sketching in Virtual Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376628

[34] Tobias Drey, Nico Rixen, Julian Karlbauer, and Enrico Rukzio. 2024. VRCreatIn: Taking In-Situ Pen and Tablet Interaction Beyond Ideation to 3D Modeling Lighting and Texturing. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia* (Stockholm, Sweden) *(MUM '24)*. Association for Computing Machinery, New York, NY, USA, 24–35. doi:10.1145/3701571.3701580

[35] Daniel Enriquez, Wai Tong, Chris North, Huamin Qu, and Yalong Yang. 2024. Evaluating Layout Dimensionalities in PC+VR Asymmetric Collaborative Decision Making. *Proceedings of the ACM on Human-Computer Interaction* 8, ISS (2024), 112–132. doi:10.1145/3698130

[36] Arthur Fages, Cédric Fleury, and Theophanis Tsandilas. 2022. Understanding Multi-View Collaboration between Augmented Reality and Remote Desktop Users. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–27. doi:10.1145/3555607

[37] Steven Feiner and Ari Shamash. 1991. Hybrid user interfaces: breeding virtually bigger interfaces for physically smaller computers. In *Proceedings of the 4th Annual ACM Symposium on User Interface Software and Technology* (Hilton Head, South Carolina, USA) *(UIST '91)*. Association for Computing Machinery, New York, NY, USA, 9–17. doi:10.1145/120782.120783

[38] Anna Maria Feit, Shane Williams, Arturo Toledo, Ann Paradiso, Harish Kulkarni, Shaun Kane, and Meredith Ringel Morris. 2017. Toward Everyday Gaze Input: Accuracy and Precision of Eye Tracking and Implications for Design. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 1118–1130. doi:10.1145/3025453.3025599

[39] Judith Friedl-Knirsch, Christian Stach, Fabian Pointecker, Christoph Anthes, and Daniel Roth. 2024. A Study on Collaborative Visual Data Analysis in Augmented Reality with Asymmetric Display Types. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2633–2643. doi:10.1109/TVCG.2024.3372103

[40] Willis Fulmer, Tahir Mahmood, Zhongyu Li, Shaoting Zhang, Jian Huang, and Aidong Lu. 2019. ImWeb: cross-platform immersive web browsing for online 3D neuron database exploration. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (Marina del Ray, California) *(IUI '19)*. Association for Computing Machinery, New York, NY, USA, 367–378. doi:10.1145/3301275.3302319

[41] Dedre Gentner. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7, 2 (1983), 155–170. doi:10.1016/S0364-0213(83)80009-3

[42] Travis Gesslein, Verena Biener, Philipp Gagel, Daniel Schneider, Per Ola Kristensson, Eyal Ofek, Michel Pahud, and Jens Grubert. 2020. Pen-based Interaction with Spreadsheets in Mobile Virtual Reality. In *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (Porto de Galinhas, Brazil). IEEE, 361–373. doi:10.1109/ismar50242.2020.00063

[43] Eric J. Gonzalez, Ishan Chatterjee, Mar Gonzalez-Franco, Andrea Colaço, and Karan Ahuja. 2024. Intent-driven input device arbitration for XR. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–5. doi:10.1145/3613905.3650758

[44] Matt Gottsacker, Mengyu Chen, David Saffo, Feiyu Lu, Benjamin Lee, and Blair MacIntyre. 2025. Examining the Effects of Immersive and Non-Immersive Presenter Modalities on Engagement and Social Interaction in Co-located Augmented Presentations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3706598.3713346

[45] Matt Gottsacker, Mengyu Chen, David Saffo, Feiyu Lu, and Blair MacIntyre. 2023. Hybrid User Interface for Audience Feedback Guided Asymmetric Immersive Presentation of Financial Data. In *2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Sydney, Australia). IEEE, 199–204. doi:10.1109/ismar-adjunct60411.2023.00046

[46] Jens Grubert, Matthias Heinisch, Aaron Quigley, and Dieter Schmalstieg. 2015. MultiFi: Multi Fidelity Interaction with Displays On and Around the Body. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 3933–3942. doi:10.1145/2702123.2702331

[47] Uwe Gruenefeld, Jonas Auda, Florian Mathis, Stefan Schneegass, Mohamed Khamis, Jan Gugenheimer, and Sven Mayer. 2022. VRception: Rapid Prototyping of Cross-Reality Systems in Virtual Reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 611, 15 pages. doi:10.1145/3491102.3501821

[48] Yunqi Guo, Jinghao Zhao, Boyan Ding, Congkai Tan, Weichong Ling, Zhaowei Tan, Jennifer Miyaki, Hongzhe Du, and Songwu Lu. 2023. Sign-to-911: Emergency Call Service for Sign Language Users with Assistive AR Glasses. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking* (Madrid, Spain) *(ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3570361.3613260

[49] Sean G. Gustafson, Bernhard Rabe, and Patrick M. Baudisch. 2013. Understanding palm-based imaginary interfaces: the role of visual and tactile cues when browsing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 889–898. doi:10.1145/2470654.2466114

[50] Rorik Henrikson, Bruno Araujo, Fanny Chevalier, Karan Singh, and Ravin Balakrishnan. 2016. Multi-Device Storyboards for Cinematic Narratives in VR. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) *(UIST '16)*. Association for Computing Machinery, New York, NY, USA, 787–796. doi:10.1145/2984511.2984539

[51] Bridger Herman, Francesca Samsel, Annie Bares, Seth Johnson, Greg Abram, and Daniel F. Keefe. 2020. Printmaking, Puzzles, and Studio Closets: Using Artistic Metaphors to Reimagine the User Interface for Designing Immersive Visualizations. In *2020 IEEE VIS Arts Program (VISAP)* (Salt Lake City, UT, USA). IEEE, 19–28. doi:10.1109/visap51628.2020.00009

[52] Leila Homaeian, Nippun Goyal, James R. Wallace, and Stacey D. Scott. 2017. Investigating Communication Grounding in Cross-Surface Interaction. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces* (Brighton, United Kingdom) *(ISS '17)*. Association for Computing Machinery, New York, NY, USA, 348–353. doi:10.1145/3132272.3132282

[53] Jinbin Huang, Shuang Liang, Qi Xiong, Yu Gao, Chao Mei, Yi Xu, and Chris Bryan. 2022. SPARVIS: Combining Smartphone and Augmented Reality for Visual Data Analytics. In *2022 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Singapore, Singapore). IEEE, 111–117. doi:10.1109/ismar-adjunct57072.2022.00030

[54] Jinbin Huang, Jonathan D. Plasencia, Dianna M. E. Bardo, Nicholas C. Huber, Erik G. Ellsworth, Steven D. Zangwill, and Chris Bryan. 2021. Phoenix Virtual Heart: A Hybrid VR-Desktop Visualization System for Cardiac Surgery Planning and Education. In *2021 IEEE Workshop on Visual Analytics in Healthcare (VAHC)* (New Orleans, LA, USA). IEEE, 36–40. doi:10.1109/vahc53616.2021.00012

[55] Sebastian Hubenschmid, Marc Satkowski, Johannes Zagermann, Julián Méndez, Niklas Elmqvist, Steven Feiner, Tiare Feuchtner, Jens Emil Grønbæk, Benjamin Lee, Dieter Schmalstieg, Raimund Dachselt, and Harald Reiterer. 2026. Hybrid User Interfaces: Past, Present, and Future of Complementary Cross-Device Interaction in Mixed Reality. *IEEE Transactions on Visualization and Computer Graphics* (2026), 1–20. doi:10.1109/TVCG.2026.3683941

[56] Sebastian Hubenschmid, Jonathan Wieland, Daniel Immanuel Fink, Andrea Batch, Johannes Zagermann, Niklas Elmqvist, and Harald Reiterer. 2022. ReLive: Bridging In-Situ and Ex-Situ Visual Analytics for Analyzing Mixed Reality User Studies. In *CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3491102.3517550

[57] Sebastian Hubenschmid, Johannes Zagermann, Simon Butscher, and Harald Reiterer. 2021. STREAM: Exploring the Combination of Spatially-Aware Tablets with Augmented Reality Head-Mounted Displays for Immersive Analytics. 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, 1–14. doi:10.1145/3411764.3445298

[58] Griffin J. Hurt, Talha Khan, Nicolás Matheo Kass, Anthony Tang, Edward Andrews, and Jacob Biehl. 2025. MR-MDEs: Exploring the Integration of Mixed Reality into Multi-display Environments. *Proceedings of the ACM on Human-Computer Interaction* 9, 8 (2025), 371–392. doi:10.1145/3773074

[59] Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. 1985. Direct Manipulation Interfaces. *Human-Computer Interaction* 1, 4 (1985), 311–338. doi:10.1207/s15327051hci0104_2

[60] Yhonatan Iquiapaza, Jorge Wagner, and Luciana Nedel. 2023. DeAR: Combining Desktop and Augmented Reality for Visual Data Analysis. In *Symposium on Virtual and Augmented Reality* (Rio Grande, Brazil) *(SVR '23)*. Association for Computing Machinery, New York, NY, USA, 233–237. doi:10.1145/3625008.3625021

[61] Nuwan Janaka, Runze Cai, Ashwin Ram, Lin Zhu, Shengdong Zhao, and Kai Qi Yong. 2024. PilotAR: Streamlining Pilot Studies with OHMDs from Concept to Insight. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 3 (2024), 1–35. doi:10.1145/3678576

[62] Pascal Jansen, Julian Britten, Alexander Häusele, Thilo Segschneider, Mark Colley, and Enrico Rukzio. 2023. AutoVis: Enabling Mixed-Immersive Analysis of Automotive User Interface Interaction Studies. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–23. doi:10.1145/3544548.3580760

[63] Haiyan Jiang, Dongdong Weng, Zhenliang Zhang, Yihua Bao, Yufei Jia, and Mengman Nie. 2018. HiKeyb: High-Efficiency Mixed Reality System for Text Entry. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Munich, Germany). IEEE, 132–137. doi:10.1109/ismar-adjunct.2018.00051

[64] Xinhui Jiang, Jussi P. P. Jokinen, Antti Oulasvirta, and Xiangshi Ren. 2022. Learning to Type with Mobile Keyboards: Findings with a Randomized Keyboard. *Computers in Human Behavior* 126 (Jan. 2022), 106992. doi:10.1016/j.chb.2021.106992

[65] Tero Jokela, Jarno Ojala, and Thomas Olsson. 2015. A Diary Study on Combining Multiple Information Devices in Everyday Activities and Tasks. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 3903–3912. doi:10.1145/2702123.2702211

[66] Sungchul Jung, Yuanjie Wu, Stephan Lukosch, Heide Lukosch, Ryan Douglas Mckee, and Robert W. Lindeman. 2023. Cross-Reality Gaming: Comparing Competition and Collaboration in an Asymmetric Gaming Experience. In *29th ACM Symposium on Virtual Reality Software and Technology* (Christchurch, New Zealand) *(VRST 2023)*. Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3611659.3615698

[67] Andrew Kang, Hanspeter Pfister, and Tica Lin. 2024. PanoCoach: Enhancing Tactical Coaching and Communication in Soccer with Mixed-Reality Telepresence. In *Companion Proceedings of the 2024 Conference on Interactive Surfaces and Spaces* (Vancouver, BC, Canada) *(ISS '24)*. Association for Computing Machinery, New York, NY, USA, 13–17. doi:10.1145/3696762.3698043

[68] Taiga Kashima, Daichi Suzuo, and Fabrice Matulic. 2025. Enhancing Touch Input in Desktop VR: A Hybrid Approach Combining Headset-Based Hand Tracking and Touch Detection with a Smartphone. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3706599.3719768

[69] Jake Knapp, John Zeratsky, and Braden Kowitz. 2016. *Sprint: How to Solve Big Problems and Test New Ideas in Just Five Days.* Simon & Schuster.

[70] Panayiotis Koutsabasis and Panagiotis Vogiatzidakis. 2019. Empirical Research in Mid-Air Interaction: A Systematic Review. *International Journal of Human–Computer Interaction* 35, 18 (2019), 1747–1768. doi:10.1080/10447318.2019.1572352

[71] Balint Istvan Kovacs, Ingrid Erb, Hannes Kaufmann, and Peter Ferschin. 2023. MR.Sketch. Immediate 3D Sketching via Mixed Reality Drawing Canvases. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (Sydney, Australia). IEEE, 10–19. doi:10.1109/ismar59233.2023.00015

[72] Jan Krieglstein, Jan Kolberg, Aimée Sousa Calepso, Werner Kraus, and Michael Sedlmair. 2025. A Hybrid User Interface Combining AR, Desktop, and Mobile Interfaces for Enhanced Industrial Robot Programming. In *2025 IEEE International Conference on Robotics and Automation (ICRA)* (Atlanta, GA, USA). IEEE, 4738–4745. doi:10.1109/icra55743.2025.11128291

[73] Katja Krug, Wolfgang Büschel, Konstantin Klamka, and Raimund Dachselt. 2022. CleAR Sight: Exploring the Potential of Interacting with Transparent Tablets in Augmented Reality. In *Proceedings of the 2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 196–205. doi:10.1109/ISMAR55827.2022.00034

[74] Balasaravanan Thoravi Kumaravel, Cuong Nguyen, Stephen DiVerdi, and Bjoern Hartmann. 2020. TransceiVR: Bridging Asymmetrical Communication Between VR Users and External Collaborators. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) *(UIST '20)*. Association for Computing Machinery, New York, NY, USA, 182–195. doi:10.1145/3379337.3415827

[75] Stanislav Kyian and Robert Teather. 2021. Selection Performance Using a Smartphone in VR with Redirected Input. In *Symposium on Spatial User Interaction* (Virtual Event, USA) *(SUI '21)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3485279.3485292

[76] Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachselt. 2021. MARVIS: Combining Mobile Devices and Augmented Reality for Visual Data Analysis. 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, 1–17. doi:10.1145/3411764.3445593

[77] Juyoung Lee, Minju Baeck, Hui-Shyong Yeo, Thad Starner, and Woontack Woo. 2024. GestureMark: Shortcut Input Technique using Smartwatch Touch Gestures for XR Glasses. In *Proceedings of the Augmented Humans International Conference 2024* (Melbourne, VIC, Australia) *(AHs 2024)*. Association for Computing Machinery, New York, NY, USA, 63–71. doi:10.1145/3652920.3652941

[78] Nianlong Li, Teng Han, Feng Tian, Jin Huang, Minghui Sun, Pourang Irani, and Jason Alexander. 2020. Get a Grip: Evaluating Grip Gestures for VR Input using a Lightweight Pen. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376698

[79] Qingchuan Li, Zhao Xu, and Yang Chen. 2025. Understanding college students' cross-device learning behavior in the wild: Device ecologies, physical configurations, usage patterns, and attention issues. *Computers & Education* 229 (2025), 105259. doi:10.1016/j.compedu.2025.105259

[80] Wanwan Li. 2022. Simulating Turing Machine in Augmented Reality. In *2022 International Conference on Computational Science and Computational Intelligence (CSCI)* (Las Vegas, NV, USA). IEEE, 2015–2022. doi:10.1109/csci58124.2022.00362

[81] Hai-Ning Liang, Yuwei Shi, Feiyu Lu, Jizhou Yang, and Konstantinos Papangelis. 2016. VRMController: an input device for navigation activities in virtual reality environments. In *Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry - Volume 1* (Zhuhai, China) *(VRCAI '16)*. Association for Computing Machinery, New York, NY, USA, 455–460. doi:10.1145/3013971.3014005

[82] Shuqi Liao, Sparsh Chaudhri, Maanas K. Karwa, and Voicu Popescu. 2025. SeamlessVR: Bridging the Immersive to Non-Immersive Visualization Divide. *IEEE Transactions on Visualization and Computer Graphics* 31, 5 (2025), 2806–2816. doi:10.1109/tvcg.2025.3549564

[83] YaPeng Liu, Kai Chen, Dongliang Guo, and Fengyi Yang. 2025. DocVision: A Seamless, Cross-Device Immersive Active Reading Framework for Digital Academic Literature. In *2025 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (Daejeon, Korea, Republic of). IEEE, 847–856. doi:10.1109/ismar67309.2025.00093

[84] Rasmus Lunding, Sebastian Hubenschmid, Tiare Feuchtner, and Kaj Grønbæk. 2025. ARTHUR: authoring human–robot collaboration processes with augmented reality using hybrid user interfaces. *Virtual Reality* 29, 2, Article 73

(2025). doi:10.1007/s10055-025-01149-6

[85] Weizhou Luo, Eva Goebel, Patrick Reipschlager, Mats Ole Ellenberg, and Raimund Dachselt. 2021. Exploring and Slicing Volumetric Medical Data in Augmented Reality Using a Spatially-Aware Mobile Device. In *2021 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Bari, Italy). IEEE, 334–339. doi:10.1109/ismar-adjunct54149.2021.00076

[86] Weizhou Luo, Rufat Rzayev, Benjamin Russig, Sivanon Visutarporn, Marc Satkowski, Stefan Gumhold, and Raimund Dachselt. 2026. Beyond Links: Exploring Visual Representations of Multi-View Relations in Mixed Reality. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 637, 19 pages. doi:10.1145/3772318.3791398

[87] Allan MacLean, Richard M. Young, Victoria M. E. Bellotti, and Thomas P. Moran. 1991. Questions, Options, and Criteria: Elements of Design Space Analysis. *Human–Computer Interaction* 6, 3–4 (1991), 201–250. doi:10.1080/07370024.1991.9667168

[88] Christian Mai, Lukas Rambold, and Mohamed Khamis. 2017. TransparentHMD: revealing the HMD user's face to bystanders. In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia* (Stuttgart, Germany) *(MUM 2017)*. Association for Computing Machinery, New York, NY, USA, 515–520. doi:10.1145/3152832.3157813

[89] Akhmajon Makhsadov, Donald Degraen, André Zenner, Felix Kosmalla, Kamila Mushkina, and Antonio Krüger. 2022. VRySmart: a Framework for Embedding Smart Devices in Virtual Reality. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3491101.3519717

[90] Emeric Marguet, Marcel Borowski, Janus Bager Kristensen, Clemens Nylandsted Klokmose, and Niklas Elmqvist. 2025. WindowSpace: A Web-Based XR Window Manager for Interacting with 2D Windows in Immersive 3D Space. *Proceedings of the ACM on Human-Computer Interaction* 9, 8 (2025), 119–144. doi:10.1145/3773063

[91] Fabrice Matulic and Daniel Vogel. 2022. Terrain Modelling with a Pen & Touch Tablet and Mid-Air Gestures in Virtual Reality. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3491101.3519867

[92] Elisabeth Mayer, Francesco Chiossi, and Sven Mayer. 2024. Crossing Mixed Realities: A Review for Transitional Interfaces Design. In *Proceedings of Mensch Und Computer 2024* (Karlsruhe, Germany) *(MuC '24)*. Association for Computing Machinery, New York, NY, USA, 629–634. doi:10.1145/3670653.3677481

[93] Mark McGill, Stephen Brewster, Daniel Pires De Sa Medeiros, Sidney Bovet, Mario Gutierrez, and Aidan Kehoe. 2022. Creating and Augmenting Keyboards for Extended Reality with the Keyboard Augmentation Toolkit. *ACM Transactions on Computer-Human Interaction* 29, 2 (2022), 1–39. doi:10.1145/3490495

[94] Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica* 22, 3 (2012), 276–282.

[95] Alexandre Millette and Michael J. McGuffin. 2016. DualCAD: Integrating Augmented Reality with a Desktop GUI and Smartphone Interaction. In *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)* (Merida, Yucatan, Mexico). IEEE, 21–26. doi:10.1109/ismar-adjunct.2016.0030

[96] Roberto A. Montano-Murillo, Cuong Nguyen, Rubaiat Habib Kazi, Sriram Subramanian, Stephen DiVerdi, and Diego Martinez-Plasencia. 2020. Slicing-Volume: Hybrid 3D/2D Multi-target Selection Technique for Dense Virtual Environments. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (Atlanta, GA, USA). IEEE, 53–62. doi:10.1109/vr46266.2020.00023

[97] Robert C Nickerson, Upkar Varshney, and Jan Muntermann. 2013. A method for taxonomy development and its application in information systems. *European Journal of Information Systems* 22, 3 (2013), 336–359. doi:10.1057/ejis.2012.26

[98] Matthew J Page, Joanne E McKenzie, Patrick M Bossuyt, Isabelle Boutron, Tammy C Hoffmann, Cynthia D Mulrow, Larissa Shamseer, Jennifer M Tetzlaff, Elie A Akl, Sue E Brennan, et al. 2021. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *bmj* 372 (2021).

[99] Jiafei Pan, Dongdong Weng, and Jialong Mo. 2019. Object Manipulation: Interaction for Virtual Reality on Multi-touch Screen. In *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Beijing, China). IEEE, 280–284. doi:10.1109/ismar-adjunct.2019.00-29

[100] Zhigeng Pan, Zhipeng Pan, Tianren Luo, and Mingmin Zhang. 2022. Exploring the Use of Smartphones as Input Devices for the Mixed Reality Environment. In *Proceedings of the 18th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry* (Guangzhou, China) *(VRCAI '22)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3574131.3574451

[101] Leonardo Pavanatto, Verena Biener, Jennifer Chandran, Snehanjali Kalamkar, Feiyu Lu, John J. Dudley, Jinghui Hu, G. Nikki Ramirez, Per Ola Kristensson, Alexander Giovannelli, Luke Schlueter, Jörg Müller, Jens Grubert, and Doug A. Bowman. 2025. Working in Extended Reality in the Wild: Worker and Bystander Experiences of XR Virtual Displays in Public Real-World Settings. *IEEE Transactions on Visualization and Computer Graphics* 31, 10 (2025), 9104–9123. doi:10.1109/tvcg.2025.3589283

[102] Priyanka Pazhayedath, Pedro Belchior, Rafael Prates, Filipe Silveira, Daniel Simoes Lopes, Robbe Cools, Augusto Esteves, and Adalberto L. Simeone. 2021. Exploring Bi-Directional Pinpointing Techniques for Cross-Reality Collaboration. In *2021 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)* (Lisbon, Portugal). IEEE, 264–270. doi:10.1109/vrw52623.2021.00055

[103] Gary Perelman, Emmanuel Dubois, Alice Probst, and Marcos Serrano. 2022. Visual Transitions around Tabletops in Mixed Reality: Study on a Visual Acquisition Task between Vertical Virtual Displays and Horizontal Tabletops. *Proceedings of the ACM on Human-Computer Interaction* 6, ISS (2022), 660–679. doi:10.1145/3567738

[104] Lucas Plabst, Lena Plabst, Florian Niebling, and Francisco R. Ortega. 2025. Five-day research-in-the-wild observation of notifications on smartglasses: A double edged sword. In *Proceedings of the 2025 31st ACM Symposium on Virtual Reality Software and Technology* (Montreal, QC, Canada) *(VRST '25)*. Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3756884.3765976

[105] Carole Plasson, Renaud Blanch, and Laurence Nigay. 2022. Selection Techniques for 3D Extended Desktop Workstation with AR HMD. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (Singapore, Singapore). IEEE, 460–469. doi:10.1109/ismar55827.2022.00062

[106] Thiago Porcino, Seyed Adel Ghaeinian, Derek Reilly, and Joseph Malloch. 2024. AQUARIA: Assistive Querying in Augmented Reality for Interactive Analytics. In *OCEANS 2024 - Halifax* (Halifax, NS, Canada). IEEE, 1–10. doi:10.1109/oceans55160.2024.10753871

[107] Michael I Posner. 1980. Orienting of attention. *Quarterly journal of experimental psychology* 32, 1 (1980), 3–25.

[108] Xiangyu Qi, Dongdong Weng, Jie Hao, and Zihao Li. 2024. Text Input Through Swipe Gestures Based Personalized AI Agent. In *2024 3rd International Conference on Automation, Robotics and Computer Engineering (ICARCE)* (China). IEEE, 435–439. doi:10.1109/icarce63054.2024.00088

[109] Leping Qiu, Erin Seongyoon Kim, Sangho Suh, Ludwig Sidenmark, and Tovi Grossman. 2025. MaRginalia: Enabling In-person Lecture Capturing and Note-taking Through Mixed Reality. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3706598.3714065

[110] Carlos Quijano-Chavez, Nina Doerr, Benjamin Lee, Dieter Schmalstieg, and Michael Sedlmair. 2024. Brushing and Linking for Situated Analytics. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)* (Orlando, FL, USA). IEEE, 597–603. doi:10.1109/vrw62533.2024.00116

[111] Luis Quintero, Panagiotis Papapetrou, and John E. Munoz. 2019. Open-Source Physiological Computing Framework using Heart Rate Variability in Mobile Virtual Reality Applications. In *2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)* (San Diego, CA, USA). IEEE, 126–1267. doi:10.1109/aivr46125.2019.00027

[112] Roman Rädle, Hans-Christian Jetter, Nicolai Marquardt, Harald Reiterer, and Yvonne Rogers. 2014. HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces*. ACM, New York, NY, USA, 45–54. doi:10.1145/2669485.2669500

[113] Adrian Ramcharitar and Robert J. Teather. 2018. EZCursorVR: 2D Selection with Virtual Reality Head-Mounted Displays. In *Proceedings of Graphics Interface 2018 (GI '18)*. Canadian Human-Computer Communications Society, Toronto, Ontario, Canada, 123–130. doi:10.20380/gi2018.17

[114] Umar Rashid, Miguel A. Nacenta, and Aaron Quigley. 2012. Factors influencing visual attention switch in multi-display user interfaces: a survey. In *Proceedings of the 2012 International Symposium on Pervasive Displays* (Porto, Portugal) *(PerDis '12)*. Association for Computing Machinery, New York, NY, USA, Article 1, 6 pages. doi:10.1145/2307798.2307799

[115] Tobias Rau, Tobias Isenberg, Andreas Koehn, Michael Sedlmair, and Benjamin Lee. 2025. Traversing Dual Realities: Investigating Techniques for Transitioning 3D Objects between Desktop and Augmented Reality Environments. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3706598.3713949

[116] Muhammad Raza, Vachiraporn Ketsoi, and Derek Reilly. 2026. A Scoping Survey on Augmented Display Systems. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 706, 26 pages. doi:10.1145/3772318.3790680

[117] Carolin Reichherzer, Jack Fraser, Damien Constantine Rompapas, and Mark Billinghurst. 2021. SecondSight: A Framework for Cross-Device Augmented Reality Interfaces. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3411763.3451839

[118] Patrick Reipschläger, Frederik Brudy, Raimund Dachselt, Justin Matejka, George Fitzmaurice, and Fraser Anderson. 2022. AvatAR: An Immersive Analysis

Environment for Human Motion Data Combining Interactive 3D Avatars and Trajectories. In *CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3491102.3517676

[119] Patrick Reipschläger and Raimund Dachselt. 2019. DesignAR: Immersive 3D-Modeling Combining Augmented Reality with Interactive Displays. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces* (Daejeon, Republic of Korea) *(ISS '19)*. Association for Computing Machinery, New York, NY, USA, 29–41. doi:10.1145/3343055.3359718

[120] Patrick Reipschläger, Burcu Kulahcioglu Ozkan, Aman Shankar Mathur, Stefan Gumhold, Rupak Majumdar, and Raimund Dachselt. 2018. DebugAR: Mixed Dimensional Displays for Immersive Debugging of Distributed Systems. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal, QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3170427.3188679

[121] Jun Rekimoto and Masanori Saitoh. 1999. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 378–385. doi:10.1145/302979.303113

[122] Mauro Del Rio, Vittorio Meloni, Francesca Frexia, Francesco Cabras, Roberto Tumbarello, Sabrina Montis, Andrea Marini, and Gianluigi Zanetti. 2018. Augmented Reality for Supporting Real Time Telementoring: an Exploratory Study Applied to Ultrasonography. In *Proceedings of the 2nd International Conference on Medical and Health Informatics* (Tsukuba, Japan) *(ICMHI '18)*. Association for Computing Machinery, New York, NY, USA, 218–222. doi:10.1145/3239438.3239444

[123] David Saffo, Andrea Batch, Cody Dunne, and Niklas Elmqvist. 2023. Through Their Eyes and In Their Shoes: Providing Group Awareness During Collaboration Across Virtual Reality and Desktop Platforms. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3544548.3581093

[124] Houssem Saidi, Emmanuel Dubois, and Marcos Serrano. 2021. HoloBar: Rapid Command Execution for Head-Worn AR Exploiting Around the Field-of-View Interaction. 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, 1–17. doi:10.1145/3411764.3445255

[125] Daniel Schneider, Alexander Otte, Travis Gesslein, Philipp Gagel, Bastian Kuth, Mohamad Shahm Damlakhi, Oliver Dietz, Eyal Ofek, Michel Pahud, Per Ola Kristensson, Jörg Müller, and Jens Grubert. 2019. ReconViguRation: Reconfiguring Physical Keyboards in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 25, 11 (2019), 3190–3201. doi:10.1109/tvcg.2019.2932239

[126] Dennis Schneider, Julian Seifert, and Enrico Rukzio. 2012. MobIES: extending mobile interfaces using external screens. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia* (Ulm, Germany) *(MUM '12)*. Association for Computing Machinery, New York, NY, USA, Article 59, 2 pages. doi:10.1145/2406367.2406438

[127] Ephraim Schott, Elhassan Belal Makled, Tony Jan Zoeppig, Sebastian Muehlhaus, Florian Weidner, Wolfgang Broll, and Bernd Froehlich. 2023. UniteXR: Joint Exploration of a Real-World Museum and its Digital Twin. In *29th ACM Symposium on Virtual Reality Software and Technology* (Christchurch, New Zealand) *(VRST 2023)*. Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3611659.3615708

[128] Jan-Henrik Schröder, Daniel Schacht, Niklas Peper, Anita Marie Hamurculu, and Hans-Christian Jetter. 2023. Collaborating Across Realities: Analytical Lenses for Understanding Dyadic Collaboration in Transitional Interfaces. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3544548.3580879

[129] Mohammad Rajabi Seraji, Parastoo Piray, Vahid Zahednejad, and Wolfgang Stuerzlinger. 2024. Analyzing User Behaviour Patterns in a Cross-Virtuality Immersive Analytics System. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (2024), 2613–2623. doi:10.1109/tvcg.2024.3372129

[130] Wei Shi, Wanxin Xiong, and Yoshihiro Okada. 2025. A System Supporting Sport and Dance Learning via Real-Time Skeleton Estimation. In *2025 IEEE International Symposium on Consumer Technology (ISCT)* (Denpasar, Bali, Indonesia). IEEE, 327–331. doi:10.1109/isct66099.2025.11297391

[131] Swati Shilaskar, Mrunal Shinde, Sameer Tadavi, and Prasad Shevale. 2024. VR Based Medical Procedure Simulator with Haptic Feedback Gloves. In *2024 International Conference on Emerging Systems and Intelligent Computing (ESIC)*. 533–537. doi:10.1109/ESIC60604.2024.10481573

[132] Richard Skarbez, Missie Smith, and Mary C. Whitton. 2021. Revisiting Milgram and Kishino's Reality-Virtuality Continuum. *Frontiers in Virtual Reality* 2 (2021). doi:10.3389/frvir.2021.647997

[133] Maximilian Speicher, Brian D. Hall, and Michael Nebeling. 2019. What is Mixed Reality?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland, UK) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3290605.3300767

[134] Lorenzo Stacchio, Giuseppe Di Maria, and Gustavo Marfia. 2024. Flying in XR: Bridging Desktop Applications in eXtended Reality through Deep Learning. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)* (Orlando, FL, USA). IEEE, 471–476. doi:10.1109/vrw62533.2024.00093

[135] Junwei Sun, Wolfgang Stuerzlinger, and Bernhard E. Riecke. 2018. Comparing input methods and cursors for 3D positioning with head-mounted displays. In *Proceedings of the 15th ACM Symposium on Applied Perception* (Vancouver, British Columbia, Canada) *(SAP '18)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3225153.3225167

[136] Xiaowen Sun, Xiangxu Meng, Yafang Wang, Gerard de Melo, Wei Gai, Yuliang Shi, Lu Zhao, Yulong Bian, Juan Liu, and Chenglei Yang. 2017. Enabling Participatory Design of 3D Virtual Scenes on Mobile Devices. In *Proceedings of the 26th International Conference on World Wide Web Companion - WWW '17 Companion* (Perth, Australia) *(WWW '17 Companion)*. ACM Press, New York, New York, USA, 473–482. doi:10.1145/3041021.3054173

[137] Zhi Sun, Ruiqing Zhou, Kai Wang, and Zhe Li. 2024. Enhancing Mobile Interaction for Individuals With Tremors via Optical See-Through Augmented Reality. *IEEE Access* 12 (2024), 123946–123955. doi:10.1109/access.2024.3449880

[138] Hemant Bhaskar Surale, Aakar Gupta, Mark Hancock, and Daniel Vogel. 2019. TabletInVR: Exploring the Design Space for Using a Multi-Touch Tablet in Virtual Reality. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland, UK) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300243

[139] Ibrahim A. Tahmid, Francielly Rodrigues, Alexander Giovannelli, Lee Lisle, Jerald Thomas, and Doug A. Bowman. 2023. CoLT: Enhancing Collaborative Literature Review Tasks with Synchronous and Asynchronous Awareness Across the Reality-Virtuality Continuum. In *2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (Sydney, Australia). IEEE, 831–836. doi:10.1109/ismar-adjunct60411.2023.00183

[140] Kazuhiro Takeda and Hiroyuki Manabe. 2024. Screen Augmentation Technique Using AR Glasses and Smartphone without External Sensors. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–5. doi:10.1145/3613905.3648664

[141] 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. doi:10.1145/3411764.3445497

[142] Wai Tong, Haobo Li, Meng Xia, Wong Kam-Kwai, Ting-Chuen Pong, Huamin Qu, and Yalong Yang. 2025. Exploring Spatial Hybrid User Interface for Visual Sensemaking. *IEEE Transactions on Visualization and Computer Graphics* 31, 10 (2025), 7062–7077. doi:10.1109/tvcg.2025.3538771

[143] Esen K. Tütüncü, Mar Gonzalez-Franco, and Eric J Gonzalez. 2025. HandOver: Enabling Precise Selection & Manipulation of 3D Objects with Mouse and Hand Tracking. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology* (Busan, Republic of Korea) *(UIST '25)*. Association for Computing Machinery, New York, NY, USA, 1–11. doi:10.1145/3746059.3747757

[144] Arda Ege Unlu and Robert Xiao. 2021. PAIR: Phone as an Augmented Immersive Reality Controller. In *Proceedings of the 27th ACM Symposium on Virtual Reality Software and Technology* (Osaka, Japan) *(VRST '21)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3489849.3489878

[145] Katja Vock, Sebastian Hubenschmid, Johannes Zagermann, Simon Butscher, and Harald Reiterer. 2021. IDIAR: Augmented Reality Dashboards to Supervise Mobile Intervention Studies. In *Mensch und Computer 2021* (Ingolstadt, Germany) *(MuC '21)*. Association for Computing Machinery, New York, NY, USA, 248–259. doi:10.1145/3473856.3473876

[146] Julius von Willich, Frank Nelles, Wen-Jie Tseng, Jan Gugenheimer, Sebastian Günther, and Max Mühlhäuser. 2025. A Qualitative Investigation of User Transitions and Frictions in Cross-Reality Applications. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713921

[147] Yuto Wada, Myungguen Choi, Kaoru Shirane, and Buntarou Shizuki. 2025. Flick-in: Japanese Text Entry Method for Indirect Touch Using Bezel-Initiated Swipe. In *Proceedings of the 2025 ACM Symposium on Spatial User Interaction* (Montreal, QC, Canada) *(SUI '25)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3694907.3765927

[148] Nanjia Wang, Daniel Zielasko, and Frank Maurer. 2024. User Preferences for Interactive 3D Object Transitions in Cross Reality - An Elicitation Study. In *Proceedings of the 2024 International Conference on Advanced Visual Interfaces* (Arenzano, Genoa Italy) *(AVI 2024)*. Association for Computing Machinery, New York, NY, USA, 1–9. doi:10.1145/3656650.3656698

[149] Quanyu Wang, Zhi Wang, Bei Li, and Dejian Wei. 2021. An Improved YOLOv3 Object Detection Network for Mobile Augmented Reality. In *2021 IEEE 7th International Conference on Virtual Reality (ICVR)*. 332–339. doi:10.1109/ICVR51878.2021.9483829

[150] Johann Wentzel, Fraser Anderson, George Fitzmaurice, Tovi Grossman, and Daniel Vogel. 2024. SwitchSpace: Understanding Context-Aware Peeking Between VR and Desktop Interfaces. 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, 1–16. doi:10.1145/3613904.3642358

[151] Jonathan Wieland, Hyunsung Cho, Sebastian Hubenschmid, Akihiro Kiuchi, Harald Reiterer, and David Lindlbauer. 2024. Push2AR: Enhancing Mobile List Interactions Using Augmented Reality. In *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (Bellevue, WA, USA). IEEE, 671–680. doi:10.1109/ismar62088.2024.00082

[152] Sixuan Wu, Jiannan Li, Maurício Sousa, and Tovi Grossman. 2023. Investigating Guardian Awareness Techniques to Promote Safety in Virtual Reality. In *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)* (Shanghai, China). IEEE, 631–640. doi:10.1109/vr55154.2023.00078

[153] Rui Xavier, José Luís Silva, Rodrigo Ventura, and Joaquim Armando Pires Jorge. 2024. Pseudo-Haptics Survey: Human-Computer Interaction in Extended Reality and Teleoperation. *IEEE Access* 12 (2024), 80442–80467. doi:10.1109/ACCESS.2024.3409449

[154] Seda Yilmaz, Colleen M. Seifert, and Richard Gonzalez. 2010. Cognitive Heuristics in Design: Instructional Strategies to Increase Creativity in Idea Generation. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 24, 3 (2010), 335–355. doi:10.1017/S0890060410000235

[155] Ye Yuan, Nathalie Riche, Nicolai Marquardt, Molly Jane Nicholas, Teddy Seyed, Hugo Romat, Bongshin Lee, Michel Pahud, Jonathan Goldstein, Rojin Vishkaie, Christian Holz, and Ken Hinckley. 2022. Understanding Multi-Device Usage Patterns: Physical Device Configurations and Fragmented Workflows. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 64, 22 pages. doi:10.1145/3491102.3517702

[156] Li Zhang, Weiping He, Huidong Bai, Qianyuan Zou, Shuxia Wang, and Mark Billinghurst. 2023. A hybrid 2D–3D tangible interface combining a smartphone and controller for virtual reality. *Virtual Reality* 27, 2 (2023), 1273–1291. doi:10.1007/s10055-022-00735-2

[157] Lixiang Zhao, Tobias Isenberg, Fuqi Xie, Hai-Ning Liang, and Lingyun Yu. 2025. SpatialTouch: Exploring Spatial Data Visualizations in Cross-Reality. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (2025), 897–907. doi:10.1109/tvcg.2024.3456368

[158] Yuchen Zhao, Jay David Bolter, Kexin Chen, Xuerui Wang, Xintong Jiang, and Zhao Ni. 2025. BioVR: An Exploratory Study of Biofeedback-Driven Adaptive VR for Personalized and Sustained Fitness Intensity. In *Companion Publication of the 2025 ACM Designing Interactive Systems Conference* (Funchal, Portugal) *(DIS '25)*. Association for Computing Machinery, New York, NY, USA, 586–591. doi:10.1145/3715668.3736380

[159] Qian Zhou, George Fitzmaurice, and Fraser Anderson. 2022. In-Depth Mouse: Integrating Desktop Mouse into Virtual Reality. In *CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3491102.3501884

[160] Xiaoyan Zhou, Benjamin Lee, Francisco R. Ortega, Anil Ufuk Batmaz, and Yalong Yang. 2024. Lights, Headset, Tablet, Action: Exploring the Use of Hybrid User Interfaces for Immersive Situated Analytics. *Proceedings of the ACM on Human-Computer Interaction* 8, ISS (2024), 517–539. doi:10.1145/3698147

[161] Fengyuan Zhu and Tovi Grossman. 2020. BISHARE: Exploring Bidirectional Interactions Between Smartphones and Head-Mounted Augmented Reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376233

[162] Fengyuan Zhu, Xun Qian, Daniel Kalmar, Mahdi Tayarani, Eric J. Gonzalez, Mar Gonzalez-Franco, David Kim, and Ruofei Du. 2025. Beyond the Phone: Exploring Phone-XR Integration through Multi-View Transitions for Real-World Applications. In *2025 IEEE Conference Virtual Reality and 3D User Interfaces (VR)* (Saint Malo, France). IEEE, 770–780. doi:10.1109/vr59515.2025.00099

A Search Query

The search combined two sets of terms capturing cross-device systems with a set of terms describing immersive reality technologies. A small number of terms that inherently refer to hybrid or cross-reality interfaces were searched independently and therefore were not combined with the Reality Set. In abstract form, the query was:

```
((Main Set OR Supplementary Set) AND Reality Set) OR Standalone Set.
```

A.1 Main Set.

The primary cross-device search terms were:

```
"Hybrid Virtual Environment" OR "Hybrid Interaction"
OR "Cross-Device" OR "Cross Device" OR "Cross-Surface"
OR "Cross Surface" OR "Multi-Device" OR "Multi Device"
OR "Multi-Display" OR "Multi Display" OR "Distributed
User Interface" OR "Distributed User Interfaces" OR
"Transitional Interface" OR "Transitional Interfaces"
```

Standalone Terms. The following terms inherently describe hybrid or cross-reality interfaces and were therefore searched independently, without being combined with the Reality Set:

```
"Hybrid User Interface" OR "Hybrid User Interfaces"
OR "Complementary Interface" OR "Complementary Interfaces"
OR "Augmented Display" OR "Augmented Displays" OR "Cross-Reality"
OR "Cross Reality"
```

A.2 Supplementary Set.

To capture studies that described cross-device configurations primarily through the devices or interaction arrangements involved, we additionally searched:

```
"phone" OR "phones" OR "smartphone" OR "smartphones"
OR "tablet" OR "tablets" OR "tabletop" OR "tabletops"
OR "laptop" OR "laptops" OR "desktop" OR "desktops"
OR "smartwatch" OR "smartwatches" OR "smart watch" OR
"smart watches" OR "keyboard" OR "keyboards" OR "mouse"
OR "mice" OR "pen" OR "pens" OR "pencil" OR "pencils" OR
"stylus" OR "styluses" OR "asymmetric interaction" OR
"asymmetric interactions" OR "asymmetrical interaction"
OR "asymmetrical interactions" OR "companion device" OR
"companion devices"
```

A.3 Reality Set.

Both the Main Set and Supplementary Set were combined with the following reality-related terms:

```
"Augmented Reality" OR "AR" OR "Extended Reality" OR
"XR" OR "Mixed Reality" OR "MR" OR "Virtual Reality"
OR "VR"
```