How Hand Constraints Influence User Defined Gestures in Mixed Reality

Alexander Barquero[*]
Department of Computer & Information Science & Engineering
University of Florida
Gainesville, Florida, USA
a.barquero@ufl.edu

Niriksha Regmi
Department of CISE
University of Florida
Gainesville, Florida, USA
n.regmi@ufl.edu

Oluwatomisin Obajemu
Department of CISE
University of Florida
Gainesville, Florida, USA
o.obajemu@ufl.edu

Rohith Venkatakrishnan[†]
University of Central Florida
Orlando, Florida, USA
rohith.venkatakrishnan@ucf.edu

Christina Boucher
Department of Computer & Information Science & Engineering
University of Florida
Gainesville, Florida, USA
cboucher@cise.ufl.edu

Lisa Anthony
Department of Computer & Information Science & Engineering
University of Florida
Gainesville, Florida, USA
lisa.anthony@mail.ufl.edu

Jaime Ruiz
Department of Computer & Information Science & Engineering
University of Florida
Gainesville, Florida, USA
jaime.ruiz@ufl.edu

Abstract

How do user-defined gestures for mixed reality change when users' hands are engaged in tasks? To address this question, we conducted a gesture elicitation study to understand user preferences and the characteristics of gestures conceptualized in three scenarios with varying levels of hand constraints, namely: "both hands free", "one hand fixed", and "both hands busy". We analyzed these gestures across multiple dimensions and compared our findings with those from prior research. Our results indicate that when both hands are occupied, users tend to favor head gestures over those involving other body parts, such as the eyes or legs. Additionally, we found that most of the proposed gestures were metaphorical, with many influenced by legacy bias. These insights enhance our understanding of how hand constraints influence gesture choices in mixed reality scenarios.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; **Interaction techniques**; **Gestural input**;

[*]Currently at Universidad de Costa Rica, San José, CR.
[†]Work conducted while at the University of Florida, Gainesville, USA.

GI '25, Okanagan, BC, Canada

ACM ISBN 979-8-4007-1876-2/25/05
https://doi.org/10.1145/3769872.3769888

Keywords

gesture interaction, elicitation study, mixed reality, augmented reality, task guidance systems

ACM Reference Format:
Alexander Barquero, Niriksha Regmi, Oluwatomisin Obajemu, Rohith Venkatakrishnan, Christina Boucher, Lisa Anthony, and Jaime Ruiz. 2025. How Hand Constraints Influence User Defined Gestures in Mixed Reality. In *Graphics Interface 2025 (GI '25), May 26–29, 2025, Okanagan, BC, Canada.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3769872.3769888

1 Introduction

In the evolving landscape of mixed reality (MR) technology, developing intuitive interaction methods is crucial for enhancing user experiences across various applications. MR technology enables environments that integrate physical and digital worlds, allowing users to interact with virtual elements naturally while remaining aware of their surroundings. A key user interaction component in these systems involves performing gestures to communicate within the MR environment [4, 10].

Gesture-based interaction provides an intuitive input method that eliminates the need for traditional hardware interfaces, proving advantageous in diverse settings such as maintenance work [16], vehicles [23, 27], museums [39], electronic devices [40], smarthomes [20], and healthcare [18, 42]. In these environments, external factors like noise or distractions can limit voice-based interaction, emphasizing the value of reliable non-verbal alternatives. The success of gestural interfaces in MR systems depends on their design; gestures must be easy to perform, remember, and interpret [42], while also being responsive and accurate to avoid user frustration [18]. Gestures must also minimize physical and cognitive strain to prevent fatigue [12] and avoid hindering task performance [33].

An added challenge occurs when gestures are restricted by the user's environment, particularly in industrial contexts where hands may be occupied [29]. This limitation is significant in MR task guidance systems that guide users through tasks with real-time instructions, aiming to reduce cognitive load, lower error rates, and decrease task completion times [31, 34]. Effective gestures can be essential in these systems, yet it becomes challenging when users' hands are engaged with tools or materials, limiting feasibility.

Given these constraints, this study addresses the research question: *How do hand constraints impact the design and selection of intuitive gestures in MR interactions?* Furthermore, it explores how users adapt gestures between different hand constraint scenarios, such as using one hand versus both hands or when tools or environmental factors limit dexterity.

To examine how hand constraints influence user-defined gestures in MR settings, we conducted a study with 20 participants using the HoloLens 2. Participants proposed gestures for 13 action referents relevant to general MR use and task-specific commands, like navigating steps or requesting demonstrations. The study included four scenarios: one control scenario with both hands available (*Control*), one with the non-dominant hand occupied (*One-Hand-Fixed*), another with both hands occupied (*Two-Hands-Busy*), and a final scenario for refining gestures (*Refinement*).

Our findings reveal that when hands are occupied, users prefer head gestures over other body parts, demonstrating adaptability to hand constraints. Most proposed gestures were metaphorical, influenced by legacy biases or real-world experiences, indicating a reliance on familiar interactions (e.g., swiping and tapping) or common social conventions (e.g., raising a thumb for yes). Notably, even when participants knew their hands were free for refined gestures, they continued exploring head gestures, acknowledging potential hand occupation in MR environments. Simple, one-step gestures that are easy to perform and remember were favored, with common actions like *Accept/Yes* and *Decline/No* showing high consensus, while abstract referents displayed more variability.

Our findings contribute critical insights for designing adaptable MR interfaces, emphasizing the necessity of supporting diverse, low-effort gestures to ensure robust interaction under realistic hand constraints. Furthermore, our results reaffirm the methodological importance of incorporating constraint-based scenarios and a refinement phase in gesture elicitation studies.

2 Related Work

Gesture elicitation studies are widely used to design gesture-based interfaces [42]. By asking users to propose gestures for specific actions (referents), researchers can identify consensus and generate user-defined gesture sets. While only a small number of gestures typically achieve high agreement rates (AgR) [37], the gestures proposed by users provide valuable insights for design, as they tend to be easier to perform, learn, and remember than those created by designers [26]. Consequently, elicitation methods have been applied for diverse technologies and varied contexts, including MR [2, 6, 9, 42, 46].

Early work in MR space has identified user-defined gesture sets and ways to categorize AR gestures [32]. Studies have also examined how AR gestures are influenced by object size and scale [30], how users manipulate single and multiple objects using both speech and gesture interactions [43, 50], and how different gesture modalities (surface and motion) affect AR gesture outcomes [14].

Despite their benefits, elicitation studies also have limitations. *Legacy bias* [25] can restrict the discovery of new gestures for emerging technologies, while *performance bias* [37], may limit the long-term viability of elicited gestures. These biases have been observed in MR-related studies [27, 30]. Attempts to mitigate legacy bias, such as those by Ortega et al. [28], have yielded mixed results. On the other hand, some researchers argue that legacy bias can enhance gesture memorability and familiarity [14, 50].

Another challenge for elicitation studies is ensuring the practical usability of gestures in real-world contexts. For example, an increasing number of studies have explored gesture elicitation in hands-free or constrained scenarios, where users must interact while holding objects, multitasking, or using wearable devices [17, 19]. Some researchers have focused on body-specific solutions, such as head gestures [49] or microgestures (i.e., subtle finger movements) [5, 7, 22], while others have explored object-based inputs [38]. Our research extends these findings by examining how physical constraints, specifically hand occupancy, shape gesture preferences in MR. Unlike prior work, we do not prime participants toward specific interaction modalities, microgestures, or body parts, allowing us to observe how users naturally adapt their gestures when facing limitations. This holistic approach offers broader insights for designing intuitive MR interfaces.

Category	Referent Action
System Commands	1. Accept / Yes 2. Decline / No 3. Help
System Menu	4. Open Menu 5. Hide Display
Task Navigation	6. Show All Steps 7. Next Step 8. Previous Step 9. Show Current Step
Task Assistance	10. Identify 11. Verify 12. Demonstrate 13. Alternative

Table 1: System actions associated with referents for gesture elicitation, listed in order of appearance.

3 Methods

3.1 Referent and Scenario Selection and System Development

To compile a comprehensive list of referents across various actions in a hand-constrained MR context, we reviewed and selected referents related to MR from previous elicitation studies [16, 21, 28, 32, 40, 43]. Furthermore, to better capture real-world situations in which hands can be occupied, we framed the referents and our constraint scenarios in the context of the user needs for task guidance systems presented by Barquero et al. [3], which emphasize critical factors, such as task coordination and user safety, or limitations, such as having to hold an object in place to cut it, where hands play a significant role. Table 1 shows all referents, contexts, and appearance order.

We implemented the referents using Unity and C#, integrating a Microsoft HoloLens 2 MR head-mounted display (HMD) to provide an authentic MR environment. We designed the system to display the referents as visual prompts within the participant's field of view. For each scenario, we presented the "before" and "after" states on the HMD visor to illustrate task transitions, using a local server to manage and progress the referents remotely.

3.2 Participants

We recruited 24 adult participants from a local university. Data from four participants were excluded due to session interruptions or protocol deviations. As a result, the final analysis included data from 20 participants (10 self-identified as men, 10 as women), ranging in age from 18 to 30 years (mean: 20.95, standard deviation: 3.15). All participants were right-handed. Participants were given a choice of $30 or class credit as compensation.

3.3 Experimental Setup

We carried out the study in a room equipped with a computer, a HoloLens 2 device, a fixed object at waist level to hold as a constraint during one of the scenarios, two handheld objects used as constraints for a different scenario, and a whiteboard displaying subjective rating statements. We used a front-facing camera to record participants' full-body interactions, ensuring comprehensive data capture for analysis.

We first asked participants to sign a consent form and gave them an overview of the study. Next, we told the participants that the system is capable of recognizing any gesture. To prevent priming, we avoided any gestures, touch, or mention of examples while providing instructions. We began each session with the *Control* scenario, where both hands were free. For each referent, we showed the "before" and "after" screens, and then asked participants to propose a gesture, explain their reasoning using a think-aloud protocol, and rate the gesture on a 10-point scale based on *goodness of fit* ("*The gesture I picked is a good match for its intended use*"), *ease of use* ("*The gesture I picked is easy to perform*"), *memorability* ("*The gesture I picked is easy to remember*"), and *frequency of use* ("*Assuming the gesture exists, how often would I use it?*"). This rating system is aligned with established practices in other elicitation studies [36, 45]. After the ratings, we asked participants to repeat the gesture to ensure clear and accurate recordings. We repeated this process for all referents in the *Control* scenario. We administered the referents in the order given in Table 1.

Next, we guided the participants through the *One-Hand-Fixed*, *Two-Hands-Busy*, and *Refinement* scenarios. The design of these four sequential scenarios exposed participants to different constraints, allowing them to adapt their gestures according to their encountered limitations. By maintaining a fixed order of scenarios, we observed how each constraint influenced participants' perceptions of new limitations and their adaptations in gesture selection.

In the *Control* scenario, we allowed participants to propose gestures without constraints. In the *One-Hand-Fixed* scenario, we required participants to hold a fixed object with their non-dominant hand, restricting its movement, while keeping their dominant hand free for gestures; to parallel the real-world scenario of holding an object in place during a task. The fixed object was a standing camera tripod, which we instructed participants to hold by its head, adjusted to their waist level. This required participants to grasp the tripod using the *Power Disk* grasp style from the GRASP taxonomy [15]. In the *Two-Hands-Busy* scenario, participants occupied both hands with handheld objects while still allowing free arm movement; to parallel the real-world scenario of holding two objects during a task. The handheld objects were two small, soft, and lightweight American football-shaped stress balls, requiring grasping consistent with the *Medium Wrap* style from GRASP. These limitations allowed us to observe how participants adapted gestures in response to each condition. Importantly, we did not explicitly instruct participants to account for constraints when proposing gestures. This deliberate decision further prevented priming and ensured their adaptations emerged naturally.

In the final *Refinement* scenario, we asked participants to design gestures for the same referents, without constraints. Once again, we refrained from instructing them to reflect on constraint scenarios to avoid priming effects. However, we allowed participants to refine their previous gestures or introduce new ones if they wished. We anticipated the cumulative experience of different constraints would prompt a conceptual change in their gesture choices. Therefore, the gestures that persisted in the *Refinement* scenario should reflect the participants' ability to adapt to various constraints, synthesizing their understanding of feasible actions in these environments.

At the end of the study, we asked participants to complete a demographic questionnaire, compensated them, and debriefed them. All data was securely stored on institutional platforms. The protocol was approved by our Institutional Review Board.

3.4 Data Analysis

From the 20 participants, across 4 scenarios and 13 gestures per scenario, we collected a total of 1,040 gestures. We analyzed these gestures through a multi-phase process using video recordings from the study sessions. First, we documented each gesture and participants' reasoning, following a detailed description protocol developed and refined through multiple joint sessions. In the second phase, two researchers independently coded all documented gestures. We reviewed and resolved coding discrepancies through group discussions and consensus during team working sessions. We then classified the gestures based on five dimensions: *Body Part*,

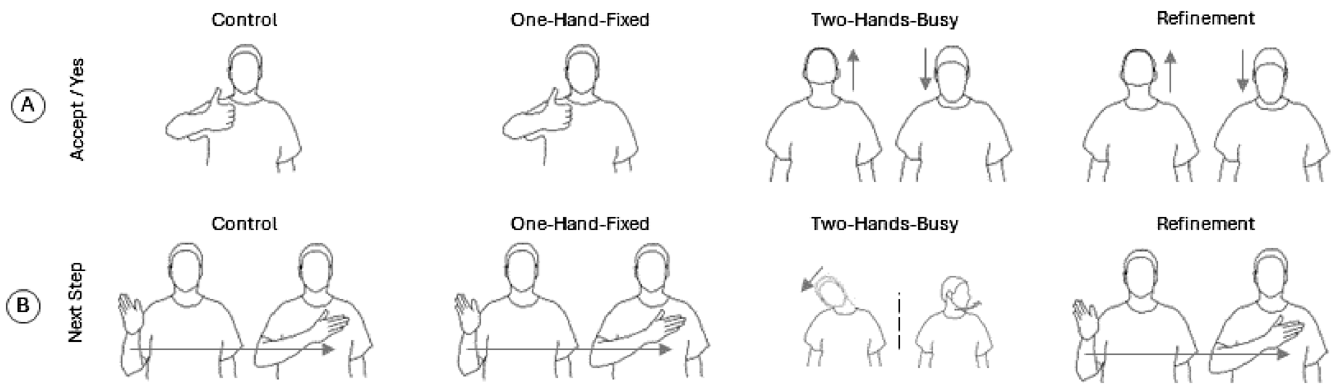

Figure 1: Most common gestures for select referents across all scenarios. (A): For *Accept/Yes,* a shifting pattern is observed from hand gestures to head gestures. (B): For *Next Step,* a returning pattern from head gestures to hand gestures emerges. Hand gestures executed by participants in the *Two-Hands-Busy* scenario were performed while holding objects in their hands.

Temporal, *Context*, *Complexity*, and *Nature*, as described in the *Results* section. These dimensions and their possible encodings were adapted from previous works on gesture works [11, 13, 36, 45, 47]. Two researchers independently assigned classifications to all gestures, and differences were resolved through group discussions until a consensus was reached. By the end of this process, we successfully classified all gestures.

Following the approach of Wobbrock et al. [44] and Ruiz et al. [36], we calculated agreement rates (AgR) between user-defined gestures within each scenario using the AGreement Analysis Toolkit (AGATe) [41]. To account for natural variations in gesture execution, we followed the approach of Piumsomboon et al. [32] on user-defined gestures for AR, relaxing grouping constraints, and prioritizing the most essential characteristics of each gesture. We also cross-verified participants' intent using their reasoning to avoid oversimplification in our groupings. Regarding subjective ratings for gestures, we conducted quantitative analyses to compare scores.

4 Results

4.1 Agreement Rates (AgR)

Participants proposed a total of 194 different gestures across all referents and scenarios. When examining the occurrence of unique gestures in each scenario, we observed a slightly higher number of unique gestures in the *One-Hand-Fixed* (96) and *Two-Hands-Busy* (101) scenarios compared to the *Control* (90) scenario. Interestingly, the number of unique gestures did not decrease greatly in the *Refinement* (99) scenario. However, when examining the uniqueness of gestures across referents, considerably fewer unique gestures were proposed for the *Accept/Yes* (15), *Decline/No* (13), *Next Step* (13), and *Previous Step* (12) referents compared to the total number of unique gestures proposed for the following *Task Assistance* referents: *Demonstrate* (51), and *Alternative* (42), and *Verify* (39).

When looking at gesture repetitions through the AgR lens, our results indicate that some AgR vary when constraints are introduced. For instance, we observed reduced AgR for referents such as *Help* (*Control* = .153 vs. *One-Hand-Fixed* = .026) and *Hide Display* (*Control* = .184 vs, *Two-Hands-Busy* = .037). Conversely, other AgR remained consistently low across all scenarios, including *Verify* (mean = .025, SD = .008), *Demonstrate* (mean = .011, SD = .010), and *Alternative* (mean = .017, SD = .005). In addition, we encountered instances where multiple gestures tied with the highest AgR for certain referents, such as *Open Menu* in the *Control* (2 gesture groups = .068) and *One-Hand-Fixed* (3 gesture groups = .095) scenarios, introducing conflicts in determining a single consensus gesture. A complete list of the most frequent gesture groups for every referent and scenario has been included in Appendix A.

Certain referents evoked highest AgR. In particular, the user-defined gestures for the *Accept/Yes*, *Decline/No*, *Next Step*, and *Previous Step* referents consistently scored the highest AgR across all four scenarios, but with interesting variations. For instance, Figure 1 Section A illustrates the most common gestures for the *Accept/Yes* referent across scenarios. In the *Control* and *One-Hand-Fixed* scenario, "Thumb Up" was the most frequent gesture for *Accept/Yes*. However, in the *Two-Hands-Busy* and *Refinement* scenario, the most frequent gesture changed to "Head Nod". "Head Nod" also demonstrated a higher AgR in *Refinement* and *Two-Hands-Busy* scenarios compared to "Thumb Up" in *Control*.

In contrast, the *Next Step* referent exhibited a distinct pattern, as shown in Figure 1 Section B. While "Swipe Left" was the most frequent gesture in both *Control* and *One-Hand-Fixed*, two head gestures, "Head Swipe Left" and "Head Tilt Right", emerged during *Two-Hands-Busy* tying for highest AgR. However, the original hand gesture from *Control* reasserted itself in *Refinement*.

Overall, AgR and gesture variety across scenarios highlight notable differences in user-defined gestures based on the referent and type of constraint.

4.2 Gesture Characterization

To further analyze gestures, we categorized them into five dimensions: *Body Part*, *Temporal*, *Context*, *Complexity*, and *Nature*. We used two distinct grouping methods to analyze our results. The first method was scenario-based, examining the attributes of gestures produced by participants within each scenario [37]. The second grouping method focused on selecting the most common gestures for each referent and scenario [32]. A complete list of dimensions and codes can be observed in Table 2.

Category	Gesture Description
Body Part	
Arms	Involving arms, hands, and fingers
Head	Using head movements
Eyes	Based on eye movement
Legs	Involving leg movements
Temporal	
Discrete	Action occurs after gesture completion
Continuous	Action occurs during gesture
Context	
Context-dependent	Requires a specific context
Context-independent	Does not require context
Complexity	
Simple	A single-action gesture
Compound	Composed of multiple actions
Nature	
Symbolic	Based on pre-defined symbols
Metaphorical	Represents a real-world metaphor
Deictic	Pointing or interacting with objects
Abstract	No specific or predefined meaning

Table 2: Taxonomy of gestures categorized by body part, temporal characteristics, context, complexity, and nature.

4.2.1 Body Part. This dimension classifies gestures based on the specific body parts involved in their execution. The categories emerged organically during the classification process. We decided to group all gestures performed with fingers, hands, shoulders, and arms under a single category: arm gestures, similar to what Fleiner et al. [16] did for hand gestures. In addition to arms (whether left, right, or both), we also identified the use of the head, eyes, and legs (left, right, and both), and gestures that combined multiple body parts. Notably, none of the participant's proposed gestures could be categorized as object-centric gestures.

When using the second grouping method (most common gestures), we identified 69 top-performing gesture groups across 52 possible combinations (4 scenarios × 13 referents). This number exceeded the total combinations because some referents had multiple gestures that tied for the highest AgR.

As shown in Figure 2 Section B, when counting gestures using the first method (i.e., their appearance across different scenarios), we observed a preference for gestures involving arms in the *Control*, *One-Hand-Fixed*, and *Refinement* scenarios. However, there was a substantial shift in the *Two-Hands-Busy* scenario, with head gestures surpassing all others. This is noteworthy, given that participants still had considerable freedom with their fingers, hands, wrists, and arms to propose gestures while their hands were holding the objects. In fact, 15 out of 20 participants still used their arms to propose multiple gestures in the *Two-Hands-Busy* scenario, yet head gestures remained predominant. In contrast, when analyzing body part usage with the second method, the majority were arm gestures. Of the 69 top-performing gesture groups, only 18 involved the head: 1 in the *One-Hand-Fixed* scenario, 15 in the *Two-Hands-Busy* scenario, and 2 in the *Refinement* scenario, specifically the *Head Nod* for *Yes* and *Head Shake* for *No*.

4.2.2 Nature. This dimension categorized gestures based on the intrinsic characteristics that inspired their creation. Specifically, the following codes were applied:

- *Symbolic*: gestures that use agreed symbols influenced by cultural or social conventions. For example, drawing an "X" in mid-air to indicate rejection or nodding to signal affirmation.
- *Metaphorical*: gestures acting as metaphors for real-world situations or objects, like a gesture mimicking opening a book. This also covers legacy bias gestures, such as swiping to navigate, reminiscent of touchscreen behaviors.
- *Deictic*: gestures involving pointing, gazing at, or touching objects based on their physical or virtual location.
- *Abstract*: gestures that do not fit other categories, often used when participants run out of ideas or cannot justify a gesture.

Several gestures were classified under multiple nature categories. For example, a gesture where a participant points to an object and then air-taps in its direction was categorized as both *deictic* (pointing to the object's location) and *metaphorical* (point-and-click metaphor).

When analyzing gestures based on their nature, we found that a majority (61.25%) were either purely metaphorical or had a metaphorical component. Distribution of gesture types was relatively consistent across scenarios, except the *Two-Hands-Busy* scenario, where abstract gestures became more common and metaphorical gestures decreased. For the gestures with the highest AgR, *metaphorical* and *symbolic* gestures predominated, sometimes including a *deictic* component. Interestingly, no abstract gestures were present in the set of gestures with the highest AgR.

4.2.3 Complexity. This dimension pertains to whether a proposed gesture is classified as a simple or compound gesture. A compound gesture can be broken down into simpler gestures by segmenting pauses in motion, changes in body parts, repetitions, or sequences of entirely different gestures. For example, a single hand swipe is a simple gesture, while a hand swipe followed by an air tap is compound. Approximately 86% of all gestures were classified as simple. Compound gestures were more commonly observed in *Task Assistance* actions. For instance, 27 out of 80 gestures proposed for *Verify* were compound. However, when observing complexity on the group with the highest AgR, we found that all were simple gestures, without exception.

4.2.4 Context. This dimension classifies gestures based on whether they depend on a specific context for their use. For example, a *Tap On* gesture aimed at a message on the screen would be context-dependent, while a head swipe upward to open the menu would be context-independent. During analysis, we found that 77% of all gestures were classified as context-independent. The *Tap On* gesture was the only one that predominantly required a specific context. In contrast, gestures such as *Raise Hand* and *Thumb Down*, were always context-independent. No other differences were observed when analyzed across scenarios or within the highest AgR group.

4.2.5 *Temporal.*
The temporal dimension distinguishes between discrete gestures, where the action occurs after the gesture is fully completed, and continuous, where the system reacts to the gesture as it is being performed. We found that approximately 98% of gestures were discrete, meaning the intended action occurred after the gesture was finished. The remaining 2% were mostly swipe-like movements in which participants stated that objects in screen would drag along their gesturing trajectory. No additional details were observed.

4.2.6 *Summary.*
Figure 2 Section A illustrates the characterization of the 1,040 gestures performed by participants in our study. As previously indicated, totals for the Body Part and Nature dimensions exceed the total number of gestures due to gestures being classified under multiple categories.

4.3 Effects of Hand Constraints on Perceptions of Goodness of Fit, Ease of Use, Memorability, and Frequency of Use

Parametric analyses were conducted on the participants' subjective ratings of the perceived levels of *goodness of fit*, *ease of use*, *memorability*, and *frequency of use* associated with the gestures that users conceptualized over the four different constraint scenarios, after verifying that the assumptions for each test were appropriately met. This included ensuring normal distribution of the data across the levels of the categorical predictor (i.e., each scenario) by using Shapiro Wilk tests of normality, homogeneity of variance using Levene's tests, and validating that the error variances in groups of samples were equivalent using Mauchly's test of sphericity. When the assumptions of sphericity were violated, the effects of the categorical predictor were interpreted after applying Greenhouse Geisser corrections. Post-hoc pairwise comparisons between levels of the within-subjects variable (scenario) were conducted using Bonferroni corrections.

The dependent measures of *goodness of fit*, *ease of use*, *memorability*, and *frequency of use* of the gestures were each separately analyzed using a repeated measures ANOVA, with the scenario (*Control* vs *One-Hand-Fixed* vs *Two-Hands-Busy* vs *Refinement*) as a categorical predictor. Significant effects in each analysis are presented with effect size measures. Table 3 shows the results.

A significant effect of scenario was found for perceived *goodness of fit* ($p < 0.01$). Post hoc pairwise comparisons revealed that gestures in the *Refinement* scenario were perceived as having significantly better *goodness of fit* compared to both *Control* and *Two-Hands-Busy*. Additionally, *goodness of fit* was significantly higher in the *One-Hand-Fixed* scenario compared to *Two-Hands-Busy*. Similarly, scenario had a significant effect on perceived *ease of use* ($p < .001$). Post hoc pairwise comparisons revealed that perceived *ease of use* was significantly higher in the *Refinement* scenario compared to *Two-Hands-Busy*. Additionally, *ease of use* was rated significantly higher in the *Control* scenario than in *Two-Hands-Busy*. Perceived *memorability* also showed a significant effect of scenario ($p < .001$). Post hoc pairwise comparisons revealed that gestures in the *Refinement* scenario were perceived as significantly more memorable than those in *Two-Hands-Busy*. The *memorability* of gestures was also significantly higher in the *Control* scenario compared to *Two-Hands-Busy*. Finally, a significant effect of scenario was also found for perceived *frequency of use* ($p = .008$). Post hoc comparisons revealed gestures in *Refinement* were perceived as significantly more likely to be used than those in *Two-Hands-Busy*.

5 Discussion

Our study revealed how hand constraints shape user-defined gestures and preferences in MR. This section contextualizes these findings within prior research and broader MR applications.

5.1 Metaphorical Gestures and Legacy Bias in Constrained MR Environments

Metaphorical gestures, which mimic real-world interactions, dominated user-defined gestures, reinforcing their role as intuitive inputs in MR [44]. However, in the *Two-Hands-Busy* scenario, users resorted to more abstract gestures, suggesting that physical constraints can push users beyond familiar metaphors [32]. Despite this, metaphorical gestures remained prevalent, underscoring their importance for simplifying MR interactions. Legacy bias adaptations also varied by scenario. For example, in *Control*, traditional screen-like interactions, like pointing to select and dragging to move, were common. However, in *Two-Hands-Busy*, participants adapted these gestures to rely on head movements instead of hands. This aligns with works showing physical constraints driving creative adaptations [1, 37, 48], and confirms legacy bias persistence even when trying to omit it [28]. Furthermore, gestures varied greatly between scenarios, but there were notable exceptions. For example, gestures anchored in a clear spatial schema were very resilient, mainly concerning their *Nature*. In particular, *Next* and *Previous Step* prompted a side-to-side movement metaphor that would be performed with a free hand, with an object in hand, or occasionally with a head movement. In contrast, commands lacking a spatial axis (e.g.: *Verify*, *Demonstrate*) were quickly reinvented, oftentimes with new compound or sometimes abstract motions. Altogether, these findings suggest that constraints amplify the role of motor comfort and legacy bias. Gestures that already enjoy a strong metaphorical anchor merely adjust to constraints or migrate to another body part, while abstract commands evolve until users find an alternate form.

Another key finding was the increased reliance on head gestures when hands were constrained. Gestures like the *Head Nod* for *Accept/Yes* and *Head Shake* for *Decline/No* became dominant in *Two-Hands-Busy*, reflecting users' preference for low-effort alternatives. This supports previous research on minimizing physical strain in interaction design [24]. Even in the *Refinement* phase, users sometimes continued favoring head gestures, highlighting their practicality in MR systems where hand availability is inconsistent.

Interestingly, none of our participants spontaneously proposed any object-centric gestures [38] or microgestures [7], even though such approaches have been highlighted in prior work for hands-free or constrained settings. In our study, all finger-based gestures also involved larger arm movements (e.g., pointing, swiping), and participants never used small-scale finger motions exclusively. This outcome is notable because we did not prime participants toward or away from microgestures or object-centric interactions. Thus, when

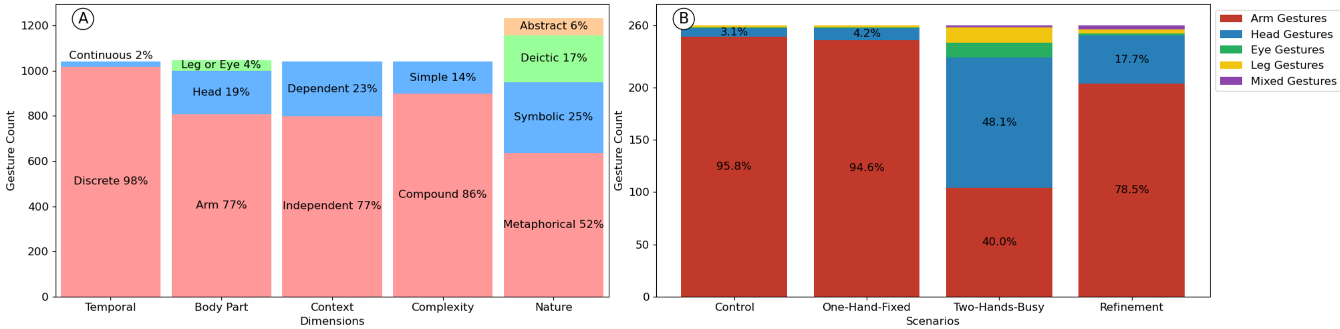

Figure 2: Bar charts for gesture counts across the whole study. (A): A stacked bar chart categorizing all 1040 gestures across dimensions. Gestures can be under multiple codes in the *Body Part* and *Nature* dimensions. (B): A stacked bar chart detailing only the *Body Part* distribution across all four scenarios.

	Control	One-Hand-Fixed	Two-Hands-Busy	Refinement
Goodness of Fit	F(3,57) = 14.375, $p < .001$, $\eta_p^2 = 0.431$			
	7.892 ± 0.230	8.285 ± 0.238	7.600 ± 0.245	8.492 ± 0.222
Ease of Use	F(2.35,44.70) = 9.67, $p < .001$, $\eta_p^2 = 0.337$			
	8.962 ± 0.175	8.835 ± 0.194	8.473 ± 0.204	9.088 ± 0.183
Memorability	F(3,57) = 10.87, $p < .001$, $\eta_p^2 = 0.364$			
	8.365 ± 0.213	8.273 ± 0.234	7.908 ± 0.244	8.581 ± 0.215
Frequency of Use	F(1.877,35.66) = 5.648, $p = .008$, $\eta_p^2 = 0.229$			
	8.342 ± 0.257	8.335 ± 0.302	7.973 ± 0.301	8.504 ± 0.280

Table 3: Summary of F-statistics, p-values, η_p^2, and means ± standard errors for the four scenarios.

not explicitly prompted, users defaulted to more prominent head, arm, or even leg and eye-based solutions. These findings underscore that, in real-world usage, designers may need to introduce or train users for microgestures or object-centric gestures to take advantage of their potential benefits.

5.2 Agreement Rates (AgR) and Gesture Variability

Certain symbolic gestures, such as *Head Shake* for *Decline/No* and *Raise Hand* for *Help*, exhibited high specificity, meaning they were consistently mapped to a single referent. In contrast, other common symbolic gestures like *Wave* were proposed for multiple referents, likely due to the absence of an explicitly related action, such as a greeting. This suggests that designers of MR systems should leverage symbolic gestures only when their meaning clearly aligns with actual system actions. Analysis of agreement rates (AgR) across scenarios revealed strong consensus for familiar actions like *Accept/Yes*, *Decline/No*, *Next Step*, and *Previous Step*, suggesting that these actions are especially well-suited for consensus-driven gesture sets. Notably, the "Head Nod" and "Head Shake" gestures for *Accept/Yes* and *Decline/No* in both the *Two-Hands-Busy* and *Refinement* scenarios yielded higher AgR than those reported for similar referents in previous works [8, 21, 32]. This may be because when our hand-constraint scenarios prompted participants to change their *Control* gestures, they might have opted to rely on these simple and intuitive alternatives. The emergence of these low-effort options aligns with Ruiz and Vogel's [37] findings, showing that introducing soft constraints such as wrist weights encourages users to create a more diverse yet sustainable gesture set. From a methodological perspective, their work, together with our results, underscores the importance of multi-phase elicitation protocols that include realistic constraints and a refinement pass to identify gestures that remain robust when users are encumbered.

Conversely, abstract referents like *Verify* and *Demonstrate* had lower AgR, reinforcing challenges in designing intuitive gestures for ambiguous actions [32]. Interestingly, the *Refinement* scenario did not consistently alter overall AgR, though some specific referents and gestures shifted (Figure 1). This suggests that users are adaptable yet maintain strong preferences for certain intuitive gestures across conditions.

Regarding participants' ratings, we observed how hand constraints negatively impacted subjective ratings of *goodness of fit, ease of performance, memorability*, and *frequency of use*. Gestures in *Two-Hands-Busy* were rated lower across all metrics compared to *Control* and *Refinement*, suggesting that constraints impose additional strain. However, the *Refinement* phase showed a notable increase in ratings, indicating that allowing users to refine gestures enhances confidence and usability. These findings align with prior work showing the benefits of iterative gesture refinement [24, 37].

6 Limitations, Future Work, and Conclusion

This study highlights how hand constraints influence user-defined MR gestures, with participants favoring head gestures when hands were occupied, emphasizing the need for alternative inputs. Legacy bias and familiar metaphors shaped gesture selection, while simple, low-effort gestures were consistently preferred. High-specificity symbolic gestures, such as head nods for confirmation, reinforce the importance of socially ingrained inputs. These findings stress the need for flexible MR systems that support multiple, context-aware gestures for hands-busy environments.

Several limitations should be acknowledged. Participants were primarily university-affiliated, limiting generalizability to diverse populations with varying MR familiarity, social and cultural backgrounds, and professional expertise. Future research should incorporate users from industries such as manufacturing to capture broader perspectives. Additionally, the controlled study environment with the HoloLens 2 may not fully reflect real-world complexities, where noise and social context can impact gesture performance [35]. Testing gestures in dynamic, real-world settings could provide deeper insights into environmental influences on MR interactions.

This study also explored only two constrained-hand scenarios. In contrast, real-world constraints vary, such as using medical tools, or wearing protective gear, and the many different ways of grasping objects which could change the experience. Different contexts may affect gesture adoption, emphasizing the need for adaptable, context-aware alternatives. While referents were chosen for task-guidance relevance, additional actions warrant further exploration.

Not all user-defined gestures are feasible or reliably recognized outside controlled environments. Moreover, high AgR may not necessarily mean gestures are intuitive or self-revealing, nor do they ensure seamless integration with existing MR interactions (e.g., gaze-based pointing, air taps). Future work should, therefore, explore strategies to enhance gesture recognition reliability and mitigate usability challenges in real-world conditions. For instance, exploring multimodal interactions, like integrating constrained gestures with voice, eye-tracking, or virtual-touch controls, could better help understand limitations and enhance adaptability in constrained scenarios.

Further research should also investigate MR gestures in more complex environments with dynamic constraints and varied tasks. For example, longitudinal studies could reveal how user proficiency, fatigue, and constrained gesture preference evolve over time, refining MR system design. Additionally, AI-driven MR systems that learn from user behavior could further optimize interactions, predicting intent and possibly reducing cognitive and physical strain. By addressing these areas, MR systems can better accommodate real-world constraints, ensuring low-effort and adaptable gesture-based interactions that enhance usability across diverse applications.

Acknowledgments

The authors thank the reviewers of current and prior versions of this document whose comments helped us to significantly improve the paper. This work was partially funded by the U.S. Defense Advanced Research Projects Agency under contract #HR00112220004. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect these agencies' views.

References

[1] Bashar Altakrouri, Daniel Burmeister, Dennis Boldt, and Andreas Schrader. 2016. Insights on the impact of physical impairments in full-body motion gesture elicitation studies. In *Proceedings of the Nordic Conference on Human-Computer Interaction*. ACM, Gothenburg, Sweden, 1–10.

[2] Ronald T Azuma. 1997. A Survey of Augmented Reality. *Presence: Teleoperators and Virtual Environments* 6, 4 (1997), 355–385.

[3] Alexander Barquero, Rodrigo Luis Calvo, Daniel Alexander Delgado, Isaac Wang, Lisa Anthony, and Jaime Ruiz. 2024. Understanding User Needs for Task Guidance Systems Through the Lens of Cooking. In *Proceedings of the ACM Designing Interactive Systems Conference*. ACM, Copenhagen, Denmark, 2006–2018.

[4] Mark Billinghurst, Adrian Clark, and Gun Lee. 2015. A Survey of Augmented Reality. *Computers & Graphics* 39 (2015), 1–12. https://doi.org/10.1016/j.cag.2014.11.028

[5] Adrien Chaffangeon Caillet, Alix Goguey, and Laurence Nigay. 2025. Microgesture+ Grasp: A journey from human capabilities to interaction with microgestures. *International Journal of Human-Computer Studies* 195 (2025), 103398.

[6] Nabil Al Nahin Ch, Diana Tosca, Tyanna Crump, Alberta Ansah, Andrew Kun, and Orit Shaer. 2022. Gesture and Voice Commands to Interact With AR Windshield Display in Automated Vehicle: A Remote Elicitation Study. In *Proceedings of the International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, Seoul, Korea, 171–182. https://doi.org/10.1145/3543174.3545257

[7] Edwin Chan, Teddy Seyed, Wolfgang Stuerzlinger, Xing-Dong Yang, and Frank Maurer. 2016. User Elicitation on Single-hand Microgestures. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, San Jose, USA, 3403–3414. https://doi.org/10.1145/2858036.2858589

[8] Zhen Chen, Xiaochi Ma, Zeya Peng, Ying Zhou, Mengge Yao, Zheng Ma, Ci Wang, Zaifeng Gao, and Mowei Shen. 2018. User-defined gestures for gestural interaction: extending from hands to other body parts. *International Journal of Human–Computer Interaction* 34, 3 (2018), 238–250.

[9] Kadek Ananta Satriadi Christopher R. Austin, Barrett Ens and Bernhard Jenny. 2020. Elicitation study investigating hand and foot gesture interaction for immersive maps in augmented reality. *Cartography and Geographic Information Science* 47, 3 (2020), 214–228. https://doi.org/10.1080/15230406.2019.1696232 arXiv:https://doi.org/10.1080/15230406.2019.1696232

[10] Pietro Cipresso, Irene A.C. Giglioli, M. Angeles Raya, and Giuseppe Riva. 2018. The Past, Present, and Future of Virtual and Augmented Reality Research: A Network and Cluster Analysis of the Literature. *Frontiers in Psychology* 9 (2018), 2086. https://doi.org/10.3389/fpsyg.2018.02086

[11] Hélène Cochet and Jacques Vauclair. 2013. Deictic gestures and symbolic gestures produced by adults in an experimental context: Hand shapes and hand preferences. *Laterality: Asymmetries of Body, Brain and Cognition* 19, 3 (June 2013), 278–301. https://doi.org/10.1080/1357650x.2013.804079

[12] Andreea Danielescu and David Piorkowski. 2022. Iterative design of gestures during elicitation: Understanding the role of increased production. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans, USA, 1–14.

[13] Nem Khan Dim, Chaklam Silpasuwanchai, Sayan Sarcar, and Xiangshi Ren. 2016. Designing Mid-Air TV Gestures for Blind People Using User- and Choice-Based Elicitation Approaches. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, Brisbane, Australia, 204–214. https://doi.org/10.1145/2901790.2901834

[14] Ze Dong, Thammathip Piumsomboon, Jingjing Zhang, Adrian Clark, Huidong Bai, and Rob Lindeman. 2020. A Comparison of Surface and Motion User-Defined Gestures for Mobile Augmented Reality. In *Extended Abstracts of the Conference on Human Factors in Computing Systems*. ACM, Honolulu, USA, 1–8. https://doi.org/10.1145/3334480.3382883

[15] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M Dollar, and Danica Kragic. 2015. The grasp taxonomy of human grasp types. *IEEE Transactions on human-machine systems* 46, 1 (2015), 66–77.

[16] Christian Fleiner, Till Riedel, Michael Beigl, and Marcel Ruoff. 2021. Ensuring a Robust Multimodal Conversational User Interface During Maintenance Work. In *MuC '21: Proceedings of Mensch und Computer 2021*. ACM, Ingolstadt, Germany, 79–91. https://doi.org/10.1145/3473856.3473871

[17] Weilun Gong, Stephanie Santosa, Tovi Grossman, Michael Glueck, Daniel Clarke, and Frances Lai. 2023. Affordance-Based and User-Defined Gestures for Spatial Tangible Interaction. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 1500–1514.

[18] Glebys Gonzalez, Naveen Madapana, Rahul Taneja, Lingsong Zhang, Richard Rodgers, and Juan P Wachs. 2018. Looking beyond the gesture: Vocabulary acceptability criteria for gesture elicitation studies. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 62. SAGE Publications, Philadelphia, USA, 997–1001.

[19] Seongkook Heo, Michelle Annett, Benjamin J Lafreniere, Tovi Grossman, and George W Fitzmaurice. 2017. No Need to Stop What You're Doing: Exploring No-Handed Smartwatch Interaction.. In *Graphics Interface*. Canadian Human-Computer Communications Society, Waterloo, CAN, 107–114.

[20] Fabian Hoffmann, Miriam-Ida Tyroller, Felix Wende, and Niels Henze. 2019. User-defined interaction for smart homes: voice, touch, or mid-air gestures?. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia*.

ACM, Pisa, Italy, 1–7.

[21] Masoumehsadat Hosseini, Tjado Ihmels, Ziqian Chen, Marion Koelle, Heiko Müller, and Susanne Boll. 2023. Towards a consensus gesture set: A survey of mid-air gestures in HCI for maximized agreement across domains. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg, Germany, 1–24.

[22] Nikhita Joshi, Parastoo Abtahi, Raj Sodhi, Nitzan Bartov, Jackson Rushing, Christopher Collins, Daniel Vogel, and Michael Glueck. 2023. Transferable Microgestures Across Hand Posture and Location Constraints: Leveraging the Middle, Ring, and Pinky Fingers. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) *(UIST '23)*. Association for Computing Machinery, New York, NY, USA, Article 103, 17 pages. https://doi.org/10.1145/3586183.3606713

[23] Keenan R May, Thomas M Gable, and Bruce N Walker. 2017. Designing an in-vehicle air gesture set using elicitation methods. In *Proceedings of the International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, Oldenburg, Germany, 74–83.

[24] Meredith Ringel Morris, AJ Bernheim Brush, and Brian R Meyers. 2008. A field study of knowledge workers' use of interactive horizontal displays. In *Proceedings of the International Workshop on Horizontal Interactive Human Computer Systems*. IEEE, Amsterdam, Netherlands, 105–112.

[25] Meredith Ringel Morris, Andreea Danielescu, Steven Drucker, Danyel Fisher, Bongshin Lee, M. C. Schraefel, and Jacob O. Wobbrock. 2014. Reducing legacy bias in gesture elicitation studies. *Interactions* 21, 3 (May 2014), 40–45. https://doi.org/10.1145/2591689

[26] Miguel A. Nacenta, Yemliha Kamber, Yizhou Qiang, and Per Ola Kristensson. 2013. Memorability of pre-designed and user-defined gesture sets. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Paris, France, 1099–1108. https://doi.org/10.1145/2470654.2466142

[27] Nabil Al Nahin, Diana Tosca, Tyanna Crump, Alberta Ansah, Andrew Kun, and Orit Shaer. 2022. Gesture and Voice Commands to Interact With AR Windshield Display in Automated Vehicle: A Remote Elicitation Study. In *Proceedings of the International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, Seoul, Republic of Korea, 171–182. https://doi.org/10.1145/3543174.354525

[28] Francisco R. Ortega, Katherine Tarre, Mathew Kress, Adam S. Williams, Armando B. Barreto, and Naphtali D. Rishe. 2019. Selection and Manipulation Whole-Body Gesture Elicitation Study In Virtual Reality. In *Proceedings of the Conference on Virtual Reality and 3D User Interfaces*. IEEE, Osaka, Japan, 1723–1728. https://doi.org/10.1109/VR.2019.8798105

[29] Jorge-Luis Pérez-Medina, Santiago Villarreal, and Jean Vanderdonckt. 2020. A gesture elicitation study of nose-based gestures. *Sensors* 20, 24 (2020), 7118.

[30] Tran Pham, Jo Vermeulen, Anthony Tang, and Lindsay MacDonald. 2018. Scale impacts elicited gestures for manipulating holograms: Implications for AR gesture design. In *Proceedings of the Designing Interactive Systems Conference*. ACM, Hong Kong, China, 227–240. https://doi.org/10.1145/3196709.3196719

[31] Leon Pietschmann, Paul Zürcher, Erik Bubik, Zhutian Chen, Hanspeter Pfister, and Thomas Bohné. 2023. Quantifying the Impact of XR Visual Guidance on User Performance Using a Large-Scale Virtual Assembly Experiment. In *Proceedings of the International Conference Visualization and Visual Analytics*. IEEE, Melbourne, Australia, 211–215.

[32] Thammathip Piumsomboon, Adrian Clark, Mark Billinghurst, and Andy Cockburn. 2013. User-Defined Gestures for Augmented Reality. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Paris, France, 955–960. https://doi.org/10.1145/2468356.2468527

[33] Wim TJL Pouw, Jacqueline A De Nooijer, Tamara Van Gog, Rolf A Zwaan, and Fred Paas. 2014. Toward a more embedded/extended perspective on the cognitive function of gestures. *Frontiers in Psychology* 5 (2014), 359.

[34] Umair Rehman and Shi Cao. 2020. Comparative evaluation of augmented reality-based assistance for procedural tasks: a simulated control room study. *Behaviour & Information Technology* 39, 11 (Nov. 2020), 1225–1245. https://doi.org/10.1080/0144929X.2019.1660805 Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/0144929X.2019.1660805.

[35] Julie Rico and Stephen Brewster. 2010. Usable gestures for mobile interfaces: evaluating social acceptability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, Georgia, USA, 887–896. https://doi.org/10.1145/1753326.1753458

[36] Jaime Ruiz, Yang Li, and Edward Lank. 2011. User-defined motion gestures for mobile interaction. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Vancouver, Canada, 197–206.

[37] Jaime Ruiz and Daniel Vogel. 2015. Soft-constraints to reduce legacy and performance bias to elicit whole-body gestures with low arm fatigue. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Seoul, Korea, 3347–3350.

[38] Adwait Sharma, Joan Sol Roo, and Jürgen Steimle. 2019. Grasping microgestures: Eliciting single-hand microgestures for handheld objects. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–13.

[39] Milka Trajkova, A'aeshah Alhakamy, Francesco Cafaro, Rashmi Mallappa, and Sreekanth R Kankara. 2020. Move your body: Engaging museum visitors with human-data interaction. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu, Hawai'i, 1–13.

[40] Radu-Daniel Vatavu. 2012. User-Defined Gestures for Free-Hand TV Control. In *Proceedings of the European Conference on Interactive TV and Video*. ACM, Berlin, Germany, 45–48. https://doi.org/10.1145/2325616.2325626

[41] Radu-Daniel Vatavu and Jacob O Wobbrock. 2015. Formalizing agreement analysis for elicitation studies: new measures, significance test, and toolkit. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Seoul, South Korea, 1325–1334.

[42] Santiago Villarreal-Narvaez, Arthur Sluÿters, Jean Vanderdonckt, and Radu-Daniel Vatavu. 2024. Brave new gesture elicitation studies world: A systematic literature review of gestures and referents in gesture elicitation studies. *Comput. Surveys* 56, 5 (2024), 1–55.

[43] Adam S. Williams and Francisco R. Ortega. 2020. Understanding Gesture and Speech Multimodal Interactions for Manipulation Tasks in Augmented Reality Using Unconstrained Elicitation. *Proceedings of the ACM on Human-Computer Interaction* 4, 202 (2020), 1–21. https://doi.org/10.1145/3427330

[44] Jacob O Wobbrock, Htet Htet Aung, Brandon Rothrock, and Brad A Myers. 2005. Maximizing the guessability of symbolic input. In *Extended Abstracts on Human Factors in Computing Systems*. ACM, Portland, USA, 1869–1872.

[45] Jacob O. Wobbrock, Meredith Ringel Morris, and Andrew D. Wilson. 2009. User-defined gestures for surface computing. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, New York, USA, 1083–1092. https://doi.org/10.1145/1518701.1518866

[46] Huiyue Wu, Yu Wang, Jiali Qiu, Jiayi Liu, and Xiaolong (Luke) Zhang. 2018. User-defined gesture interaction for immersive VR shopping applications. *Behaviour & Information Technology* 38, 7 (2018), 726–741. https://doi.org/10.1080/0144929X.2018.1552313

[47] Jiang Xu, Patrick J. Gannon, Karen Emmorey, Jason F. Smith, and Allen R. Braun. 2009. Symbolic gestures and spoken language are processed by a common neural system. *Proceedings of the National Academy of Sciences* 106, 49 (Dec. 2009), 20664–20669. https://doi.org/10.1073/pnas.0909197106

[48] Momona Yamagami, Alexandra A Portnova-Fahreeva, Junhan Kong, Jacob O Wobbrock, and Jennifer Mankoff. 2023. How do people with limited movement personalize upper-body gestures? considerations for the design of personalized and accessible gesture interfaces. In *Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility*. ACM, New York, USA, 1–15.

[49] Yukang Yan, Chun Yu, Xin Yi, and Yuanchun Shi. 2018. Headgesture: Hands-free input approach leveraging head movements for hmd devices. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 4 (2018), 1–23.

[50] Xiaoyan Zhou, Adam Sinclair Williams, and Francisco Raul Ortega. 2022. Eliciting Multimodal Gesture+Speech Interactions in a Multi-Object Augmented Reality Environment. In *Proceedings of the Symposium on Virtual Reality Software and Technology*. ACM, Tsukuba, Japan, 1–10. https://doi.org/10.1145/3562939.3565637

A Most Common Gesture Groups per Referent And Scenario

	Control (avg. AgR = .136)		One-Hand-Fixed (avg. AgR = .138)		Two-Hands-Busy (avg. AgR = .127)		Refinement (avg. AgR = .163)	
Referent	**Gesture**	**Count (AgR)**	**Gesture**	**Count (AgR)**	**Gesture**	**Count (AgR)**	**Gesture**	**Count (AgR)**
ACCEPT/YES	Thumb Up	9 (.221)	Thumb Up	7 (.163)	Head Nod	15 (.558)	Head Nod	14 (.495)
DECLINE/NO	Thumb Down	6 (.189)	Thumb Down Head Shake	6 (.163)	Head Shake	15 (.558)	Head Shake	13 (.442)
HELP	Raise Hand	8 (.153)	Raise Hand	5 (.058)	Raise Hand	3 (.026)	Raise Hand	3 (.021)
OPEN MENU	Closed Fist Closed Pinch	4 (.068)	Closed Fist Swipe Up Closed Pinch	4 (.095)	Head Swipe Up	4 (.047)	Swipe Down	4 (.058)
SHOW ALL STEPS	Swipe Down	5 (.084)	Swipe R	4 (.053)	Head Draw 'Circle' Hands Spread Head Circle	2 (.016)	Hands Spread	3 (.026)
NEXT STEP	Swipe L	11 (.316)	Swipe L	14 (.484)	Head Swipe L Head Tilt R	5 (.137)	Swipe L	13 (.432)
PREVIOUS STEP	Swipe R	11 (.316)	Swipe R	14 (.500)	Head Swipe R Swipe R	4 (.111)	Swipe R	13 (.432)
CURRENT STEP	Stop	4 (.053)	Swipe Up Tap On	3 (.042)	Head Swipe Up	4 (.058)	Tap On	4 (.042)
HIDE DISPLAY	Point	8 (.184)	Swipe Down	6 (.121)	Head Swipe Down-L	5 (.037)	Swipe Down	4 (.063)
IDENTIFY	Question Wave	7 (.116)	Point	5 (.074)	Head Tilt Forward	4 (.058)	Point	4 (.053)
VERIFY	Draw Circle	3 (.032)	Thumb Up	3 (.032)	Head Swipe Up Knock On Head Tilt L	2 (.016)	Thumb Up Tap On Tap On x2	2 (.021)
DEMONSTRATE	Go On	2 (.021)	N/A	0 (.000)	Head Draw Circle	2 (.005)	Draw Circle	3 (.016)
ALTERNATIVE	Go On Draw 'X' Draw 'Semi-Circle'	2 (.016)	Go On Draw 'X'	2 (.011)	Tap On Head Tilt R Draw 'X'	2 (.021)	Go On	3 (.021)