

Labeling Out-of-View Objects in Immersive Analytics to Support Situated Visual Searching

Tica Lin, Yalong Yang, Johanna Beyer and Hanspeter Pfister

Abstract—Augmented Reality (AR) embeds digital information into objects of the physical world. Data can be shown *in-situ*, thereby enabling real-time visual comparisons and object search in real-life user tasks, such as comparing products and looking up scores in a sports game. While there have been studies on designing AR interfaces for situated information retrieval, there has only been limited research on AR object labeling for visual search tasks in the spatial environment. In this paper, we identify and categorize different design aspects in AR label design and report on a formal user study on labels for out-of-view objects to support visual search tasks in AR. We design three visualization techniques for out-of-view object labeling in AR, which respectively encode the relative physical position (height-encoded), the rotational direction (angle-encoded), and the label values (value-encoded) of the objects. We further implement two traditional in-view object labeling techniques, where labels are placed either next to the respective objects (situated) or at the edge of the AR FoV (boundary). We evaluate these five different label conditions in three visual search tasks for static objects. Our study shows that out-of-view object labels are beneficial when searching for objects outside the FoV, spatial orientation, and when comparing multiple spatially sparse objects. Angle-encoded labels with directional cues of the surrounding objects have the overall best performance with the highest user satisfaction. We discuss the implications of our findings for future immersive AR interface design.

Index Terms—Object Labeling, Mixed / Augmented Reality, Immersive Analytics, Situated Analytics, Data Visualization

1 INTRODUCTION

SEARCHING for elements of interest is an essential component for almost all visual analytics tasks, such as *identifying* outliers and clusters, *comparing* values encoded in visual elements, and *summarizing* properties for a group of objects. Designers of desktop visualization tools have grappled with this problem for decades and have developed sophisticated techniques to support visual search. To make search results pop out in a visualization, it is common to show icons or labels overlaid on top of the original visualization (e.g., restaurants on a map). To avoid occluding local visualization features, external labeling techniques [1] place labels outside the visualization, and use leader lines to connect the visual elements with their associated labels. Other techniques allow users to zoom out to an overview to locate the searched for elements, and then zoom in to check the details. These techniques are well studied and widely used in desktop visualization tools.

With the rapid development in technology, consumer-level head-mounted displays (HMDs) for virtual reality (VR) and augmented reality (AR) have become mature and affordable. In response to this revolution, immersive analytics has emerged as a new research field, focusing on the use of engaging, embodied analysis tools in VR and AR, to support data understanding and decision making [2]. VR/AR removes the boundaries of physical screens and facilitates new human-computer interaction experiences. In particular, AR can embed digital information in the physical world in almost any environment. Thereby, AR provides opportunities to support situated visual search for real-life tasks [3], such as showing ratings of nearby restaurants [4], adding

historical information for artifacts in a museum [5], or annotating player names and stats in a basketball game [6], [7], all without diverging users' attention from the real-world environment.

In these AR applications, additional information for the physical objects is commonly rendered in small “virtual canvases” or labels. However, compared to the desktop, visual searching in AR has two distinct differences. First, interface designers cannot manipulate everything in the AR display space. While labels can be moved freely in AR, physical objects (e.g., the locations of the restaurants, the basketball players on the court) usually cannot be manipulated. As a result, users have to physically move or rotate to get to or face the physical targets. Second, unlike desktop displays where visual elements can be scaled to fit the screen, physical objects in the real world cannot be scaled (e.g., again, a restaurant or a basketball player). Due to these two distinct differences, some desktop techniques for visual searching, such as zooming, cannot be used in AR applications, and more emphasis needs to be put on efficient object labeling techniques. Therefore, further studies are needed to understand **how to place labels in AR to effectively support situated visual search**. Currently, there are still many open questions, such as whether labels should be placed as close to their linked objects as possible, or whether labels should also encode spatial information of objects that are inside or outside the user's FOV to help reduce a user's search effort.

To address these questions, we first thoroughly review existing labeling techniques in AR. Previous studies on labeling objects in AR focused on algorithms to place labels [8], [9], [10], [11], [12] and filter data [13] to avoid visual clutter on limited AR FoV. Most labeling techniques only show labels for objects that are within the user's FoV, so-called in-view objects. In many real-world AR scenarios, however, the objects of interest are not just in front of the user. Sometimes the user may not even know the locations of the target objects. Without any additional clues, the user will then have to physically scan the entire surrounding environment to find

- Tica Lin, Johanna Beyer, and Hanspeter Pfister are with John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, 02138. E-mail: {mlin, jbeyer, pfister}@g.harvard.edu
- Yalong Yang is with the Department of Computer Science, Virginia Tech, Blacksburg, VA, 24060. E-mail: yalongyang@vt.edu

Manuscript received April 19, 2005; revised August 26, 2015.

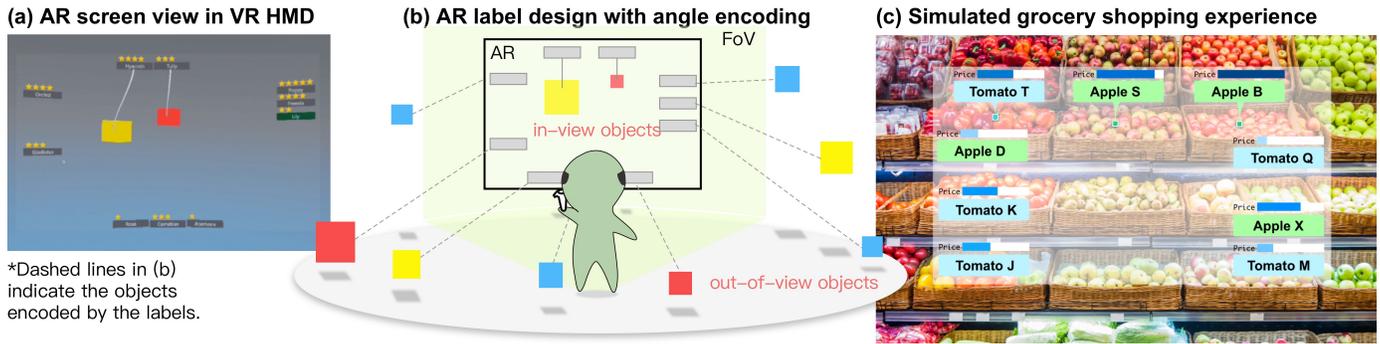


Figure 1: AR label design for objects outside the field-of-view (FOV). (a) Our user study uses a VR HMD to simulate consistent AR conditions. (b) The user is surrounded by spatially sparse objects and can see labels for in-view and out-of-view objects on the AR screen. Labels for out-of-view objects are placed on the boundary to support embodied navigation. (c) A simulated grocery shopping experience with AR labels.

an object and its label, which is a time-consuming and fatiguing task. To support efficient situated visual search tasks, labels for out-of-view objects (i.e., physical objects that are not within the FoV) is thus necessary. Although previous work studied visual encoding techniques for out-of-view objects in AR [14], [15], [16], [17], [18], not much research has focused on labels and label design for out-of-view objects. In this work, we explore different visual encodings for AR labels and quantitatively compare them for situated analysis tasks in a user study.

The motivation for this paper is to investigate in which situations out-of-view labels are helpful, and whether certain tasks would benefit from only providing in-view labels. We specifically focus on labeling static objects as a preliminary step towards answering this question. Labeling static objects has many real-world use scenarios, such as goods in grocery stores or restaurants around a person (see Sec. 3.3). Labeling moving objects, on the other hand, requires other design considerations, which we briefly discuss in our future work (Sec. 7.3).

Our first contribution is a **systematical exploration of the design space of labels in AR applications, with an emphasis on labeling out-of-view objects**. We analyzed previous works on AR labeling [8], [12], embodied interaction [19], and leader line placement [11], and identified four different aspects in AR labeling systems (label proxy, proximity, encoding, and interaction), which can guide the design of future AR labeling systems.

Our second contribution is a **controlled user study comparing five representative AR labeling techniques**. In two conditions we show labels for in-view objects only: *situated*, where labels are placed slightly above the physical object, and *boundary*, where labels are pushed to the boundaries of the FoV to reduce visual clutter. We also designed three conditions for labeling of out-of-view objects: *height*, where labels are placed at the same height as their linked objects will appear once the user rotates toward them; *angle*, where label positions indicate the angle between the user's viewing direction and the object, using a top-down view metaphor; and *value*, where labels are ordered by their associated values (e.g., restaurant ratings) on the left of the FoV. Following previous work in AR research [20], [21], [22], to easily manipulate the locations of physical objects, we used VR to simulate AR applications in our study. To achieve this, we created an AR FoV in VR, and render labels and leader lines only when they are within this AR FoV. In our study, we limit the number of labels to 10 and 20 objects. This is based on the number of AR labels used in other user studies [11], [23], and also follows the design guideline of avoiding information overload in

AR applications [24], [25], [13]. For labeling a larger number of objects, label clustering and filtering techniques have been proposed [13], but this is outside the scope of our study.

We evaluated our five different label conditions in three visual search tasks. Overall, we found that labels for out-of-view objects helped users in getting a quick overview, provided helpful directional cues, and was beneficial for visual comparisons in AR. We found that the *angle* condition was the overall winner; it was faster than *value* in all tested tasks, faster than *situated* and *boundary* labels in value comparisons, and faster than *height* in summarizing multiple data points. The *angle* condition was also preferred by our participants overall. Our work contributes to the growing body of knowledge of effective AR interface design to support situated analysis in the real world.

2 RELATED WORK

In the following, we review the literature on situated visual search to motivate the use case of AR labels, and point out the gap in previous AR labeling and visualization techniques.

2.1 Situated Visual Search in AR

Visual search is a fundamental perceptual task that people perform to identify a particular target among other distractors in a visual environment [26]. In real-world visual search tasks, information relevant to the location or object can be placed *situated* to support target comparison and spatial navigation [27]. Many design frameworks and use cases have been proposed that exploit the situatedness [4], [28] and embodiment [19] of AR visualizations. Buschel et al. [29] proposed a conceptual framework for annotating real-world objects with situated labels and accessing information through interaction with the labels. Bach et al. [30] proposed AR-canvas, a framework for designing embedded visualizations for visual search in realistic use scenarios. Attention management research [31], [32] has shown that high information density can lead to information overload, especially in real-world use cases [33], [34]. Managing information density on AR screens can be achieved with spatial and knowledge-based filtering [25], [24], and data clustering [13]. Building on the existing AR information retrieving workflow, our study is a concrete step towards designing labels in immersive environments to support situated visual search.

2.2 Labeling Objects in AR

In AR settings, we distinguish between labels for static and labels for dynamic objects that change over time.

Earlier studies have focused on managing label placement for static 3D landscape models [8] viewed from a top-down viewing angle. Subsequent work improved algorithmic performance to create embedded [9] and external labels [35] for landscape models in interactive 3D virtual space. Azuma et al. [10] proposed a cluster-based algorithm to place dense occlusion-free AR labels of static in-view objects in real-time. Tatzgern et al. [11] designed a depth-encoded hedgehog labeling technique to connect labels to their targets. Using similar techniques, Madsen et al. [23] studied the effect of label rendering space and update frequency on locating labels. Both studies used around ten labels for in-view objects on an AR tablet.

For dynamically changing scenes, Makita et al. [36] studied view management for annotating a small number of moving objects (less than five) in wearable AR devices. Orlosky et al. [37] resolved annotations overlapping with objects in AR using image detection and halo layout design for up to six labels. Grasset et al. [12] proposed image analysis to place location-based external labels on mobile AR browsers for up to 30 labels, but working best with 10 labels to avoid visual clutter. Furthermore, Tatzgern et al. [13] proposed dynamic data clustering and filtering to manage a larger number of annotations with adaptive label displays in AR.

2.3 Out-of-View Object Visualization

While labeling for out-of-view objects has only been preliminary explored, approaches for visualizing out-of-view objects have been widely studied. A main difference between the two is that labels generally support a two-way retrieval workflow (i.e., object-to-label and label-to-object look-up), while out-of-view object visualization focuses on a one-way retrieval to allow users to find the object quickly. Out-of-view object visualization can be classified into *overview+detail*, *focus+context*, and *detail-in-context* techniques [38]. For AR HMD environments *detail-in-context* approaches are most suitable, which includes object labels. *Overview+detail* with an extra miniature map can introduce visual overload on the AR screen and *focus+context* methods with view distortions [39] can cause misinterpretation and difficulty in locating the target [40].

2D visualization. Previous *detail-in-context* visualization approaches have used various visual proxies to encode spatial information for out-of-view objects for large graph navigation and notifications [41], [42]. Wedge [43], Halo [42], [44], Arrow [44] and EdgeRadar [45] are well-known 2D visualization techniques that use abstract shapes as proxies to encode spatial information of off-screen targets. A comparable concept for designing labels for out-of-view objects is encoding targets as insets on a map or large network visualization [46], [47], [48]. Ghani et al. [47] designed dynamic insets to encode out-of-view targets on the boundary.

AR/VR visualization. Out-of-view object visualization techniques in AR/VR environments build upon 2D visualization techniques, such as 3D Arrows [49], [50], [51], [52], 3D Halo [53], attention funnel [27], and radar projection [17], [15].

Several user studies have applied and compared these techniques in target searching tasks in AR/VR. Petford et al. [54] compare wedge visualizations to other attention-guiding techniques such as WIM [55] in a projector-based display environment for a single out-of-view target. Gruenefeld et al. [56] further apply Halo [42] and Wedge [43] techniques to AR and VR HMD and evaluate their performance in multiple (up to eight) out-of-view target search and direction estimation tasks. Bork

et al. [18] investigate visual guiding techniques for searching spatial virtual objects in MR with AR HMD. They compared six out-of-view object visualization techniques in a controlled user study, including four previous techniques, 3D Arrows [14], AroundPlot [15], EyeSee360 [17], sideARs [16], and two novel techniques, 3D Radar and Mirror Ball, where users are asked to use the visualization to collect up to eight out-of-view objects that are spatially distributed. They found that 3D Radar and EyeSee360 are fastest when guiding users to collect objects outside FoV. Gruenefeld et al. [57] used radar-like visualizations to encode out-of-view objects on AR HMDs. They compared different methods for showing locations for in-view objects, out-of-view objects, or both, but focused on the actual objects and did not use labels. They evaluated direction guiding and object selection tasks for eight objects. Most existing out-of-view object visualization techniques suffer from occlusion and edge clutter with multiple out-of-view objects, and are not suitable to encode labels. They also primarily focus on attention guiding and searching for a single or a few out-of-view targets. Other visual analysis tasks such as value comparison have not been considered.

In our study, we explore both in-view and out-of-view object visualization with AR labels to support situated analytic tasks. Two in-view object labeling techniques (SITUATED and BOUNDARY) are based on external labeling techniques. We designed three out-of-view object labels building upon previous out-of-view object visualization techniques. HEIGHT label extends the concept of dynamic insets [47] and Wedge [43] to map an object's physical position onto the AR screen boundary. ANGLE label can be seen as a top-down view of radar projection mapped onto the AR screen boundary. VALUE label encodes an object's data value and ranks labels by values to support situated analysis.

3 LABEL DESIGN SPACE & TASKS

A *label* is a visual representation of a text-based annotation attached to an *anchor* on the target object. In the following, we discuss the most important properties of label design for immersive analytics, define goals and tasks for labels in situated search tasks, and outline some usage scenarios of AR labels.

3.1 Properties of Labels for Situated Analysis

We define four distinct factors of label design in AR for situated analytics based on the relationship between labels, objects and the user, shown in Fig. 2. In addition, we identify other considerations that should be taken into account when designing labels in AR, such as object attributes and subjective design considerations.

Previous work [12], [58] has outlined rule-based properties of label layout in AR/VR scenes, such as that labels should be occlusion-free and placed close to the referred objects. While these are properties required to create a visually aesthetic and useful label layout, we focus on label properties that need to be considered in situated analytics tasks.

Label Proxy. Labels for objects in space need a visual representation. Proxies can include textual or visual information of the linked objects. Labels can be shown for visible objects (in-view), invisible objects (out-of-view), or both. In immersive environments, a label can be *embedded* inside the target or linked to the target through a leader line as an *external* label. However, embedding labels accurately into objects can be computationally expensive due to the heterogeneous visual appearance of real-world objects and fluctuating view conditions [59]. Furthermore,

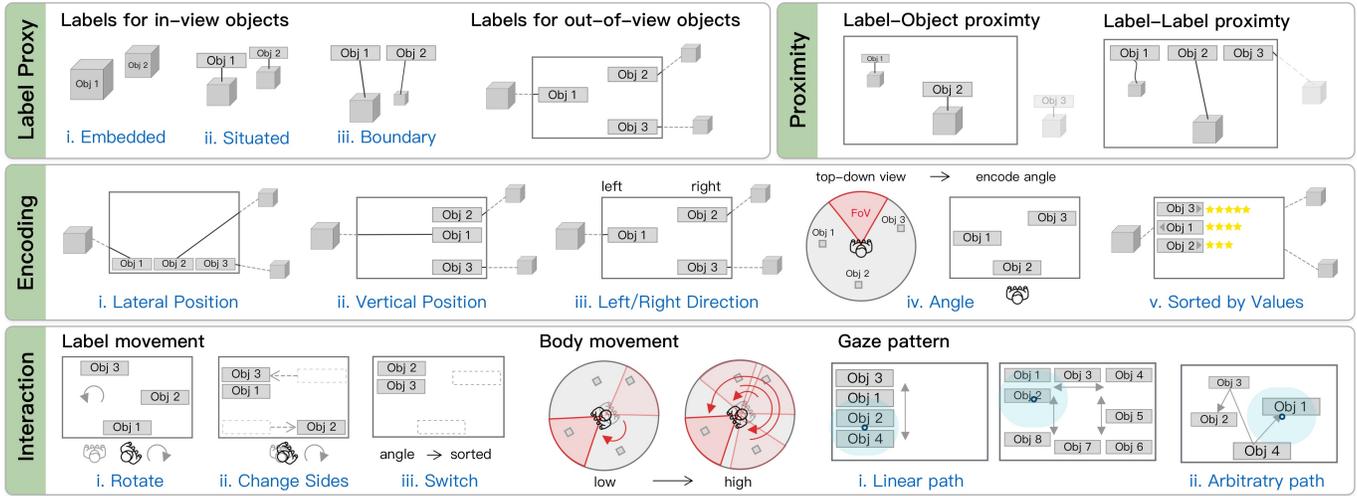


Figure 2: Label design space in immersive analytics. We categorize four factors of label design for situated target search, including label proxy, proximity, encoding, and interaction.

label embedding cannot be applied to out-of-view objects. In our study, we focus on *external* labels to explore a generalizable label design for real-world situated visual search tasks.

Proximity. Proximity describes the closeness between two labels or between a label and its object. Traditional AR/VR labels for in-view objects are optimized for label-object proximity. A higher proximity between label and object allows easier linking between the two. However, labels can be designed to prioritize label-label proximity to support efficient visual comparison, such as comparing the values on all labels.

Encoding. Label positions can encode additional attributes such as the lateral or vertical position of an object with respect to the user’s orientation/gaze. Alternatively, a label position can also encode the left/right direction (i.e., whether the object is to the left or to the right of the current view), angle (i.e., the angle between the user’s view direction and the object), or the data value of the object.

Interaction. To support spatial search, labels are designed to respond to a user’s interaction and movement. Label movement describes if the label placement remains static or updates based on user movement. On the other hand, to use the labels in visual search tasks, the user has to perform body or gaze movements. Body movement reflects the degree of the user’s physical movement required in each label design; gaze pattern classifies the direction of gaze movement in order to access the labels, including linear and arbitrary paths. For example, a user’s gaze pattern is linear if labels are aligned vertically or horizontally, and arbitrary if labels are situated at the object’s position.

Other Considerations for AR label design. In addition to the factors outlined above, there are other considerations that influence the design of efficient labels that focus on *object attributes* and *subjective measures*. **Object attributes** include the number of objects (sparse versus dense), the spatial distribution of objects (full 3D space versus restricted to certain areas such as the half-space above ground), and object movement (static versus dynamic). **Subjective measures** include the user’s familiarity with a chosen label design as well as the predictability of label locations. For example, if labels are always located at the lower screen boundary, users will have an easier time locating them.

3.2 Goals & Tasks

Inspired by previous work in AR information retrieval and situated analytics [28], [27], [29], [30], we have identified two primary goals of label design in AR and for situated visual search tasks.

G1 - visualize the spatial relationship between the objects and the user to support situated spatial search. Users should be able to explore all objects and locate a target in space more easily with the provided labels, for in-view as well as out-of-view objects.

G2 - provide extra information for objects to support visual analysis tasks in the real world that would otherwise be extremely tedious or impossible to perform. Such visual search tasks include but are not limited to identifying outliers, comparing multiple targets, or making overall estimations.

We identify common visualization tasks to support situated spatial and visual search. Based on the typology of abstract visualization tasks by Brehmer and Munzner [60] “the user must find elements of interests in the visualization”. Labels in AR already encode an object’s *identity* and point towards the *location* of the referred object. Therefore, labels can be seen as the result of an information retrieval process [61], [29], and situated visual search requires interpreting available labels in the AR scene. According to Brehmer and Munzner’s typology [60], labels for situated analysis need to support the three low-level querying actions, IDENTIFY, COMPARE, and SUMMARIZE.

T1: IDENTIFY a single target among all objects of interest. Users need to be able to identify a single target from all objects of interest in the environment. Typically, users need to physically examine each object in the environment to decide on a single target. Labels in AR can provide a quick overview of all objects to help identify a single target (e.g., a product with the lowest price), and spatial cues to guide the user to the desired target.

T2: COMPARE between multiple targets. Users need to perform comparisons between multiple targets in the environment. Typically, users have to examine and memorize the object-specific data for each object. With labels, this task can be done with visual comparison from the labels, such as comparing the prices and calories of similar products.

T3: SUMMARIZE data across all targets. To get an overview of the scene or perform comparisons at a larger scale, users need to summarize data across targets. This task is very challenging

in a real-world scenario, such as comparing the average prices of one brand against another in a store, or house prices in two neighborhoods. With labels, users can summarize the trend or identify outliers among all available targets.

3.3 Usage Scenario

We exemplify the identified tasks of Sec. 3.2 with real-world usage scenarios of AR labels.

Grocery shopping. A person shopping in the supermarket needs to know the location of each item on the shopping list. Instead of looking for items aisle by aisle, they look up the labels on an AR HMD. Labels show the categories of items and encode the directions of each category that is outside of the current view. They *identify* “Fruits” and follow the label to directly locate the aisle. Similar to Fig. 1 (c), AR labels show the price of different fruits and varieties with color-coded bar charts. The shopper *compares* their prices right in place without having to examine each item. They also estimate the price difference between apples and tomatoes by *summarizing* the prices shown on the labels. With the labels to support situated analysis, they can make data-driven decisions without much effort.

Spatial navigation. A traveler standing in front of a tourist site is deciding where to go for lunch in the surrounding area. The search results are shown as labels on the AR HMD and are encoded with directions. Among all labels, the tourist *summarizes* the average ratings of Japanese and Italian restaurants. They decide on Japanese food and use the labels to *compare* the prices of individual restaurants. While making a decision, they are able to continue to explore the surrounding area. Finally, they *identify* the target restaurant and follow the label direction to the destination.

4 LABEL PLACEMENT DESIGN

Based on our label design space, we explore different combinations of label placement properties as shown in Table 1, including two variations of labels for in-view objects – SITUATED and BOUNDARY, and three variations of labels for out-of-view objects – HEIGHT, ANGLE and VALUE. Each label design consideration is discussed below and shown in Fig. 3 and the supplemental video.

4.1 SITUATED

SITUATED is the most common labeling technique in existing AR/VR applications [8], [35], [12], such as a name tag and health bar placed above an avatar in a video game. We design SITUATED labels as a baseline condition in our user study. SITUATED labels are placed directly above their referred objects with a straight leader line connecting the object center to the bottom center of the label (Fig. 3 (1)). When objects are too close, we adjust the vertical label positions to avoid overlaps. The close proximity between labels and objects makes it clear to which object a label is connected to and reduces time and effort to identify the object. Due to the familiarity of similar SITUATED labels in existing 3D environments, it feels natural to users. Therefore, SITUATED labels have the advantage of **high label-object proximity** and **high familiarity**.

On the other hand, SITUATED labels require **high body movement** because the label placement does not encode any spatial information of the objects. It requires users to do a full-space scan to search for out-of-view objects. When scanning the entire surrounding space, it is also difficult to find labels because they

Factor	<i>Situated</i>	<i>Boundary</i>	<i>Height</i>	<i>Angle</i>	<i>Value</i>
Label					
In-View	✓	✓	✓	✓	✓
Out-of-View	✗	✗	✓	✓	✓
Proximity					
Label-Label	low	high	low/med.	high	high
Label-Object	high	med.	high	med.	low
Encoding					
Lateral position	✓/ -	✓/ -	✓/ ✗	✓	✗
Vertical position	✓/ -	✗/ -	✓	✗	✗
Direction	✗	✗	✓	✓	✓
Angle	✗	✗	✗	✓	✗
Value	✗	✗	✗	✗	✓
Interaction					
Label movement	✗	✗	✓	✓	✗
Body movement	high	high	med.	low	med.
Gaze pattern	arb.	linear	arb./lin.	linear	linear
Subjective					
Familiarity	high	low	med.	low	med.
Predictability	low	high	low	high	high

*arb. = arbitrary, lin. = linear, med. = medium

Table 1: Characteristics (rows) of five AR label placement designs (columns) in our study. The slash indicates different properties of labels for in-view and out-of-view objects.

appear at different vertical heights and require more **arbitrary gaze patterns**. In particular, when comparing multiple spatially sparse objects, users have to remember the information on the labels as well as the locations of the corresponding objects. This **low predictability** of label locations makes the comparison between multiple targets error-prone and inefficient.

4.2 BOUNDARY

BOUNDARY labels are an alternative design of in-view labels, with a focus on visual search and comparison tasks. As shown in Fig. 3 (2), we align BOUNDARY labels horizontally at the bottom edge of the AR screen and encode the horizontal positions of the linked in-view objects. We prioritize objects at the center of the view and push overlapping labels towards the left or right to avoid overlaps when necessary. The one-dimensional label placement makes it easy to quickly scan all labels in one direction without having to change the vertical gaze pattern to find labels. Labels are also closer together, which makes label-to-label comparison more efficient. This helps users to get a faster overview and reduces gaze movement. Label locations are more predictable as labels are always placed at the bottom.

We originally designed the labels to be aligned at the top, but initial user feedback suggested that users found it uncomfortable to constantly have to look up with their eyes. Therefore, we place boundary labels at the bottom to mimic the conventional task bar or dock placement on a computer or mobile device. Summarizing, BOUNDARY labels have the advantage of **label-label proximity**, a **linear gaze pattern**, and **high predictability**.

Similar to SITUATED, BOUNDARY labels also require **high body movement** to complete visual search and comparison tasks. Furthermore, they require a separate step to link the label to the object due to **lack of label-object proximity**. The longer leader lines make it more difficult for users to retrieve the linked object and can cause visual clutter when there are multiple objects in view and lines are overlapping with the objects in the scene.

4.3 HEIGHT

In our study, we use three encodings for out-of-view object labels. Compared to label design for in-view object labels, they all require

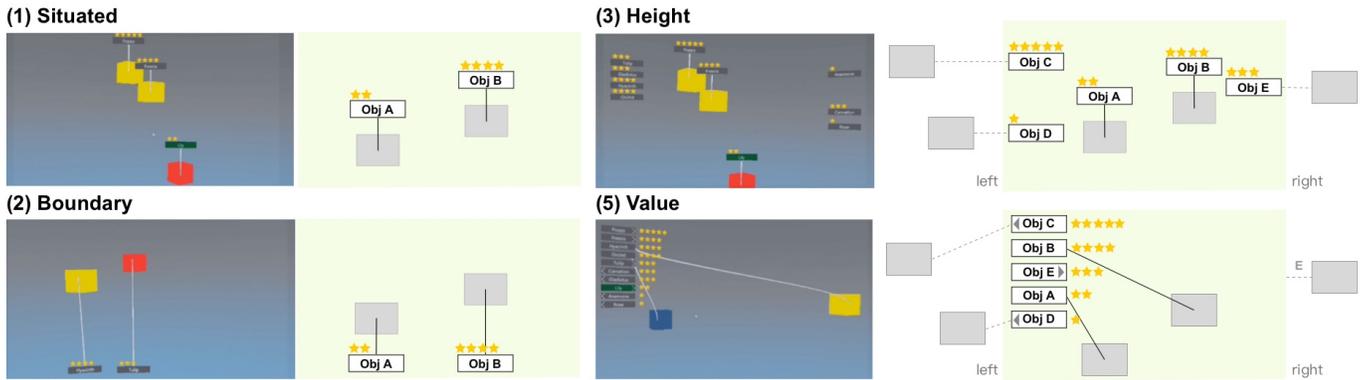


Figure 3: AR screen views (left) and diagrams (right) of the four label design in Sec. 4. ANGLE design is shown in Fig. 1 (a) and (b).

less body movement and support a linear gaze pattern to track labels, but more labels also result in higher visual density.

As shown in Fig. 3 (3), HEIGHT labels are placed at the left and right boundaries of the view, based on the closer lateral direction of the target object. Furthermore, their y position encodes the relative *height* at which the linked object will come into the screen as the user rotates. Showing labels for all available objects in the scene allows users to search for and compare targets without physically moving around. The height-encoded label placement provides cues to both the lateral direction and vertical position of an object, and thus users can know immediately which way to turn to when searching for a single target. For example, seeing a label in the top left border means the user will turn left and look upward to find the target. When the object appears in view, the height-encoded label will become SITUATED and directly lead the users to the target object. HEIGHT labels provide an **encoding for the vertical position and lateral direction** of out-of-view objects and exhibit **high label-object proximity**.

Since HEIGHT labels are arranged along the left and right borders of the view, based on which border the target object is closer to, their placement on the screen is changing as people turn. For example, objects that were closer to the left border might become closer to the right border as the user rotates. While **label movement** provides real-time directional cues on objects, it requires users to search for and constantly track out-of-view labels (i.e., labels for out of view objects) while rotating and moving their gaze. Due to the lack of distance information (i.e., how far outside the field of view an object is), users have to follow a label at all times to find the object, which can cause higher mental stress and slightly **more body movement**. Furthermore, when comparing multiple out-of-view objects, it takes extra mental effort to search for and compare labels, as they are not ordered by angles or values.

4.4 ANGLE

ANGLE labels are placed along all four edges of the user's view boundary and encode the *angular* direction of all available objects as seen from a top-down view, as shown in Fig. 1 (b). The top boundary represents the front-facing direction (i.e., the direction the user is currently looking at) and the bottom boundary represents the area behind the user. Therefore, all labels for in-view objects are naturally placed along the top boundary and linked to their target objects with leader lines. Labels for out-of-view objects are placed along the left, bottom, and right view boundary, depending on their position in space. As users turn around, the ANGLE labels rotate along the boundary to reflect the relative

angular direction of the objects to the user. Therefore, ANGLE provides a precise encoding of the angular distance between the user and the target object, as well as between objects. This can increase the spatial awareness of users and reduces mental and physical efforts when locating objects. Similar to BOUNDARY, ANGLE label placement is very predictable and aligned vertically and horizontally, which allows easier label-to-label comparison. As users move, ANGLE labels rotate in either clockwise or counter-clockwise direction and hence do not distract users as much as in the HEIGHT condition. ANGLE labels provide an encoding for an object's **lateral direction and angle**, and exhibit **label-label proximity**, and **high predictability**.

On the other hand, users might have **low familiarity** with angle-encoded labels, which can be difficult to interpret without prior training. When looking for a target object, the longer leader lines between label and object might cause more cognitive effort to link the object and the label due to **lack of label-object proximity**. Similar to HEIGHT labels, when comparing multiple labels, it takes time to find and sort the label values.

4.5 VALUE

As shown in Fig. 3 (5), VALUE labels are placed stationary along the left view boundary and ordered vertically by the label's *value*. In order to stack more labels and allow more direct value comparisons similar to a horizontal bar chart, we chose a parallel label layout where the label's value is shown to the right of the label's text. We indicate the direction of the linked objects with left/right arrow icons on the labels. When the linked object is in view, the arrow icon disappears and a leader line appears. Label placement is very predictable as people only have to look at the left boundary to look for labels and then follow the arrow to find the object. When comparing multiple labels, VALUE has the advantage of efficient value comparisons due to the fixed label layout and implicit ranking. Users only have to move their gazes vertically to make cross-label comparisons. Therefore, VALUE labels provide an encoding for an object's **value and direction**, and exhibit **label-label proximity**, and **high predictability**.

While VALUE labels have a stable and predictable placement, it requires an extra step to interpret the arrow icons and follow the leader line to find the objects due to **lack of label-object proximity**. Encoding spatial information in icons is more subtle than encoding it as label positions, which might require more mental effort to interpret and **more body movement** than ANGLE labels. Furthermore, leader lines of different labels might exhibit

line crossings due to the label’s vertical position being independent of the object’s vertical position.

4.6 AR Interface Implementation

We developed and implemented all label designs using Unity v2020.1.12 [62] and the Mixed Reality Toolkit v2.5.1 [63]. In order to vary object distribution and achieve higher tracking fidelity of the objects in the study, we use VR to simulate the real-world object space and AR screen. Simulating an AR environment with VR is a common method to create a controlled environment in a user study [20], [21], [22].

We set up our AR interface as a semi-transparent canvas placed 1.8m in front of the user with a $35^\circ \times 25^\circ$ FoV. We display the AR screen in the users’ FoV at a fixed location, and, therefore, the labels on the AR screen fall on the periphery of the user’s vision (see Fig. 1 (a)). Our chosen AR screen size falls within the FoV of the current state-of-the-art AR HMDs, Magic Leap 1 ($40^\circ \times 30^\circ$) and HoloLens ($43^\circ \times 29^\circ$). We apply the tooltip component of the Mixed Reality Toolkit UX building block [64]. The leader lines are designed to curve based on the depth between the object and the label on the AR screen to provide a sense of depth, such as in Fig. 3 (5). We place all objects around the user in 360 degrees. Details on data generation are discussed in Sec. 5.2.

5 USER STUDY

We evaluate the usefulness of our five label designs in supporting situated visual search tasks in a controlled user study. In particular, we want to investigate in which situations out-of-view labels are helpful, whether in-view labels have any advantages, and the trade-offs between the two.

5.1 Experimental Design

Labeling conditions. We evaluated the five label conditions described in Sec. 4, including SITUATED, BOUNDARY, HEIGHT, ANGLE, and VALUE. In all conditions, each label shows the unique name of the object and a rating ranging from 1 to 5 stars.

Data sizes. We generated datasets with two data sizes – 10 and 20 objects. Previous studies on label design for in-view objects in tablet AR used around 10 labels [11], [23]. We also constrain the amount of information in the AR FoV to avoid information overload and occlusion [33], [34]. 10 and 20 objects are chosen for different real-world use cases, such as comparing a fixed set of targets and exploring a broader set of options, respectively.

Experiment Set-Up. We conducted our study in an indoor lab space of approximately 385ft^2 . We used an Oculus Quest virtual reality headset, with a $96^\circ \times 94^\circ$ FoV, 1440×1600 pixel resolution, and a weight of 571g to render objects, labels, and the semi-transparent AR screen. The headset was connected to a PC through a 5m USB3 Type-C cable. The PC had an Intel i7-9700F 3.00GHz processor and NVIDIA GeForce RTX 2070 graphics card.

To reduce motion sickness and maintain orientation in VR, we rendered a $4 \times 4\text{m}$ virtual floor and an arrow at the center of the virtual floor indicating the user’s starting position (i.e., front-facing direction). The study had a stationary set-up and required minimum body movement. Users did not need to walk around but only had to turn, move their head and use one controller to interact with the question pop-up windows. Participants were required to return to the front-facing direction before each trial.

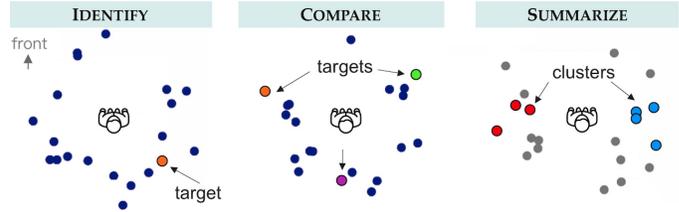


Figure 5: Example datasets for all three user tasks showing the spatial distribution of objects. Each small circle represent the object’s x and y location in the trial in a top-down view (object height is not shown). The selected target(s) for each task are highlighted.

5.2 Task & Data

We used three visual search tasks (see Sec. 3.2) based on the typology of abstract visualization tasks by Brehmer and Munzner [60].

IDENTIFY: *What is the color of the object linked to the green label?* We colored the target object’s label in green and all others in grey. All objects were randomly colored in either red, yellow, or blue. Participants had to identify the color of the object linked to the green label and select the answer in the pop-up menu.

COMPARE: *What is the color of the object with the highest rating?* Three labels were colored in green, all others in grey. All objects were randomly colored in either red, yellow or blue, while each of the three linked objects was assigned a different color. Participants had to compare the ratings of the three green labels and find the color of the object with the highest rating.

SUMMARIZE: *In the two colored clusters, which cluster has a higher average rating?* We designed two spatially separate clusters, with three objects in the red cluster and four objects in the blue cluster. Initially, the clusters were not shown, and all objects were colored in grey. Two selected labels were colored in red and blue, the others in grey. Participants first had to locate and click on the two objects linked to the blue and red labels to reveal the clusters. The labels and the objects of each cluster were colored in red and blue. They then determined which cluster had the higher average rating: red, blue or equal. The task represents a potential user flow of identifying two targets of interest and expanding the search to their immediate neighborhood.

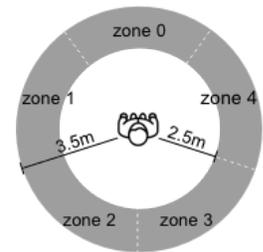


Figure 4: Object locations.

We generated a distinct dataset with 10 or 20 spatially sparse objects for each trial in our study. Each object is represented as a $20\text{cm} \times 20\text{cm}$ cube, with a minimum distance of 40cm between cubes. All objects were randomly placed within 2.5m to 3.5m from the user at a viewing height between 0.5m to 2m . To distribute the objects in space, we divided the circular space around a user into five zones (see Fig. 4) and ensured that each zone contained a similar number of objects. Each object’s label was randomly assigned a unique name selected from a list of items, such as fruits and flowers, and a rating value between 1 and 5. To control for the different reading speeds of our participants, we used color to highlight the targets in our tasks instead of relying on text labels.

We show sample data for each task in Fig 5. For the IDENTIFY task, we randomly selected an object in zone 2 or 3 as the target object to avoid showing the target immediately in front of the user. For the COMPARE task, we randomly selected three objects from

different zones and assigned a non-repetitive rating value from 1 to 5 and a distinct color from red, yellow and blue. We intentionally spread out the target objects in space so that no two objects being compared could be seen in the AR view at the same time. We also did not assign 5-star ratings to the targets being compared during the trials to make the task more challenging. For the SUMMARIZE task, we selected two objects from different zones to be labeled with red and blue labels at the beginning of the task. We selected objects close to the red and blue object as the red and blue clusters respectively, which resulted in three red objects and four blue objects in each trial. We then assigned rating values to each object so that the resulting average of each cluster was an integer.

5.3 Participants & Procedures

Participants. We recruited 15 participants from the University mailing list. Due to the COVID-19 pandemic and restrictions to campus access, we limited participation to university students and staff only. Participants ranged in age from 18-34 years. Five identified as female and ten as male. Seven participants had prior experience with VR HMDs, and one had experienced AR HMDs. **Procedure.** We followed a full-factorial within-subject study design to account for individual differences between participants. We balanced the order of label conditions using a Latin square (5 groups), and fixed the task order by increasing task complexity: IDENTIFY, COMPARE, and SUMMARIZE.

Each condition (label \times task) included 2 training trials and 6 timed study trials. Each participant completed 90 study trials: 5 label conditions \times 6 trials \times 3 tasks. Participants first filled out a consent form and were introduced to the study procedure and all five label designs with slides. Next, the instructor helped participants set up the VR headset and showed them how to select an answer with the trigger button. Before each task, the instructor introduced the task and reminded participants to perform the task *as precisely and as quickly* as possible during the timed trials. Participants were encouraged to spend as much time as needed on training. They could take breaks between each trial and task. After each task, the instructor collected participants' oral feedback. After the final task, participants filled out a post-study questionnaire. The whole study took 1.5 hours to complete, and each participant was compensated with a \$20 gift card.

5.4 Measures

We recorded the *time* performance of each trial from starting the visualization to the time participants double-clicked to input their answers. For *accuracy*, we compared their answer to the correct answer. We also collected participants' subjective feedback for each task and their overall evaluation in the end with a post-study questionnaire, including a standard NASA-TLX survey, qualitative feedback, and subject rankings for the five labeling conditions. To extract insights from the qualitative feedback, we derived a set of codes in an open coding session among three authors for the first five participant responses. The first author then applied the set of codes to the feedback of the remaining 10 participants.

5.5 Statistical Analysis

For dependent variables or their transformed values that can meet the normality assumption (i.e., time), we used *linear mixed modeling* to evaluate the effect of independent variables on the dependent variables [65]. Compared to repeated measure ANOVA,

linear mixed modeling is capable of modeling more than two levels of independent variables and does not have the constraint of sphericity [66, Ch. 13]. We modeled all independent variables (i.e., label placement techniques and data sizes) and their interactions as fixed effects. A within-subject design with random intercepts was used for all models. We evaluated the significance of the inclusion of an independent variable or interaction terms using log-likelihood ratio. We then performed Tukey's HSD post-hoc tests for pair-wise comparisons using the least square means [67]. We used predicted vs. residual and Q — Q plots to graphically evaluate the homoscedasticity and normality of the Pearson residuals, respectively. For dependent variables that cannot meet the normality assumption (i.e., accuracy, NASA-TLX ratings), we used a *Friedman* test to evaluate the effect of the independent variable, as well as a *Wilcoxon-Nemenyi-McDonald-Thompson* test for pair-wise comparisons. Significance values are reported for $p < .05(*)$, $p < .01(**)$, and $p < .001(***)$, respectively, abbreviated by the number of stars in parenthesis.

6 RESULTS

We did not find significant effects of label placement techniques and data sizes on accuracy. User accuracy was similar and high across all conditions (see supplemental material). Therefore, we focus our analysis on time (see Fig. 6), subjective ratings, rankings (see Fig. 7), and qualitative feedback.

6.1 Completion Time

IDENTIFY: We found that label placement techniques had a significant effect on task completion time in IDENTIFY (***) . VALUE was slower than all other techniques (all ***). There was no difference between label techniques for in-view and out-of-view objects. We found no significant effect in data size and the interaction between the two factors (i.e., label placement technique \times data size) in IDENTIFY.

COMPARE: We also found that label placement techniques had a significant effect on task completion time in COMPARE (***) . Having labels for out-of-view objects (HEIGHT, ANGLE, and VALUE) was faster than techniques that only provided labels for in-view objects (SITUATED and BOUNDARY), all ***. For labels of out-of-view objects, ANGLE was faster than VALUE (*), and ANGLE tended to be faster than HEIGHT, but not statistically significant ($p = 0.060$). For labels of in-view objects, there was no difference between SITUATED and BOUNDARY. Data sizes had a significant effect on time in the COMPARE task (***) with large data sizes increasing completion time. No significant effect was found in the interaction between the two factors.

SUMMARIZE: Label placement techniques had a significant effect on task completion time in SUMMARIZE (***) . Like in IDENTIFY, VALUE was slower than other techniques (all ***). We also found HEIGHT was slower than SITUATED (***) and ANGLE (*). No significant effect was found in data size, and the interaction between the two factors was marginally significant ($p = 0.057$).

6.2 Subjective Ratings and Ranking

Our NASA-TLX results (Fig. 7 (a)) show that label placement techniques had a significant effect on the perceived *satisfaction* (**), *difficulty* (*), and *frustration* (*). **ANGLE received higher satisfaction ratings than SITUATED, BOUNDARY, and VALUE (**)**. SITUATED was rated more difficult than ANGLE (*). VALUE

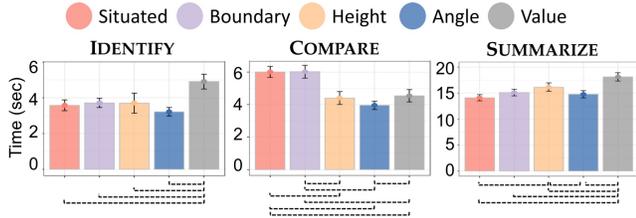


Figure 6: Results of task completion *time* by tasks. Confidence intervals indicate 95% confidence for mean values. Dashed lines in all charts indicate statistical significance for $p < .05$. All out-of-view labels outperformed in-view labels in the COMPARE task.

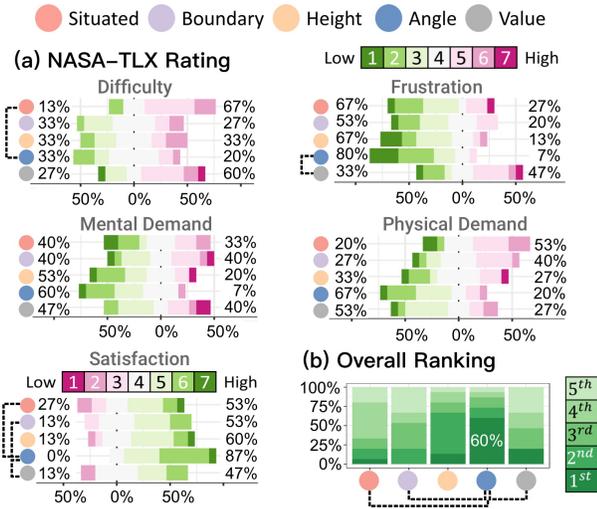


Figure 7: (a) NASA-TLX ratings. The percentage of negative (pink) and positive ratings (green) is shown next to the bars. (b) User preference ranking for each labeling condition. All dashed lines indicate $p < .05$.

was rated more frustrating than ANGLE (*). We also found that placement techniques had a significant effect on a user’s overall ranking (Fig. 7 (b)). ANGLE was more preferred than SITUATED, BOUNDARY, and VALUE (all *).

6.3 User Strategies for each Task

We observed substantial differences in user strategies for different label conditions apart from the task performance.

IDENTIFY. Participants used different strategies for in-view and out-of-view object labeling conditions. In SITUATED and BOUNDARY, all 15 participants chose an arbitrary direction to start a full-space search to find the target object. In HEIGHT, ANGLE, and VALUE, all participants scanned the AR screen for the target label before following the directional cues on the label to find the target.

COMPARE. In SITUATED and BOUNDARY, participants had to physically locate all target objects and memorize the label values when moving around. There was also a slight difference between them, where In SITUATED, participants had to constantly move their gaze up and down as they turned to find labels. In BOUNDARY, participants only moved their gaze horizontally to scan for labels and vertically to link to the object. With labels for out-of-view objects, users first compared labels and directly decided on the final target without having to physically turn around.

SUMMARIZE. In label conditions for in-view objects, all participants sequentially found clusters and calculated their group average one by one, which required them to remember the average of the previous cluster. Sometimes participants had to return to the

first cluster to double-check their answer. In out-of-view object labeling conditions, users directly looked at the labels to check the values of both clusters and estimate cluster averages. However, for the HEIGHT condition, four participants reported that finding labels on the screen took more work than physically turning back to the other cluster because label positions were less predictable and not grouped into clusters. They performed the task exactly the same as with SITUATED.

6.4 Overall Qualitative Feedback

Participants commented on the overall usefulness of each label condition for visual search in the post-study questionnaire.

Labels for In-View Objects. With SITUATED, 9 out of 15 participants commented that it was easy to link label to object. 3 of them also appreciated the lower visual density. With BOUNDARY, participants cited various advantages, including predictable label locations (4), easy matching from label to in-view objects (4), being good for comparison (3), label-to-label proximity (2), and being good for summarize tasks (2). **The major downside of labels for in-view objects was that they require a full spatial scan**, which was reported by 13 (SITUATED) and 7 (BOUNDARY) participants. Participants found that it was physically demanding (6 and 3), had a high load on memory (5 and 3), and was bad for compare tasks (4 and 2) and summarize tasks (3 and 2). They also mentioned that it was difficult to find objects for SITUATED (2) and leader line clutter for BOUNDARY (5).

Labels for Out-of-View Objects. Getting an immediate overview of objects and ratings was reported as an advantage in all three out-of-view object labeling conditions, with 5, 4, and 6 participants for HEIGHT, ANGLE, and VALUE, respectively. **Good directional cues was the major advantage for HEIGHT and ANGLE labels**, which was reported by 7 and 9 participants, respectively. In particular, 5 out of those 9 appreciated the precise directional cues provided by ANGLE labels. Participants also found it is easy to go from label to object in both conditions, 3 in HEIGHT and 6 in ANGLE. For the HEIGHT condition, participants reported predictable height locations of objects (2) and being good for comparisons (2). ANGLE was also preferred for its predictable label locations (1), similar to BOUNDARY. **VALUE was found most beneficial for comparison tasks.** Participants found VALUE good for comparisons (6), and good for compare tasks (3) and summarize tasks (6), and for providing good directional cues (2). **Leader line clutter was a common disadvantage reported for BOUNDARY, ANGLE, and VALUE.** The three conditions with longer leader line were reported to have leader line clutter, 5 in BOUNDARY and ANGLE, and 7 in VALUE. Participants also reported that it took time in the ANGLE condition to learn how to interpret the label placement (5). Unlike ANGLE, which showed the precise direction of objects, participants reported that **disadvantages for the HEIGHT and VALUE conditions were subtle directional cues** (2 and 4), **having to move slowly** (2 and 2), and **getting no distance information** (3 and 2). Furthermore, HEIGHT was reported to suffer from label clutter (3), unpredictable label locations (2), and being bad for summarize tasks (2). Another major disadvantage for VALUE was the difficulty of going from label to object (5).

6.5 Key Findings

General findings. Based on the quantitative and qualitative results, we summarize our key findings. **In-view labels** are good for

simple identify tasks; however, they do require a full spatial scan by the user. The main advantages of **out-of-view labels** are that they achieve high performance in visual comparison tasks, provide a fast overview of all objects, and can give immediate spatial cues to their linked object's location. However, when labeling many out-of-view objects, visual clutter can become a problem. We consider this an important topic for future work.

Task-specific findings. (1) **ANGLE was the overall winner.** ANGLE was one of the best in all three visual search tasks, and achieved the highest level of satisfaction compared to all other labels. (2) **Out-of-view labels were beneficial in COMPARE tasks.** All three out-of-view object labels performed better than labels for in-view objects in COMPARE tasks, for multiple spatially sparse objects. (3) **VALUE was not good for IDENTIFY and SUMMARIZE.** VALUE performed the worst in IDENTIFY and SUMMARIZE tasks, which was due to the difficulty to link leader lines to their objects. (4) **SITUATED and BOUNDARY had very similar performances.** In all task performance and subjective ratings, there is no significant difference between SITUATED and BOUNDARY. Finally, (5) **Data size only affected the performance in COMPARE.**

7 DISCUSSION

Our study aims to answer how to best place labels in AR to support situated visual search. Based on our study results with five representative label placing techniques, a short answer to the question is that placing the labels according to the **angular** direction of the object is overall the best option. Across all tasks, ANGLE performed as well as or better than other tested techniques. In addition to identifying the best performing technique, in this section, we discuss finer-grained findings.

7.1 AR Label Properties and Considerations

Visual search is a complicated process, which includes multiple components [68], [69], [70], [71]. To provide a more nuanced understanding of designing AR labels for situated visual search, we discuss the potential explanation and implications of our results with four different components in visual search as compiled by Yang et al [71]: (a) **Wayfinding** is the process of finding the destination. (b) **Travel** refers to the motor part of moving to the destination (i.e., walking, rotating). (c) **Context-switching** refers to the extra mental effort that is required to re-interpret a changed view. (d) **Number-of-travels** costs occur when completing a task may involve more than one travel, due to different capabilities of visual representations, or limited working memory.

Providing labels for out-of-view objects reduces the number of travels in comparison tasks. In COMPARE, all conditions with out-of-view object labels were faster than the conditions with only in-view object labels. We conjecture that a potential reason is that providing out-of-view object labels can reduce the number-of-travels when comparing multiple spatially sparse targets. The COMPARE task asked participants to search for the object with the highest value. With out-of-view object labels, all objects' values were accessible in the FoV, so users could first identify the target within the FoV and travel only once. In contrast, with labels only for in-view objects, participants needed to iterate through the candidates with multiple times of travel increasing the participants' memory load to keep track of all candidates' information. Our finding also aligns with the famous Shneiderman Information Seeking Mantra of "*Overview first, filtering and*

details on demand", which suggests the use of overviews to provide guided search [72]. Participants' comments support this hypothesis, as 13 participants complained that labels for only in-view objects required a full-spatial scan. Although the NASA-TLX physical demand rating was not statistically significant, we can see a trend of higher physical demand with labels for only in-view objects (Fig. 7 (a)). One may argue that automatic filtering can be used to reduce the number of candidates. However, in many cases, coming up with informed criteria for automatic filtering can be challenging due to a wide range of user goals [60]. Providing an overview of the entire information to guide the visual search is important for those scenarios.

Precise spatial cues reduce the costs in wayfinding and travel.

ANGLE was as good or better than other techniques across all tasks. In situated visual search tasks, the most fundamental interaction for users is to physically turn to face the target object. We believe that ANGLE provides the most precise spatial cue among the five tested conditions. It shows the angular direction of all objects, which allows users to estimate how much and in which direction they need to turn to face the target object. On the other hand, HEIGHT places labels on the left or right boundary of the FoV to indicate the turning direction, but users cannot estimate the turning angle. VALUE uses a more subtle cue to indicate the turning direction by placing an arrow on the labels. As a result, in HEIGHT and VALUE, users need to slowly turn towards the indicated direction and frequently check if the target object is already within their FoV. Label conditions for only in-view objects did not provide any spatial cues, and users had to fully scan the surrounding space. In summary, the ranking for spatial cue preciseness is: ANGLE > HEIGHT > VALUE > BOUNDARY = SITUATED. This ranking aligns with the good performance of ANGLE. However, HEIGHT and VALUE performed worse than BOUNDARY and SITUATED, which is not expected from the ranking. We believe their poor performance is mainly due to proximity factors, which we discuss later in this section.

Well-designed techniques reduce context-switching costs. To utilize the precise spatial cue in ANGLE, participants needed to mentally map the spatial information from the AR FoV to the surrounding physical environment. According to the VR study by Yang et al. [71], such operations are likely to introduce a high context-switching cost. However, their results were not reflected in our study, where we found ANGLE to be the best performing technique. One potential reason is that, in ANGLE, we used a top-down view to show the spatial distribution of objects with the user's AR FOV always facing up. We believe this design allows the user to easily interpret the "overview" provided by ANGLE, thus reduce the context-switching cost. Participants' ranking aligns with our finding, where ANGLE was ranked first by 60% of participants. Participants' comments also support this finding, as 9 participants found ANGLE easy to navigate.

Label-object proximity is important for spatial search. Despite providing a spatial cue, VALUE was the slowest condition for IDENTIFY and SUMMARIZE tasks. In VALUE, labels were placed in fixed positions on the left side of AR FoV. As a result, they were usually far away from the objects shown within the FoV, resulting in the lowest label-object proximity among all conditions. We believe that this far distance introduced difficulties for users to follow labels to objects. Furthermore, with increasing distance, the leader lines for connecting labels and objects become longer and can potentially introduce more line-crossings on the screen, making tracing leader lines even more difficult. This is also

reflected in the participants' comments, where 5 participants found it difficult to link between objects and labels, and 7 mentioned the line-crossings hindering them from completing the tasks.

Neighborhood consistency facilitates summarizing values by groups. In SUMMARIZE, participants needed to compare the average values to two clusters. Neighborhood consistency means that when objects are neighbors, their labels should also be neighbors to each other. This property can be shown as the lateral position encoding in Table. 1. SITUATED, BOUNDARY and ANGLE all preserve neighborhood consistency. HEIGHT only encodes lateral positions for in-view objects, and VALUE does not encode lateral positions at all. As objects in each target cluster are closer together, a condition with higher neighborhood consistency can potentially take less time for summarizing the values of that cluster. The collected performance data confirmed our analysis, where SITUATED, BOUNDARY, and ANGLE had a similar performance, and were faster than HEIGHT. VALUE had the worst performance.

Boundary labels did not demonstrate advantages over situated labels. We found a similar performance of SITUATED and BOUNDARY across all tasks. In 2D visualizations, boundary labeling is a standard way to reduce visual clutter in many scenarios, and was found to be beneficial [1]. However, we could not confirm the advantage of boundary labels in our AR study. One possible explanation is the limited FoV in AR, which makes these two conditions very similar to each other. We also did not formally control for the visual clutter or occlusions in our study. Future studies are needed to investigate the effectiveness of these two conditions for other scenarios.

7.2 Implications for Visual Search with AR Labels

Based on the above discussion, label designs for only in-view objects are suitable for applications where targets are located in the same direction, requiring less physical movement, such as browsing products on the same shelf. Furthermore, BOUNDARY is a good alternative for SITUATED when less visual clutter in the center of the FoV is desired, and tasks require more value comparison, such as comparing the prices of books on the shelf.

Labels for out-of-view objects are suitable for situated visual search when multiple targets of interest are scattered in space. In particular, ANGLE supports location navigation tasks well when precise directional cues are desired, such as searching for coffee shops in the area. As an alternative, HEIGHT provides better guidance when the height information is important for searching targets, such as books on a multiple-layer shelf, paintings hung at various heights in the gallery, and targets located at multiple floors. On the other hand, when knowing the precise locations is not as important as value comparison on the objects, VALUE is a good option for both in-view and out-of-view objects. Use cases include comparing data attributes of dynamically changing targets in real-time, such as players' scores in a live sports game. Furthermore, combining different label placement strategies in a single application can better support specific tasks. For example, after comparing menu items with VALUE labels, users might switch to ANGLE to locate the targets.

7.3 Limitations, Generalization and Future Work

While we have followed the guidance and practice from previous work [20], [21], [22], simulating AR interactions in VR might not be fully representative of a real-world use scenario.

Visually complex backgrounds. Real-world applications may have more visually complex backgrounds, which are likely to affect the ability to precisely perceive visual information (e.g., color [73]). To minimize the influence from the background, we suggest using visual cues that do not require precise perception in labels (e.g., we used the unit visualization method [74] in our study). We believe with effective encoding, people can still easily perceive simple visual information even with complex backgrounds. Thus, our results can largely transfer to these scenarios. However, there are cases where one would want to encode multivariate data in labels. Further studies are required to investigate the influence of complex backgrounds for these scenarios. Alternatively, we can show the most critical information in a simple and effective design, like ours, and provide on-demand interactions for users to selectively investigate objects of interest.

Real-world objects. Objects in the real world can be more complicated than our controlled environment. In our study, we did not control for objects' sizes and occlusion. Therefore, targets were relatively easy to find, which might have led to ceiling effects in accuracy. Real-world objects may vary largely in size and shape in some cases. Larger objects with unique shapes may easily attract people's attention, which possibly will influence the performance in identifying the target objects. We also did not explicitly control the occlusion between objects, as this would have resulted in too many testing factors. We have conducted a post hoc investigation of the trials with occluded targets and did not find a significant difference in performance. We conjectured that is because participants can easily move in space to find a viewpoint that avoids occlusion. However, in the case of AR labels, with explicit leader lines to direct the user's attention, such influence can be limited. Further studies are required to verify this expectation with varying difficulty in the real-world settings.

Dynamic objects and participants. Another important property for some real-world objects is that they are actively moving (e.g., players on a basketball court). Labeling dynamic objects requires further design considerations. For example, situated labels should move with the objects; therefore, we need to consider peoples' cognitive capacity to track dynamic visual elements [75], e.g., by stabilizing label movement.

Although we did not limit participants' movement, we observed limited body movements in the study, possibly due to the nature of our tested tasks. In real-world scenarios, we believe users may move more frequently in some scenarios, but they move mostly to get close to one specific target. As a result, the user does not need to analyze the entire view all the time. Substantially, our results can still largely apply to these scenarios.

Scalability. There can be cases that have a larger amount of objects to be labeled. We believe our study has tested representative amounts of objects (e.g., a basketball game has 10 players, and a soccer game has 22 players). Considering the limited size of current AR displays, it is challenging to display a larger amount of labels while still keeping the FoV uncluttered. One potential solution is to use focus+context technique to selectively enlarge more important labels, which can be determined by the user's position, facing direction, and object properties. Further study is needed to evaluate the effectiveness of such designs.

In addition to the aforementioned designs and studies, properties of AR labels other than label placement and interaction, such as a label's visual properties [8], leader line design [12], [11] and update frequency [23], are essential in applying AR labels to real-world applications but are not within the scope of this study.

8 CONCLUSION

We presented a thorough design space exploration of AR labels for out-of-view objects to support situated visual search tasks. Based on our classification of the design space, we compared five AR label conditions in a user study, focusing on the aspect of label placement. Our conditions included SITUATED and BOUNDARY for in-view objects, and HEIGHT, ANGLE, and VALUE for out-of-view objects. Our main results suggest that (1) Labels for out-of-view objects are beneficial for compare tasks. (2) Angle-encoded labels showed the best performance overall in supporting situated visual searching and were preferred by participants. (3) Value-encoded labels have strengths in compare tasks but are weak in identify and summarize tasks.

We hope that our classification and quantitative evaluation of AR out-of-view label design will help researchers to design effective labels and interactions for real-life tasks and applications. We envision future AR systems that integrate AR labels into their information retrieval systems, thereby enabling users to perform object-based search or to explore multiple levels of details in data in a truly situated setup.

ACKNOWLEDGMENTS

This research is supported in part by the National Science Foundation (NSF) under NSF Award Number IIS-1901030, and the Harvard Physical Sciences and Engineering Accelerator Award.

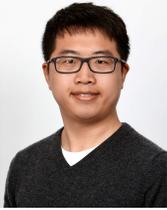
REFERENCES

- [1] M. A. Bekos, B. Niedermann, and M. Nöllenburg, "External Labeling Techniques: A Taxonomy and Survey," *Computer Graphics Forum*, vol. 38, no. 3, pp. 833–860, Jun. 2019.
- [2] K. Marriott, F. Schreiber, T. Dwyer, K. Klein, N. H. Riche, T. Itoh, W. Stuerzlinger, and B. H. Thomas, *Immersive Analytics*. Springer, 2018, vol. 11190.
- [3] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne, T. Dwyer, J. Grubert, J. H. Haga, N. Kirshenbaum, D. Kobayashi, T. Lin, M. Olaosebikan, F. Pointecker, D. Saffo, N. Saquib, D. Schmalstieg, D. A. Szafir, M. Whitlock, and Y. Yang, "Grand Challenges in Immersive Analytics," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 2021, pp. 1–17.
- [4] W. Willett, Y. Jansen, and P. Dragicevic, "Embedded Data Representations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 461–470, Jan. 2017.
- [5] T. Miyashita, P. Meier, T. Tachikawa, S. Orlic, T. Eble, V. Scholz, A. Gapel, O. Gerl, S. Arnaudov, and S. Lieberknecht, "An augmented reality museum guide," in *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*. IEEE, 2008, pp. 103–106.
- [6] T. Lin, Y. Yang, J. Beyer, and H. Pfister, "SportsXR – Immersive Analytics in Sports," in *4th Workshop on Immersive Analytics: Envisioning Future Productivity for Immersive Analytics at ACM CHI*, 2020.
- [7] CourtVision. (2020). [Online]. Available: <https://www.clippercourtvision.com/>
- [8] B. Bell, S. Feiner, and T. HSllerer, "View Management for Virtual and Augmented Reality," p. 10, 2001.
- [9] S. Maass and J. Döllner, "Dynamic Annotation of Interactive Environments using Object-Integrated Billboards," in *In 14th Int. Conf. in Central Europe on Computer Graphics, Visualization and Computer Vision (WSCG)*, 2006, pp. 327–334.
- [10] R. Azuma and C. Furmanski, "Evaluating label placement for augmented reality view management," in *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings.*, Oct. 2003, pp. 66–75.
- [11] M. Tatzgern, D. Kalkofen, R. Grasset, and D. Schmalstieg, "Hedgehog labeling: View management techniques for external labels in 3D space," in *2014 IEEE Virtual Reality (VR)*, Mar. 2014, pp. 27–32.
- [12] R. Grasset, T. Langlotz, D. Kalkofen, M. Tatzgern, and D. Schmalstieg, "Image-driven view management for augmented reality browsers," in *2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Nov. 2012, pp. 177–186.
- [13] M. Tatzgern, V. Orso, D. Kalkofen, G. Jacucci, L. Gamberini, and D. Schmalstieg, "Adaptive information density for augmented reality displays," in *2016 IEEE Virtual Reality (VR)*. Greenville, SC, USA: IEEE, Mar. 2016, pp. 83–92.
- [14] T. Schinke, N. Henze, and S. Boll, "Visualization of off-screen objects in mobile augmented reality," in *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, 2010, pp. 313–316.
- [15] H. Jo, S. Hwang, H. Park, and J.-h. Ryu, "Aroundplot: Focus+ context interface for off-screen objects in 3d environments," *Computers & Graphics*, vol. 35, no. 4, pp. 841–853, 2011.
- [16] T. Siu and V. Herskovic, "SideARs: Improving awareness of off-screen elements in mobile augmented reality," in *Proceedings of the 2013 Chilean Conference on Human - Computer Interaction - ChileCHI '13*. Temuco, Chile: ACM Press, 2013, pp. 36–41.
- [17] U. Gruenefeld, D. Ennenga, A. E. Ali, W. Heuten, and S. Boll, "Eye-See360: Designing a visualization technique for out-of-view objects in head-mounted augmented reality," in *Proceedings of the 5th Symposium on Spatial User Interaction*. ACM, Oct. 2017, pp. 109–118.
- [18] F. Bork, C. Schnellzer, U. Eck, and N. Navab, "Towards Efficient Visual Guidance in Limited Field-of-View Head-Mounted Displays," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 11, pp. 2983–2992, Nov. 2018.
- [19] K. A. Satriadi, B. Ens, M. Cordeil, B. Jenny, T. Czauderna, and W. Willett, "Augmented Reality Map Navigation with Freehand Gestures," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, Mar. 2019, pp. 593–603.
- [20] A. Marquardt, C. Trepkowski, T. D. Eibich, J. Maiero, E. Kruijff, and J. Schöning, "Comparing Non-Visual and Visual Guidance Methods for Narrow Field of View Augmented Reality Displays," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 12, pp. 3389–3401, Dec. 2020.
- [21] E. Ragan, C. Wilkes, D. Bowman, and T. Hollerer, "Simulation of Augmented Reality Systems in Purely Virtual Environments," in *2009 IEEE Virtual Reality Conference*. Lafayette, LA: IEEE, Mar. 2009, pp. 287–288.
- [22] J. Jung, H. Lee, J. Choi, A. Nanda, U. Gruenefeld, T. Stratmann, and W. Heuten, "Ensuring Safety in Augmented Reality from Trade-off Between Immersion and Situation Awareness," in *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Oct. 2018, pp. 70–79.
- [23] J. B. Madsen, M. Tatzgern, C. B. Madsen, D. Schmalstieg, and D. Kalkofen, "Temporal Coherence Strategies for Augmented Reality Labeling," *IEEE Transactions on Visualization and Computer Graphics*, vol. 22, no. 4, pp. 1415–1423, Apr. 2016.
- [24] S. Feiner, B. Macintyre, and D. Seligmann, "Knowledge-based augmented reality," *Communications of the ACM*, vol. 36, no. 7, pp. 53–62, Jul. 1993.
- [25] S. Julier, Y. Baillot, D. Brown, and M. Lanzagorta, "Information filtering for mobile augmented reality," *IEEE Computer Graphics and Applications*, vol. 22, no. 5, pp. 12–15, Sep. 2002.
- [26] J. M. Wolfe, M. L.-H. Vö, K. K. Evans, and M. R. Greene, "Visual search in scenes involves selective and nonselective pathways," *Trends in Cognitive Sciences*, vol. 15, no. 2, pp. 77–84, Feb. 2011.
- [27] F. Biocca, A. Tang, C. Owen, and F. Xiao, "Attention funnel: omnidirectional 3d cursor for mobile augmented reality platforms," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 2006, pp. 1115–1122.
- [28] N. ElSayed, B. Thomas, K. Marriott, J. Piantadosi, and R. Smith, "Situated analytics," in *2015 Big Data Visual Analytics (BDVA)*. IEEE, 2015, pp. 1–8.
- [29] W. Büschel, A. Mitschick, and R. Dachselt, "Here and now: reality-based information retrieval: perspective paper," in *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*, 2018, pp. 171–180.
- [30] B. Bach, R. Sicat, H. Pfister, and A. Quigley, "Drawing into the ar-canvas: Designing embedded visualizations for augmented reality," in *Workshop on Immersive Analytics, IEEE Vis*, 2017.
- [31] E. Horvitz, C. Kadie, T. Paek, and D. Hovel, "Models of attention in computing and communication: From principles to applications," *Communications of the ACM*, vol. 46, no. 3, pp. 52–59, Mar. 2003.
- [32] R. Vertegaal, "Designing attentive interfaces," in *Proceedings of the 2002 symposium on Eye tracking research & applications*, 2002, pp. 23–30.
- [33] D. A. Redelmeier and R. J. Tibshirani, "Association between Cellular-Telephone Calls and Motor Vehicle Collisions," *New England Journal of Medicine*, vol. 336, no. 7, pp. 453–458, Feb. 1997.

- [34] R. M. Pascoal and S. L. Guerreiro, "Information overload in augmented reality: The outdoor sports environments," in *Information and Communication Overload in the Digital Age*. IGI Global, 2017, pp. 271–301.
- [35] S. Maass and J. Döllner, "Efficient view management for dynamic annotation placement in virtual landscapes," in *International Symposium on Smart Graphics*. Springer, 2006, pp. 1–12.
- [36] K. Makita, M. Kanbara, and N. Yokoya, "View management of annotations for wearable augmented reality," in *2009 IEEE International Conference on Multimedia and Expo*, Jun. 2009, pp. 982–985.
- [37] J. Orlosky, K. Kiyokawa, T. Toyama, and D. Sonntag, "Halo Content: Context-aware Viewspace Management for Non-invasive Augmented Reality," in *Proceedings of the 20th International Conference on Intelligent User Interfaces - IUI '15*. ACM Press, 2015, pp. 369–373.
- [38] A. Cockburn, A. Karlson, and B. B. Bederson, "A review of overview+detail, zooming, and focus+context interfaces," *ACM Computing Surveys (CSUR)*, vol. 41, no. 1, pp. 1–31, 2009.
- [39] M. Sarkar and M. H. Brown, "Graphical fisheye views of graphs," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '92*. ACM Press, 1992, pp. 83–91.
- [40] A. Zanella, M. S. T. Carpendale, and M. Rounding, "On the effects of viewing cues in comprehending distortions," in *Proceedings of the second Nordic conference on Human-computer interaction*, 2002, pp. 119–128.
- [41] P. T. Zellweger, J. D. Mackinlay, L. Good, M. Stefik, and P. Baudisch, "City lights: contextual views in minimal space," in *CHI'03 extended abstracts on Human factors in computing systems*, 2003, pp. 838–839.
- [42] P. Baudisch and R. Rosenbaltz, "Halo: a technique for visualizing off-screen locations," in *CHI 2003*, in *ACM Conference on Human Factors in Computing Systems, CHI Letters*, vol. 5, no. 1, 2003, pp. 481–488.
- [43] S. Gustafson, P. Baudisch, C. Gutwin, and P. Irani, "Wedge: clutter-free visualization of off-screen locations," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2008, pp. 787–796.
- [44] S. Burigat, L. Chittaro, and S. Gabrielli, "Visualizing locations of off-screen objects on mobile devices: a comparative evaluation of three approaches," in *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, 2006, pp. 239–246.
- [45] S. G. Gustafson and P. P. Irani, "Comparing visualizations for tracking off-screen moving targets," in *CHI'07 Extended Abstracts on Human Factors in Computing Systems*, 2007, pp. 2399–2404.
- [46] M. S. T. Carpendale and C. Montagnese, "A framework for unifying presentation space," in *Proceedings of the 14th annual ACM symposium on User interface software and technology*, 2001, pp. 61–70.
- [47] S. Ghani, N. H. Riche, and N. Elmqvist, "Dynamic insets for context-aware graph navigation," in *Computer Graphics Forum*, vol. 30, no. 3. Wiley Online Library, 2011, pp. 861–870.
- [48] F. Lekschas, M. Behrisch, B. Bach, P. Kerpedjiev, N. Gehlenborg, and H. Pfister, "Pattern-driven navigation in 2d multiscale visualizations with scalable insets," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 1, pp. 611–621, 2019.
- [49] L. Chittaro and S. Burigat, "3d location-pointing as a navigation aid in virtual environments," in *Proceedings of the working conference on Advanced visual interfaces*, 2004, pp. 267–274.
- [50] T. Schinke, N. Henze, and S. Boll, "Visualization of off-screen objects in mobile augmented reality," in *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, 2010, pp. 313–316.
- [51] M. Tonniss and G. Klinker, "Effective control of a car driver's attention for visual and acoustic guidance towards the direction of imminent dangers," in *2006 IEEE/ACM International Symposium on Mixed and Augmented Reality*. IEEE, 2006, pp. 13–22.
- [52] S. Burigat and L. Chittaro, "Navigation in 3d virtual environments: Effects of user experience and location-pointing navigation aids," *International Journal of Human-Computer Studies*, vol. 65, no. 11, pp. 945–958, 2007.
- [53] M. Trapp, L. Schneider, C. Lehmann, N. Holz, and J. Döllner, "Strategies for visualising 3d points-of-interest on mobile devices," *Journal of Location Based Services*, vol. 5, no. 2, pp. 79–99, 2011.
- [54] J. Petford, I. Carson, M. A. Nacenta, and C. Gutwin, "A comparison of notification techniques for out-of-view objects in full-coverage displays," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–13.
- [55] R. Stoakley, M. J. Conway, and R. Pausch, "Virtual reality on a wim: interactive worlds in miniature," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 1995, pp. 265–272.
- [56] U. Gruenefeld, A. E. Ali, S. Boll, and W. Heuten, "Beyond halo and wedge: Visualizing out-of-view objects on head-mounted virtual and augmented reality devices," in *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services*, 2018, pp. 1–11.
- [57] U. Gruenefeld, I. Koethe, D. Lange, S. Weiß, and W. Heuten, "Comparing Techniques for Visualizing Moving Out-of-View Objects in Head-mounted Virtual Reality," in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, Mar. 2019, pp. 742–746.
- [58] T. Stein and X. Décoret, "Dynamic label placement for improved interactive exploration," in *Proceedings of the 6th international symposium on Non-photorealistic animation and rendering*, 2008, pp. 15–21.
- [59] E. M. Coelho, B. Macintyre, and S. J. Julier, "Osgar: A scene graph with uncertain transformations," in *In ISMAR'04: Proceedings of the IEEE/ACM International Symposium on Mixed and Augmented Reality*, 2004, pp. 6–15.
- [60] M. Brehmer and T. Munzner, "A Multi-Level Typology of Abstract Visualization Tasks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2376–2385, Dec. 2013.
- [61] B. Yoo, J.-J. Han, C. Choi, K. Yi, S. Suh, D. Park, and C. Kim, "3d user interface combining gaze and hand gestures for large-scale display," in *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, 2010, pp. 3709–3714.
- [62] Unity Technologies. (2021) Unity real-time development platform. [Online]. Available: <https://unity.com/>
- [63] Microsoft. (2021) Mixed reality toolkit. [Online]. Available: <https://docs.microsoft.com/en-us/windows/mixed-reality/mrtk-unity/>
- [64] Microsoft. (2021) Tooltip - mixed reality toolkit. [Online]. Available: <https://docs.microsoft.com/en-us/windows/mixed-reality/mrtk-unity/features/ux-building-blocks/tooltip?view=mrtkunity-2021-05>
- [65] D. Bates, M. Mächler, B. Bolker, and S. Walker, "Fitting linear mixed-effects models using lme4," *Journal of Statistical Software*, vol. 67, no. 1, 2015.
- [66] A. Field, J. Miles, and Z. Field, *Discovering statistics using R*. Sage publications, 2012.
- [67] R. V. Lenth, "Least-squares means: The R Package lsmeans," *Journal of Statistical Software*, vol. 69, no. 1, 2016.
- [68] J. J. LaViola Jr, E. Kruijff, R. P. McMahan, D. Bowman, and I. P. Poupyrev, *3D user interfaces: theory and practice*. Addison-Wesley Professional, 2017.
- [69] N. C. Nilsson, S. Serafin, F. Steinicke, and R. Nordahl, "Natural walking in virtual reality: A review," *Computers in Entertainment (CIE)*, vol. 16, no. 2, pp. 1–22, 2018.
- [70] H. Lam, "A framework of interaction costs in information visualization," *IEEE transactions on visualization and computer graphics*, vol. 14, no. 6, pp. 1149–1156, 2008.
- [71] Y. Yang, M. Cordeil, J. Beyer, T. Dwyer, K. Marriott, and H. Pfister, "Embodied Navigation in Immersive Abstract Data Visualization: Is Overview+Detail or Zooming Better for 3D Scatterplots?" *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1214–1224, 2021.
- [72] B. Shneiderman, "The eyes have it: A task by data type taxonomy for information visualizations," in *The craft of information visualization*. Elsevier, 2003, pp. 364–371.
- [73] M. Whitlock, S. Smart, and D. A. Szafrir, "Graphical perception for immersive analytics," in *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2020, pp. 616–625.
- [74] D. Park, S. M. Drucker, R. Fernandez, and N. Elmqvist, "Atom: A grammar for unit visualizations," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 12, pp. 3032–3043, 2017.
- [75] J. Heer and G. Robertson, "Animated transitions in statistical data graphics," *IEEE transactions on visualization and computer graphics*, vol. 13, no. 6, pp. 1240–1247, 2007.



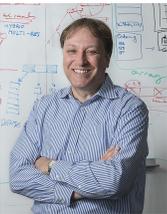
Tica Lin is a Ph.D. student at the Visual Computing Group at Harvard University. Prior to joining Harvard, she was a Data Visualization Designer at Visa and a UX Developer at NBA 76ers. Her research interests include data visualization, immersive analytics, and human-computer interaction. In particular, she explores novel visualization and interaction design in Augmented Reality.



Yalong Yang is an Assistant Professor at Virginia Tech. He was a Postdoctoral Fellow at the Visual Computing Group at Harvard University, and a Ph.D. student at Monash University, Australia. His research designs and evaluates interactive visualisations on both conventional 2D screens and in 3D immersive environments (VR/AR). He received best paper honorable mentions at VIS 2016 and CHI 2021.



Johanna Beyer is a research associate at the Visual Computing Group at Harvard University. Before joining Harvard, she was a postdoctoral fellow at the Visual Computing Center at KAUST. She received her Ph.D. in computer science at the University of Technology Vienna, Austria in 2009. Her research interests include scalable methods for visual abstractions, large-scale volume visualization, and immersive analytics.



Hanspeter Pfister is An Wang Professor of Computer Science in the John A. Paulson School of Engineering and Applied Sciences at Harvard University. His research in visual computing lies at the intersection of scientific visualization, information visualization, computer graphics, and computer vision and spans a wide range of topics, including biomedical image analysis and visualization, image and video analysis, and visual analytics in data science.