

Visual Highlighting for Situated Brushing and Linking

Nina Doerr^{id}, Benjamin Lee^{id}, Katarina Baricova^{id}, Dieter Schmalstieg^{id}, Michael Sedlmair^{id}

University of Stuttgart, Germany



Figure 1: In our user study, we present a virtual supermarket with four highlighting techniques. Center: First-person view of the supermarket while brushing a scatterplot, with products on the shelves highlighted. Top left: The *Color* technique. Bottom left: The *Outline* technique. Top right: The *Link* technique with products behind the camera selected. Bottom right: Still image of the *Arrow* technique.

Abstract

Brushing and linking is widely used for visual analytics in desktop environments. However, using this approach to link many data items between situated (e.g., a virtual screen with data) and embedded views (e.g., highlighted objects in the physical environment) is largely unexplored. To this end, we study the effectiveness of visual highlighting techniques in helping users identify and link physical referents to brushed data marks in a situated scatterplot. In an exploratory virtual reality user study (N=20), we evaluated four highlighting techniques under different physical layouts and tasks. We discuss the effectiveness of these techniques, as well as implications for the design of brushing and linking operations in situated analytics.

CCS Concepts

• *Human-centered computing* → *Empirical studies in visualization; Empirical studies in HCI; Information visualization;*

1. Introduction

Situated analytics (SitA) involves the use of situated visualizations [WF09] to improve sense-making in physical contexts [ETM*15, TWD*18]. It is facilitated predominantly by the use of augmented reality (AR) [BKT*22]. Willett et al. [WJD17] distinguish visualizations directly embedded with one particular physical referent from situated visualizations that indirectly link to a referent or the entire environment. Shin et al. [SBB*23] note that most current SitA tools use the referent as a trigger to instantiate and modify visualizations. For instance, *scanner* applications augment real-world objects with additional information as they come into view. SitA is therefore reliant on the user’s viewpoint and interactions with the real world to determine what information to display.

Although this “physical world first” approach feels natural in the context of AR, it is inherently constrained in situations with many physical referents or referents that are spread far apart [LSS23]. Consider the classic supermarket scenario introduced by ElSayed et al. [ETM*16], which features each product in the

store having its own embedded visualization. While this scenario works well for products in view, it does not scale well to products *outside* of one’s view. Even with out-of-view labeling techniques [GLH*18, LYBP23], the customer in one-to-one configuration would still need to look at and interpret the visualization of each individual product [LSS23]. The issue is exacerbated by the small field of view (FOV) of current state-of-the-art AR headsets.

Situated visualizations that are presented independently of individual referents can naturally consolidate the data pertaining to many referents in one view or dashboard [MM21, SLQS23]. Instead of looking at multiple embedded visualizations, a single situated overview visualization can be looked at instead. However, should the task require the identification or physical manipulation of referents (e.g., when shopping), then supporting the *transition* between a situated visualization of all referents to an individualized, embedded view of specific referents is required [WJD17].

We postulate that this problem is functionally equivalent to *brushing and linking*, a fundamental technique in visual analyt-

ics [Rob07]. Selections made in the situated visualization must be visually linked to the corresponding physical referents. Linking therefore serves both analytic (making selections based on certain criteria) and perceptual purposes (finding and locating objects).

In this paper, we explore multiple techniques to visually highlight physical referents in the context of brushing and linking. We focus on the *linking* between a situated visualization and its physical referent, rather than the interactions used to perform the *brushing*, as it is the visual linking that fundamentally facilitates the transition between views. While existing work has investigated many ways to highlight and guide attention in 3D environments (see Section 2), to the best of our knowledge, none have explored highlighting for brushing and linking in SitA. We selected four highlighting techniques known from either traditional brushing and linking or attention guidance in AR/VR. We conducted an exploratory user study (N=20), using a supermarket scenario. The study was carried out in virtual reality (VR) to avoid hardware and tracking limitations of current AR devices. In summary, we contribute:

- An open-source VR prototype demonstrating the use of brushing and linking in a situated analytics scenario
- An exploratory user study that seeks to compare, evaluate, and understand four common highlighting techniques
- Observations and lessons learned for the application of highlighting techniques in situated brushing and linking contexts

2. Related work

First, we consider *situated analytics* (Section 2.1) and, second, *brushing and linking* (Section 2.2) and its use in immersive and situated environments. As we focus on the linking component, we also discuss *attention guidance in AR and VR* (Section 2.3).

2.1. Immersive and situated analytics

Immersive technologies, both VR and AR, are increasingly being used for immersive analytics [MSD*18] of spatial data [FP21] and abstract data [KFS*22]. Research has shown how the 3D environment can be used as an “immersive space to think” [LDG*21], populating this space with multiple views and representations [RBR22]. However, care must be taken to consider the pitfalls of designing for immersive displays [ML14]. Designers of VR interfaces for analytics must take into account the characteristics of VR displays that differ from desktop computing, such as stereoscopic depth, angular resolution, field of view, or fatigue related to physical navigation and interaction [LKM*17].

Recent work considers immersive analytics techniques with AR in physical environments, a new direction referred to as situated analytics (SitA) [ETM*15, ETM*16]. SitA is characterized by the ability to perform higher-level analytical reasoning aided by the user’s physical environment [SBB*23], particularly when there is a semantic relationship between the data and its physical referents [WF09, WJD17]. While research has investigated how situated visualizations can be designed for SitA [LSS23, SBB*23], considering human perception to leverage efficient visual processing—particularly in situated contexts—is a grand challenge [EBC*21].

Designers of SitA systems (and AR interfaces in general) must handle the additional constraint that visual representations cannot always be positioned on a “clean slate” background. Kruijff

et al. [KSF10] list several challenges related to visual perception in AR which are caused by the environment, image capture, display, and visual representation. Satkowski et al. [SD21] study how the real-world background affects perception in AR, and Assor et al. [APHD24] propose a design space of how to handle non-visible physical referents. Similarly, AR must always consider the interplay of visual representation with the existing and immutable appearance of the real world. Therefore, any visualization pipeline targeting AR needs to deal with the fusion of real and virtual visual attributes [ZLG*21]. Addressing these challenges requires a wide range of visual encoding methods [SH16, chapter 7].

2.2. Brushing and linking

One of the core topics of this paper is brushing and linking [Rob07], which is the process of locating the corresponding items of interest in multiple views. Fundamentally, it comprises two sequential stages, wherein selections made in one visualization (brushing) are automatically shown in another visualization (linking) [Kei02]. Much research has focused on the brushing component, investigating interaction techniques to support the effective selection of data points in desktop computing [KPV*18].

In immersive analytics, brushing and linking have seen (arguably limited) use, be it to compare multiple immersive coordinated views [AWG*15, LPED20] or to support collaborative awareness [LHC*21, SBDE23, SS22]. Most methods use simple 3D selection techniques (e.g., raycasting) or pure 2D selection techniques (e.g., touchscreen inputs). While we do not propose new brushing interaction techniques, we are the first (to the best of our knowledge) to explore brushing and linking in a SitA context—particularly when using physical referents as a coordinated view.

2.3. Attention guidance in augmented and virtual reality

In SitA, we assume that the linking component targets views and referents in the real world. Hence, it can be described as a visual search task for target objects in an AR display. Note that this work considers attention guidance as a means to support visual search and identification, and not as instructional cues for movement training [YLS24] or assembly tasks [PZB*23]. Wolfe et al. [WVEG11] characterize visual search as the task of visually identifying a target among other, irrelevant objects (distractors). If there are too many distractors, a naive visual search becomes inefficient [Ver02]. We must therefore ensure that the visual search is tolerant of clutter and occlusion, since we cannot change the real world. There should be sufficient contrast and legibility between real and virtual objects. We have to be able to direct a user’s attention to the relevant objects, in particular, if these are out of view.

According to Cockburn et al. [CKB09], conventional InfoVis methods to guide visual attention and search, such as overview + detail or focus + context, occupy a lot of screen space as they enlarge objects or introduce replicas of different sizes. As a major alternative, they mention cue-based approaches, which modify object rendering styles. Lin et al. [LYBP23] note that such cues are preferred in AR, since resizing or duplicating occupies too much of the natural view space. These cue-based techniques for highlighting focal objects can select from a variety of rendering attributes. A baseline is formed by highlights, which emphasize a target by modifying its color throughout the target area or as an outline [KPV*18].

In immersive displays, several more sophisticated techniques have been proposed to ensure that a given contrast is achieved [GAM18, KVZ*13]. Contrast modulation, even below a detectable threshold [VMFS11], can pre-attentively guide the user. Other attributes that can be modified to highlight or attract attention include spatial frequencies (i.e., focus or blur) [KMH01], stereoscopy [KCWK20], or temporal frequencies [LSGB20].

Cockburn et al. [CKB09] further mention proxies as cues. Arrow-shaped proxies are commonly used in AR [GLH*18, TK06], most likely because they are simple to synthesize and easy to understand. Naturally, proxies are a popular means to guide toward out-of-view targets. In addition to arrows, various popular shapes include halos [BR03], funnels [BTOX06], compass-like [SL00] or radar-like widgets [JHPR11], and lines or trajectories [PLE*19]. Previous work also compares the effect of various visual cues in VR and AR. Several authors [BSEN18, KKO*14, TEM*19, WGRF22] investigate the performance differences of visual search with respect to the field of view of the display. Other studies compare multiple attention guidance (i.e., highlighting) techniques with respect to the effectiveness and preservation of immersion [LSGB20], the effectiveness of multimodal cues [MTE*20], or the influence of stationary and moving distractors [DARS23]. Moreover, our selection of techniques was influenced by the work of Whitlock et al. [WSS20] on the perception of fundamental graphical attributes in immersive analytics. However, none of these works investigate the use of AR or VR attention guidance in the context of brushing and linking. Therefore, our work seeks to bridge the gap between these two topics necessary to enable brushing and linking in SitA.

3. Experimental prototype for situated brushing and linking

This research aims to investigate how visual highlighting could support brushing and linking in SitA. As mentioned in Section 2, brushing and linking, as well as visual highlighting (i.e., attention guidance) are not novel ideas when considered individually, but it is the intersection in SitA that requires further study. Fortunately, existing techniques can be adapted to SitA. To ensure a more grounded and systematic approach, we draw upon popular techniques already used in brushing and linking in visual analytics, and attention guidance in AR and VR. This approach has dual benefits. First, we can immediately see if prior research directly translates into SitA. Second, we can establish a baseline understanding of brushing and linking in SitA without the confounding effects of potentially complex highlighting techniques. In this section, we describe a VR prototype that we developed to investigate situated brushing and linking, which is open source on GitHub [DL24].

3.1. Scenario

SitA is relevant in many situations involving data and its physical context [BKT*22]. Brushing and linking becomes useful when considering the actual appearance or location of physical referents. We therefore chose a supermarket scenario to situate our experimental prototype and the subsequent user study (Figure 1). A supermarket is a popular scenario in the SitA literature [AWG*15, ETM*15, EST16]. Its defining feature is an abundance of grocery products (referents) laid out on shelves, which vary in layout, making the scenario well-suited for brushing and linking. We also considered similar environments, such as a library with books or a

building façade with windows as referents [TOK*16]. We built the supermarket using VR as a kind of information-rich virtual environment [BNC*03, PKB05], serving as a simulated test environment that can be carefully controlled and replicated [RWBH09]. While using AR would be more ecologically valid, the use of VR avoids the technical challenges of tracking each product [CFSS23] or logistical challenges of renting or building a store. Previous work also suggested that findings in AR can be replicated in VR [GCOK21, LBHB09, LBBH10, LRM*13], which further motivates and validates our use of a simulated supermarket in VR.

3.2. Highlighting techniques

We implemented four different highlighting techniques in our prototype. These techniques are used in conjunction with the accompanying *brushing* component, which is described in Section 3.3. Example images of all four techniques can be seen in Figure 1. Note that all techniques use a solid yellow color where relevant in order to ensure a consistent appearance across conditions.

Color is likely the de facto highlighting technique used in regular brushing and linking [KPV*18, Rob07]. It mainly benefits from the pre-attentive “popout” effect, particularly when the new color highly contrasts with the background. Hence, we incorporate it into our prototype, changing the color of the product to a solid yellow when selected. While we had initially tried to use a semi-transparent yellow tint on the products instead, pilot tests showed that this mode was not able to sufficiently “pop-out” specific products: A red cereal box, when tinted yellow, simply looks like an orange cereal box, and not a red one that is highlighted.

Outline is another common technique found in brushing and linking [GSL*14, GGL*14]. Moreover, it is also a popular technique used in visual highlighting to show that an object is “selected” [DMT*18, SCZ*20]. Although *Color* and *Outline* have similar characteristics in how they highlight targets, the actual appearance of a product is visible for *Outline* but not for *Color*. Sometimes a more subtle, non-obtrusive highlighting can be beneficial (e.g., for safety). We show outlines on the borders of the selected grocery products in camera space. We chose a width for the outline which, after pilot testing, we felt was sufficiently wide to see the outline, but thin enough not to obstruct neighboring products.

Links are trajectories used to draw connections between entities [CC07]. In the context of brushing and linking, a link indicates that two connected data marks are the same underlying data record [KPV*18]. The same premise holds in SitA, where a data mark on the situated visualization (i.e., the virtual tablet in Figure 1) is explicitly connected to its associated referent. Particularly for the comparison and analysis across view, direct connections can facilitate tasks. To this end, we leverage a prior Unity implementation by Prouzeau et al. [PLE*19] which draws visual links between pairs of objects in a 3D scene. We adjusted some of their script’s parameters to make the links more responsive to movement, because we allow the situated visualization to be moved by the user.

Arrows are ubiquitous symbols used to direct attention. They are popular for attention guidance in AR, particularly towards objects or points of interest that are out of view [KR14, TK06, WGRF22, YLF*20]. Thus, we consider how similar arrows could be used to

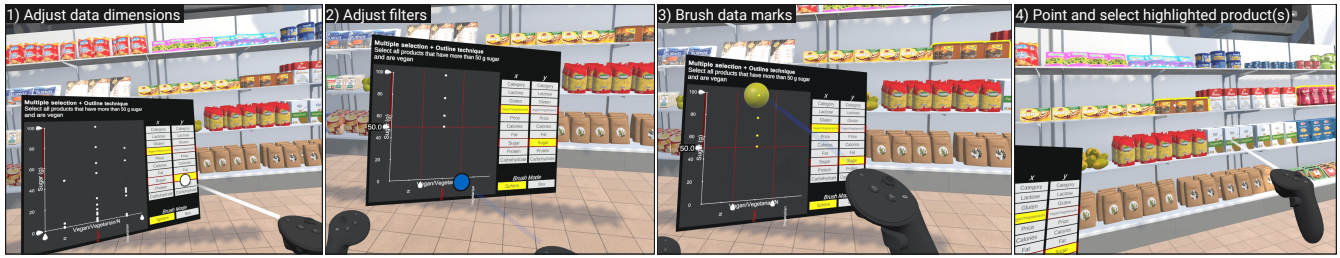


Figure 2: Sequential frames of brushing on a situated scatterplot and identifying linked products (referents) during our user study, with the virtual tablet being held at head height for illustrative purposes. Frame 1: The user adjusts the data dimensions as desired. Frame 2: The user adjusts the filters to show only the data marks within the desired ranges. Frame 3: The user selects the data marks. Frame 4: The user finds the highlighted product(s) and confirms identification by pointing and selecting with controller.

highlight objects in SitA. For our study, we extended *FlyingARrow* [GLH*18], which repeatedly instantiates 3D arrows in front of the user that then fly toward a target object. Static arrows like Wieland et al. [WGRF22] or Yu et al. [YLF*20] spawn in the viewport, causing clutter when there are multiple targets and hence might be less beneficial in real-life scenarios than dynamic arrows [GLH*18]. We also suspect that adding animations or motion makes targets more salient. Since Gruenefeld et al. [GLH*18] claim a low workload, we suspected the dynamic arrow to be a reasonable choice for highlighting in SitA. Our adapted version of *FlyingARrow* is able to present multiple arrows concurrently and dynamically change the arrows to support on-the-fly brushing.

Alternatives beyond these four fundamental highlighting techniques were considered but excluded after preliminary testing. *Size* was an obvious candidate [GCC17], but it became apparent that expanding objects in the real world is impractical because they take up too much of the physical space [LYBP23]. *Semantic depth of field* [KMH01] was another candidate, as blur can intentionally obscure unimportant parts of the scene. However, we soon realized that blur is not only visually distracting, but also potentially dangerous when certain parts of the real world are visually deteriorated.

3.3. Brushing interactions

As the premise of our work involves both *brushing* and *linking*, our prototype needs faculties to enable brushing in a situated visualization. A virtual tablet attached to the non-dominant-hand controller (Figure 1, center) presents a 2D scatterplot, generated with DXR [SLC*19]. The x and y data dimensions of the scatterplot are determined by buttons on the right side of the tablet. The tablet's size (40 cm by 31 cm) ensures that its contents are easily readable in VR, while still small enough to not be visually obtrusive.

We opt for a simple brushing approach, reminiscent of the brushing mechanism used in FIESTA [LHC*21]. Brushing is performed by pointing the dominant-hand controller at the scatterplot and pressing either the index button to add to the selection or the grip button to remove it from the selection. We provide two modes for brushing: spherical and rectangular. The mode can be changed using the corresponding buttons on the tablet. Any brushed data marks on the scatterplot turn a solid yellow. To make selections even easier, we added simple range filters along the two data axes. These filters are represented as handles which can be dragged using raycast interactions, similar to ImAxes [CCD*17]. Filtered data

marks cannot be brushed, ensuring that no unwanted points are included. Step-by-step images of this situated brushing and linking process are shown in Figure 2. The tablet could be held comfortably at hip height without obscuring any referents, as was the case for everyone during our pilot testing and subsequent user study. As our intention is to evaluate the effectiveness of the four techniques to highlight referents, we focus on the unidirectional brushing from the situated visualization to the referents.

4. User study

We conducted a user study with our VR prototype to determine how the highlighting techniques presented in Section 3.2 perform under different referent layouts as well as situated brushing and linking tasks. The overall goal of the study was to determine: (1) how well existing techniques from brushing and linking and attention guidance hold up in SitA; (2) how well each highlighting technique supports our user experience categories (Section 4.4); and (3) how brushing and linking in SitA can potentially be improved.

4.1. Study conditions

Our user study involved two independent variables:

Highlighting technique. *Color*, *Outline*, *Link*, and *Arrow*, as described in Section 3.2.

Shelf layout. *Inside-FOV* or *Outside-FOV*. For *Inside-FOV*, there are only four shelves in front of the user. *Outside-FOV* extends the layout by adding four shelves behind the user, with front and back rows spaced 4.8 m apart. This extended layout emulates situations where the highlighted referents may be positioned outside of the user's FOV. Each shelf is 2 m wide by 2 m tall, forming an overall aisle with a width of 8 m in total (Figure 3 left). Neither layout forces the user to move to resolve occlusion. This choice eliminates the need to support any special VR locomotion techniques, which would probably have confounded our results.

4.2. Tasks and data

Our study involved three main task types: single-selection, multi-selection, and statement response. All tasks required participants to first make a selection on the scatterplot and then look for one or more highlighted products on the supermarket shelves. All tasks are bivariate, with the scatterplot selection being conditioned in both dimensions. The tasks are intentionally kept low-level to allow for generalizability across various SitA applications involving

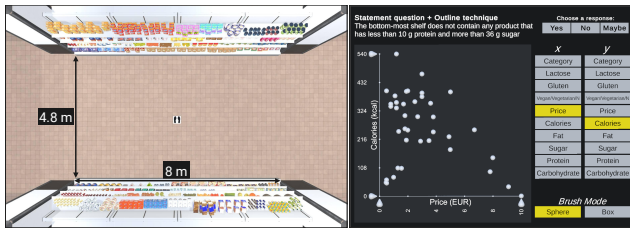


Figure 3: Left: A top-down view of the virtual environment in which the study takes place. Right: A close-up view of the panel and the user interface prompting a user response.

multiple physical referents, such as browsing cars based on mileage and price, or planning seating arrangements at a wedding based on guests’ age and relationship to the bride and groom.

Single-selection tasks. Identify, select, and locate a single product (referent) that matches a bivariate criterion, e.g., “*Select the product that has the highest fat and no sugar.*” This task helps determine a technique’s effectiveness in locating a single referent.

Multi-selection tasks. Identify, select, and locate all products that match a bivariate criterion, e.g., “*Select all products that have more than 22.5 g protein and less than 9.5 g fat.*” This task helps determine a technique’s effectiveness in locating multiple referents, as it requires highlighting to be applied to multiple products simultaneously. All questions have five products that match the given criterion to ensure consistency across trials.

Statement response tasks. Respond with “Yes”, “No”, or “Maybe” to a given statement that describes the dataset and supermarket layout, e.g., “*Vegan products that contain more than 51 g of carbohydrate are distributed only in the center-left and center-right shelves.*” This task requires participants to observe a collection of referents together and make a judgment on their spatial arrangement. The task is also reminiscent of standard brushing and linking, wherein all brushed points should be visible and considered when judging correlations and dependencies in the linked view [Kei02].

We used a curated dataset that contains 11 variables (7 quantitative, 4 categorical) for 98 products, which we placed on the supermarket shelves. This data was manually scraped from supermarket websites, and all values were checked for plausibility. Minor adjustments were made to these values to accommodate the above tasks. The *Inside-FOV* condition only uses a reduced dataset of 39 products fitting in the front-facing set of shelves.

4.3. Experimental setup

We used the experimental prototype described in Section 3. All instructions and buttons used to answer the statement response tasks (“Yes”, “No”, “Maybe”) are shown in the upper area of the tablet (see Figure 3 right). Additionally, functionality was added for the single- and multi-selection tasks to indicate that they have “identified” the required products. As we are not interested in the usability of brushing scatterplots in VR, we decided to streamline the process after pilot testing revealed it to be too cumbersome. Red visual cues were added to indicate the two required data dimensions for each task and to indicate the required value ranges along the axes. The filter handles were modified to “snap” to these range

thresholds. This simplification does not confound our primary focus on the “linking”. Figure 2 shows the steps involved in solving a multi-selection task, including indicating identified referents. In terms of movement, we neither forced participants to move nor restricted their movement. When asked for clarification, we told them that all products could be seen just by turning.

We conducted our study using a desktop PC equipped with an Intel Core i9-9980XE CPU, RTX 2080 Ti with 11 GB VRAM, and 128 GB of RAM. We used an Oculus Quest Pro headset with two hand controllers. The headset was connected via a 5 m Quest Link USB cable to the PC running our application developed in Unity version 2022.2.20f1. All questionnaires were administered on the same PC using LimeSurvey. The study was carried out in an open space of 360 cm by 220 cm that was free of obstructions.

4.4. Measures

Task Performance. We measured task performance using completion time and errors. For each task, we recorded the overall completion time, the time of the last interaction with the tablet (i.e., button presses, filter changes, and brushing), and the time when the first product was selected (Figure 2, frame 4). The overall completion time is the interval between starting the task and selecting the final product or responding to the given argument via the respective button. Incorrectly selected products are counted as errors.

User Experience. We measured user experience using two questionnaires. A questionnaire was completed after each shelf layout for each technique, and another, final questionnaire, after all conditions. The in-between questionnaire included NASA-TLX [Har06], free-text forms for comments on user strategies and experiences, as well as additional feedback. The physical environment might not be as accessible as the desktop visualizations and might introduce additional factors that influence certain highlighting techniques. Hence, we came up with the following categories that such techniques should support through internal discussions, our own experiences and findings with SitA, and reported findings in the literature. The final questionnaire included these categories using 9-point Likert scales (1=worst, 9=best, unless stated otherwise).

General experience. A subjective rating of how the highlighting technique performed in general.

Out-of-view. A highlighting technique ideally accounts for physical referents being out of view or occluded, since identifying can be equally challenging for virtual and physical objects.

Visual clutter. A highlighting technique ideally supports marking multiple referents without clutter, since scalability is often desired for both infovis [RPA*22] and SitA [BPR20].

Obstruction. A highlighting technique ideally minimizes occluding its surrounding and obstructing the task, as distracting or reducing the user’s real world awareness can be dangerous (1=not at all, 9=very).

Subtlety. A highlighting technique ideally balances between saliency and subtlety, as some SitA contexts (e.g., for safety) require a more subtle highlighting technique (1=very subtle, 9=very obtrusive).

Explicit connections. A highlighting technique ideally supports direct connections of data marks and referents, since making assumptions across views is common in visual analytics [Rob07] and potentially in SitA scenarios as well.

Recognition. A highlighting technique ideally balances between identifying and obscuring the target, as some real-world tasks (like grocery shopping) might require to be aware of the actual appearance while others do not (1=not at all, 9=very).

Enjoyment. A highlighting technique balances between usability and enjoyment resulting from an aesthetically pleasing design.

The relevance of the categories varies between use cases. However, we collected all categories throughout the experiments to obtain a full characterization of possible situated brushing and linking contexts. The final questionnaire also asked participants to describe their strategies for different techniques and shelf layout conditions, and provide general feedback.

4.5. Study design and procedure

We used a within-subjects study design to evaluate our two independent variables: highlighting technique (*Color*, *Outline*, *Link*, *Arrow*); and shelf layout (Inside-FOV, Outside-FOV). Each condition consisted of three tasks, one per each task type (single-selection, multi-selection, statement response). Each question was randomly selected from a set of eight predefined questions per type of task. Thus, each participant had to perform four highlighting techniques \times two shelf layouts \times three tasks = 24 trials. We counterbalanced the order of the highlighting technique using Latin squares. However, we used a fixed order for both the shelf layout (Inside-FOV \rightarrow Outside-FOV) and the task type (single \rightarrow multi \rightarrow statement), because these two aspects have a naturally increasing level of difficulty. Participants were compensated with €12.

Introduction (5 minutes). The participant was welcomed and introduced to the purpose of the study. After signing a consent and privacy form, they were asked for their handedness, asked to fill in a demographic questionnaire, and introduced to the VR headset and controllers. They were informed that they could take breaks or withdraw at any point in the study without consequences.

Training (10-15 minutes). The experimenter first explained the VR prototype to the participant, describing the supermarket and tablet, the available interactions and the tasks they would be performing. PowerPoint slides with pictures were also shown during the briefing. The participant then put on the VR headset and practiced brushing and linking interactions using a simplified layout of 11 products and the aforementioned *size* highlighting technique. In this phase, no task was assigned, and the participant could freely try the interactions without time pressure.

Main study (40-50 minutes). The participant was then given each of the conditions and tasks in the counterbalanced order described above. Whenever a new highlighting technique was introduced, the participant was given a short tutorial using another simplified layout of four products to let them see what the highlighting technique looked like. Between each trial, the participant was told to return to the center of the room and face the direction indicated by the image on the floor (see Figure 3 left). After completing each set of three tasks (single-selection, multi-selection, statement




Figure 4: Completion time measures using mean and 95% CIs.

response), the participant removed the VR headset to answer the in-between questionnaire. After a short break (if needed), they put the headset back on to proceed with the study.

Post study (5-10 minutes). After all trials were completed, the participant answered a final questionnaire. Then they were given compensation as a reward for their participation.

4.6. Guiding questions

For the analysis of our exploratory study, we devised four guiding questions, along with some of our prior expectations:

- **GQ1: Which techniques performs the best in terms of task performance, and which the worst?** We expect that *Link* will have the fastest completion time due to its guidance effect, even to out-of-view targets, as well as its direct connection between data marks and its referents. We also expect that *Arrow* leads to the slowest completion time for the multi-selection task due to its limited scalability. We have no expectations about error rates.
- **GQ2: Which technique is subjectively preferred or disliked?** *Link* fits many required features of our tasks and requirements (e.g., out-of-view highlighting or direct connection of data marks and their referents). We assume that it is perceived as the easiest and fastest technique and expect it to be preferred.
- **GQ3: Which technique is the least preferred if many targets are not in view?** We expect *Arrow* to be perceived worst, as it is not scalable and, therefore, clutters an environment containing many targets. For targets outside the view, we expect *Color* and *Outline* to be the least preferred, as they do not have guidance and no direct connections between data marks and referents.
- **GQ4: Which technique is perceived as the most subtle and least obstructive?** We expect *Outline* to be deemed the most subtle and least obstructive, as it obscures only the boundary of the targets and does not cause clutter even with many targets.

5. Results

We recruited 20 participants (8 female and 12 male) via university mailing lists and social media, with different age groups between 18 and 55 years (mode=26–30). 19 participants had correct or corrected-to-normal vision, while one participant had astigmatism. We had 16 participants who were right-handed, and 4 participants who were left-handed. Most of our participants had some visualization experience (16/20 between 2 and 4 on a 5-point Likert

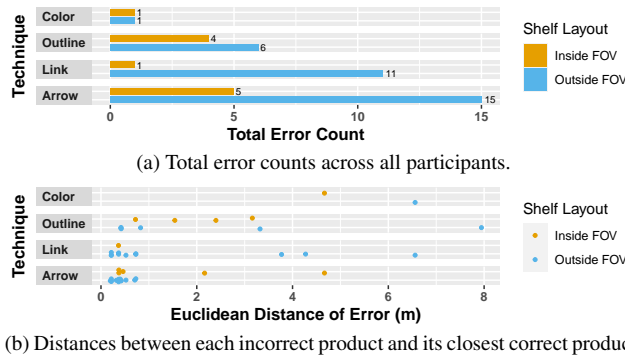


Figure 5: Error measures for the single- and multi-selection task.

scale), and only a few participants had little (4/20 between 0 and 1). More than half of the participants had some VR experience (12/20 between 2 and 4 on a 5-point Likert Scale), and the remaining participants had little (8/20 between 0 and 1).

Following the recommendations of Cumming [Cum14] and the American Psychological Association [Ame20], we visually analyze the data based on their mean and 95% confidence intervals using forest plots. In general, the evaluation and subsequent discussion focuses more on the differences between the highlighting techniques than on the shelf layout conditions.

5.1. Task performance

Completion time and errors are shown in Figure 4 and Figure 5.

Completion Time Since our analysis focuses on the linking part of the task, we calculated *linking completion time* defined as the overall completion time minus the time of the last interaction on the tablet. During the study, we noticed that some participants did not brush all data marks at once, but rather brushed one at a time when using certain techniques. For these cases, we manually calculated the brushing time and linking time from the recorded video.

For the **overall linking completion time**, which includes all tasks, we received 24 values per participant (4 techniques × 2 shelf layouts × 3 tasks). To aggregate these values for each technique and shelf layout per participant, we first averaged the values before calculating the overall mean and the respective confidence interval. The results of the overall linking completion time can be seen in Figure 4a. The **Inside-FOV** linking CT is generally faster than the **Outside-FOV** for all techniques. *Color* and *Link* perform similarly; *Link* is slightly better for the Inside-FOV, but worse for the Outside-FOV. While *Outline* resulted in similar results for the Inside-FOV, the results indicate a substantial performance drop for Outside-FOV targets. *Outline* is the only technique with a large gap between the two shelf layout conditions. *Arrow* is the worst for the Inside-FOV, and, second worst for Outside-FOV.

We also evaluated the **linking completion time for the multi-selection task**. We were specifically interested in this task, as multi-selection constitutes the "core discipline" of brushing and linking. The results are shown in Figure 4b. As before, the Inside-FOV has lower values than the Outside-FOV. Overall, *Color* performs best, closely followed by *Link*, as its Outside-FOV completion time is slightly higher. Again, *Outline* has a large gap between the shelf




Figure 6: Mean and 95% CIs of the category ratings.

layout conditions, with the worst results for the Outside-FOV. *Arrow* is the worst for the Inside-FOV. By and large, the results for the multi-selection task are very similar to the overall results above, with a slightly stronger articulation of the main characteristics.

Errors We also measured **errors** as incorrectly selected products in single- or multi-selection tasks and removed duplicate errors caused by selecting the same incorrect product repeatedly. Figure 5a shows the total count of errors among all participants. *Color* had the lowest number of errors and *Arrow* the most, with Outside-FOV being more prone to errors, especially for *Link* and *Arrow*. In addition, Figure 5b shows for each error how far the correct product was in meters. From this result, we see that most of the errors made using *Color* and *Outline* were more than 1 m away, mainly due to incorrect brushing on the scatterplot. This was confirmed by reviewing the respective video footage. In contrast, for *Link* and *Arrow*, many errors were less than 0.5 m away, suggesting that participants actually mistook the product they needed to select—with some also due to incorrect brushing. Overall, *Link* and *Arrow* were more prone to errors when identifying a highlighted product.

5.2. User experience

We now report questionnaire results of the category ratings and NASA-TLX.

Category Rating The results of the 9-point Likert scale ratings are in Figure 6. High values are better for all categories, except for obstruction, subtlety, and recognition, where low values are better.

- **General experience** *Color* and *Link* performed similarly, followed (with large gaps) by *Outline* and, lastly, *Arrow*. While *Outline* received a mediocre rating, *Arrow* was rated very low.
- **Out-of-view** *Link* was ranked the best. With a considerable gap; *Color*, *Arrow*, and *Outline* followed.
- **Visual clutter** *Color* and *Link* resulted in the best scores, followed by *Outline*. With a considerable gap, *Arrow* was the worst and considered to be the most visually cluttered.




Figure 7: Mean and 95% CIs of the NASA-TLX subscales mental demand, performance, effort and frustration.

- **Obstruction** (low better): *Color*, *Outline*, and *Link* all performed very well, while *Arrow* was worst by a large margin.
- **Subtlety** (low better): *Outline* was perceived by far as the most subtle technique. The other three were not considered subtle.
- **Explicit connections** *Link* wins with the highest score.
- **Recognition** (low better): *Link* was judged the best for recognizing the product, followed by *Outline*. *Color* and *Arrow* resulted in mediocre scores.
- **Enjoyment** Enjoyment results are similar to **General experience**, with the highest enjoyment for *Link* and *Color*, followed by *Outline* and, far behind, *Arrow*.

Color and *Link* performed the best for all ratings. *Outline* was the sole favorite in **subtlety**, and *Link* in **explicit connection**. *Arrow* was the worst in six categories, and second to last in the other two.

NASA-TLX We briefly report on four subscales of the NASA-TLX: **mental demand**, **performance**, **effort**, and **frustration** (Figure 7). Physical and temporal demand are not reported here as the task is not physically intensive nor does it have time constraints, though their scores were similar to mental demand. Please see the supplemental material for the remaining NASA-TLX data.

For **performance**, Inside-FOV scores are slightly higher than Outside-FOV scores but are still comparable. *Color* and *Link* have similar high scores, followed by *Outline* and *Arrow*. Overall, there are no large differences between the techniques. The remaining results for **mental demand**, **effort**, and **frustration** show similar patterns. In these three subscales, scores were similar between both shelf layouts. *Link* and *Color* are followed by *Outline* and *Arrow*. While *Link* and *Color* have a similarly low score for **effort**, *Link* slightly outperforms *Color* for **mental demand** and **frustration**.

Overall, the results resemble those of the category ratings: *Color* and *Link* performed the best, followed by *Outline*, and lastly *Arrow*. However, they are much less pronounced in NASA-TLX.

5.3. Qualitative analysis and comments

As mentioned in Section 4, we also included free-text forms for further feedback. In addition to mentioning strategies, there were comments on our task design, as well as on brushing and filtering methods. As this work focuses on highlighting and linking, we focus on feedback on these issues.

Tasks can have—but do not need—strategies. Regarding strategies, one participant mentioned that they first “*read the task very quickly while selecting the [components]*” (P6). For a later question, they clarified that they “*expanded [their] previous strategy with the red stripes on the axes*” (P6). Another participant (P16) described a similar strategy for the statement task, since they also

“*[looked] at named shelves before tending to the task*” (P16). For the multi-selection task, one participant stated that they “*[selected] the products on the shelves going from left to right*” (P16). Another participant claimed that they adapted their strategy “*[to select] one object at a time instead of selecting all at once*” (P3) for the multi-selection and statement task, after trying out the *Arrow* condition.

Color and Outline require more movement. A quarter of our participants said that it was easier to solve tasks with *Color*, and they were able to “*[select] multiple products at once*” (P3) with *Color* and *Outline*. However, both techniques subjectively felt to involve more movement, since there is no “*indication for objects outside the field of view*” (P12), and, hence, “*[they] had to look at the shelves to find them, and had to look twice to make sure that [they] had not forgotten any item*” (P1).

Outline is harder to perceive than Color. Some participants argued that “*for Outline, [they had] to pay more attention compared to Color because its colorful edge has a small size*” (P7) and “*[it requires] more mental processing and visual search*” (P18). Furthermore, the participants argued that some products had similar boundaries, and, hence, *Outline* was harder to perceive as the contrast was “*not always optimal*” (P17).

Link works well for out-of-view but worse in peripheral regions. According to several participants, *Link* makes multi-selection tasks easier as they can “*just follow the line [...] to select the item*” (P1) and “*needed to pay less attention to the shelves directly*” (P18). *Link* is also effective for out-of-view targets as “*the lines [helped] to indicate that the highlighted products are behind [you] without the need to turn around*” (P20). However, some participants criticized the visibility of *Link* in the peripheral regions. According to P2, these were ambiguous, and P12 said if the tablet hides some lines, they may be missed.

Arrow requires patience and strategies for many targets. Although P1 was able to follow the path of *Arrow*, most of the participants criticized that *Arrow* is “*annoying*” (P7) and “*exhausting for the eyes*” (P13). Although the general direction can be guessed, the movement of the arrow “*[seemed] to be influenced by user orientation*” (P17), and, hence, it is difficult to find the actual target. According to several participants, “*[they had] to wait to see where the arrows go, for each product separately*” (P5) and “*if [they missed] the exact moment it hits, [they] had to wait again for the next one*” (P14). Therefore, they tried to “*keep fully still*” (P7). Several participants also argued that they performed strategies for multi-selection and statement response tasks, as *Arrow* produced a “*very crowded visual field*” (P18). For example, participants “*tried to crouch to see more clearly if the arrows were flying only upward*” (P17) or “*searched for an outlier*” (P18).

6. Discussion

We conclude our paper with a summary of our observations, implications for future work, limitations, and some final remarks. Section 6.1 discusses the results based on the guiding questions.

6.1. Observations and lessons learned

Color and Link perform the best in terms of completion time, followed by Outline and Arrow. Our results show that *Color* and *Link* resulted in the lowest completion time. *Color* was slightly faster for the Outside-FOV, whereas *Link* was slightly faster for the Inside-FOV. While we expected *Link* to perform well as per GQ1, *Color* was surprisingly able to handle the Outside-FOV condition despite its lack of support for out-of-view and explicit link representations. We assume that this is due to the strong contrast of *Color*, which makes it a valid choice when a target search area or direction is known in advance. *Link* may have underperformed for Outside-FOV as it would pass directly through the participant's own body, making it difficult to visually trace it. Although *Outline* highlights a target in a similar manner as *Color*, it has similar mean completion time for the Inside-FOV, but not for the Outside-FOV. According to multiple participants (P2, P5, P7), with a more solid and better contrast of *Outline*, it would have similar benefits as *Color*. For example, P2 argued that "[Outline was] much harder to see [...] than the products that were highlighted [by Color]" (P2). Lastly, *Arrow* was clearly the slowest as expected in GQ1. Participants clearly summarized the issue: "for [Arrow] I had to wait to be sure where it goes" (P5), while other techniques such as *Color* and *Link* "made it possible to solve problems at a glance" (P18). This aligns with known reasoning for why static representations can be superior to animations [HR07, TMB02]. It is possible that a combination of static highlights with salient animations could be effective, though at the cost of increased visual clutter.

Color had the least amount of errors, with Outline, Link, and Arrow performing increasingly worse. We found that *Color* was clearly the most accurate highlighting technique, which once again may be due to its strong contrast. Note that, in several instances, errors were caused by mistakes in brushing the scatterplot or selecting the right product—likely caused by the challenging raycast interactions. However, as these issues affect all highlighting techniques uniformly, we believe that this does not invalidate our results. Therefore, the lack of errors of less than 1 m for both *Color* and *Outline* indicates that, when the product was correctly brushed, they resulted in little to no mistakes. In contrast, *Link* and *Arrow* performed much worse for errors, probably because they only indicate the center point of the target. As they do not indicate the contours of the target, like *Color* or *Outline* do, participants instead selected similar-looking nearby variants (e.g., different yogurt flavors). Therefore, we argue that, when a task involves similar varieties or distractors, *Color* and *Outline* can reduce ambiguities over *Link* and *Arrow*—especially if it is unclear if the referents stand alone individually or are grouped as one.

Color and Link are subjectively preferred. We found that *Color* was generally preferred over *Link*. However, *Color* and *Link* differ significantly from *Outline* and *Arrow*. Similar results can be seen for enjoyment. We argue for GQ2 that both *Color* and *Link* are preferred over the remaining techniques. Given that *Color* and *Link*

address different properties for a highlighting technique in situated analytics, it seems that the subjective preference for either of them would depend on the use case. A distinct advantage of *Color* is that it is easy to understand and has a bright color. One participant stated that "[Color] seemed the easiest [technique] to me" (P20). Unlike *Color*, *Link* supports finding targets which are out of view, as one can solely follow the line from its beginning to its end. Several participants mentioned that *Link* highlighted targets behind them well (P12, P17, P18). However, in the periphery, it can be difficult to accurately determine the end point of a link (P2). Hence, we speculate that increasing the link thickness with increasing distance or using a combination of *Color* and *Link* may increase overall satisfaction, especially for real-world awareness.

Arrow is the least preferred despite its ability to direct to out-of-view targets. Since *Arrow* clearly has the lowest results for general experience and enjoyment, we can declare it the least preferred technique for GQ2 and GQ3, at least, for our scenario. This observation is also reflected in the comments of our participants, who described *Arrow* as "irritating" (P1), "annoying" (P7) or even "exhausting for the eyes" (P13). Gruenfeld et al. [GLH*18] had similar task performance results, but found a lower workload for *Arrow*, which could be due to our scenario with its many targets. Our animated *Arrow* resulted in worse results than the static arrows of Wieland et al. [WGRF22] and Yu et al. [YLF*20], which could be due to its usage as a baseline condition based on familiarity reasons. Since Wieland et al. [WGRF22] only considered a single target at a time, we assume that their arrow approach would also result in clutter for multiple targets. One reason might be the clutter introduced by *Arrow* that affects its ability to highlight many targets. P12 observed that "[Arrow is] obstructing too much of the view, especially when there are multiple targets" (P12). However, *Link* was not perceived as cluttered, suggesting that animation might cause poor scalability. Additionally, all previous approaches [GLH*18, WGRF22, YLF*20] included estimation tasks that could have led to better results for *Arrow* if applied in our context. One could modify *Arrow* such that it is not affected by head movements if many targets are present. However, the clutter of *Arrow* remains a major drawback.

Outline is the most subtle and least obstructing. While *Outline* has the lowest mean rating in terms of subtlety, all methods except *Arrow* rank similarly in terms of obstruction. However, comments from participants, such as "I had to look at each product to be really sure it is not highlighted" (P5), suggest that *Outline* is indeed the most subtle regarding our GQ4. We assume that the contours of *Outline* are perceived as part of the real scene and not as artificial highlighting. Although we expected *Link* to be subtle as well, the end of a line appears more artificial. One reason for the results of *Arrow* may be the animations, which potentially obstruct the field of view. While subtlety may not be a virtue for fast visual search or classical brushing and linking applications, there are scenarios where a subtle, non-obstructive technique with high real-world is beneficial. For example, we want to avoid walking into an object or another person while moving toward the selected referent.

Brushing in situated analytics matters. We also received several comments on the brushing component of our prototype. Participants agreed that it was easy to brush, as the "red marked areas on

the axis [helped] a lot" (P12), indicating that their cognitive load was reduced as intended. Participants however still struggled with adjusting the filter handles using the raycast interactions, which "affected [their] frustration" (P17). P16 also wanted "to see how many [data marks] are selected on the tablet", which would be needed in a proper SitA application. Any such improvements to brushing would naturally benefit any immersive application, not just SitA.

6.2. Example scenarios and applications

Based on our results, we see the following practical usage for our techniques. *Color* is valid to use for targets behind the user due to its bright color. However, the scenarios should not require the perception of the actual appearance, but only the silhouette of the target. We further assume that the general direction of the target must be known in advance. Unlike the more obtrusive *Color*, *Outline* is perceived as more subtle, but harder to see. However, one can see the actual referent's appearance. For example, guiding a trainee to a desired tool during service. In contrast to *Color* and *Outline*, *Link* might be a better choice for targets that have a greater distance and are potentially occluded by other objects. We suspect *Link* to be a valid choice in wide search spaces such as a library in which books are distributed across many shelves and floors. Based on the results, *Arrow* required more work to find the correct target. We assume that scenarios which primarily require direction guidance benefit from a limited number of arrows.

6.3. Future directions

We intentionally used baseline highlighting techniques established in brushing and linking and attention guidance. Based on our results, there are clear avenues for improving each of the four techniques. Although we did not directly evaluate the matter, *Color* lacks a way to perceive the actual target (since it is distorted by the color). A modified *Outline* with increased contrast and outline width could inherit most of the benefits of *Color*, while not (fully) occluding the underlying targets. For *Link*, we suggest that adjusting size based on distance, as well as adding a way of perceiving lines behind the tablet or user, would further increase the already encouraging results. However, care has to be taken to avoid excessive visual clutter [HRD*19]. We assume that *Arrow* might be more beneficial if the user's movements did not have an impact on its path, or if used as a supplementary indicator combined with *Color* or *Outline*. In general, it seems likely that a hybrid technique can overcome the weaknesses of individual techniques, provided that we can avoid too much clutter. However, the optimal choice will likely depend on the scenario. For example, applying *Color* to large objects can be visually overwhelming, even if *Color* was the most favored overall. Future work may seek to devise an adaptive highlighting technique to tune the visual characteristics of the highlighting to the current viewing situation.

Future work should also investigate brushing. If brushing in reverse, i.e., brushing physical referents, an analyst could compare the values of nearby referent after physically navigating to an interesting real-world area, such as inspecting the bargain bin in the supermarket. Several techniques could support brushing in this manner, including spatial filtering [FMS93, JLB*00], aggregation [TOK*16], and selection [SCZ*20] techniques. We leave the exploration of this inverse brushing for future work.

While we believe the supermarket is generalizable to other scenarios, our study lacked several conditions that should be explored in future work. First, our shelf layouts did not force participants to move because of occlusions. We assume out-of-view highlighting techniques such as *Link* are more beneficial here, especially, since they can also help with wayfinding. Second, we chose only static referents, whereas referents can also be dynamic (e.g., autonomous robots, animals). Third, our referents are at a comfortable human scale when compared to buildings or landscapes. These characteristics not present in our scenario require further evaluation.

6.4. Limitations

Some limitations can potentially affect the interpretation of the findings reported in our work. First, our small sample size and the chosen tasks could affect the results and design considerations of our work. In addition, we adjusted and calculated our linking completion time based on the last selection on the tablet, assuming that participants would first brush all data marks and then select the respective targets on the shelves. While we manually adjusted the timestamps for participants who obviously used a different strategy, it is possible that some participants missed a data mark during the multi-selection task. Hence, some timestamps may be non-representative. Since we used an open experimenter study design, we assume that this rarely happened and does not limit our results.

Another limitation is our use of VR to simulate AR, even if transferability from VR to AR has been demonstrated for other tasks [LBHB09, LRM*13]. We assume that any differences observed between the highlighting techniques in our study will still translate into AR, and the reasons for using one technique over the other remain the same. However, it is likely that their effectiveness in AR will largely be dictated by the technical capabilities of the AR headset and tracking solution [CFSS23]. For example, tracking drift may cause noticeable misalignment in object silhouettes when using *Color*, which may justify the use of other techniques less sensitive to this issue like *Link* or *Arrow*. The study of these considerations lies firmly outside the scope of this work.

7. Conclusions

We present an initial exploration of brushing and linking for situated analytics in AR. We describe four highlighting techniques, representing conventional approaches to brushing and linking and attention guidance. We used these techniques in an exploratory user study and found that *Color* is still a reasonable choice as a highlighting technique for situated brushing and linking. We also saw that *Arrow*, a highlighting technique that requires spatial interpretation, performed worst in terms of both task performance and user experience. Based on our overall results, we propose that a combination of highlighting techniques might reap the benefits while overcoming individual weaknesses. We consider our work a starting point to bring brushing and linking into the real and merging it with previous techniques designed for "reading" the environment.

Acknowledgments. We thank DFG (495135767 and EXC 2120/1-390831618), Carl-Zeiss-Stiftung (P2016-03-004), FWF (I5912-N), and the Alexander von Humboldt Foundation funded by the German Federal Ministry of Education and Research.

References

- [Ame20] AMERICAN PSYCHOLOGICAL ASSOCIATION (WASHINGTON, DISTRICT OF COLUMBIA) (Ed.): *Publication Manual of the American Psychological Association*, seventh edition ed. 2020. 7
- [APHD24] ASSOR A., PROUZEAU A., HACHET M., DRAGICEVIC P.: Handling Non-Visible Referents in Situated Visualizations. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 30, 1 (2024), 1336–1346. doi:10.1109/TVCG.2023.3327361. 2
- [AWG*15] AHN J., WILLIAMSON J., GARTRELL M., HAN R., LV Q., MISHRA S.: Supporting Healthy Grocery Shopping via Mobile Augmented Reality. *ACM Transactions on Multimedia Computing, Communications, and Applications* 12, 1s (2015), 16:1–16:24. doi:10.1145/2808207. 2, 3
- [BKT*22] BRESSA N., KORSGAARD H., TABARD A., HOUBEN S., VERMEULEN J.: What's the Situation with Situated Visualization? A Survey and Perspectives on Situatedness. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 28, 1 (2022), 107–117. doi:10.1109/TVCG.2021.3114835. 1, 3
- [BNC*03] BOWMAN D. A., NORTH C., CHEN J., POLYS N. F., PYLA P. S., YILMAZ U.: Information-rich virtual environments: Theory, tools, and research agenda. In *ACM Symposium on Virtual Reality Software and Technology (VRST)* (2003), p. 81–90. doi:10.1145/1008653.1008669. 3
- [BPR20] BÜTTNER S., PRILLA M., RÖCKER C.: Augmented Reality Training for Industrial Assembly Work - Are Projection-based AR Assistive Systems an Appropriate Tool for Assembly Training? In *ACM Conference on Human Factors in Computing Systems* (2020), pp. 1–12. doi:10.1145/3313831.3376720. 5
- [BR03] BAUDISCH P., ROSENHOLTZ R.: Halo: A technique for visualizing off-screen objects. In *ACM Conference on Human Factors in Computing Systems* (2003), p. 481–488. doi:10.1145/642611.642695. 3
- [BSEN18] BORK F., SCHNELZER C., ECK U., NAVAB N.: Towards efficient visual guidance in limited field-of-view head-mounted displays. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 24, 11 (2018), 2983–2992. doi:10.1109/TVCG.2018.2868584. 3
- [BT0X06] BIOCCA F., TANG A., OWEN C., XIAO F.: Attention funnel: Omnidirectional 3d cursor for mobile augmented reality platforms. In *ACM Conference on Human Factors in Computing Systems* (2006), p. 1115–1122. doi:10.1145/1124772.1124939. 3
- [CC07] COLLINS C., CARPENDALE S.: VisLink: Revealing Relationships Amongst Visualizations. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 13, 6 (2007), 1192–1199. doi:10.1109/TVCG.2007.70521. 3
- [CCD*17] CORDEIL M., CUNNINGHAM A., DWYER T., THOMAS B. H., MARRIOTT K.: ImAxes: Immersive Axes as Embodied Affordances for Interactive Multivariate Data Visualisation. In *Annual ACM Symposium on User Interface Software and Technology* (2017), pp. 71–83. doi:10.1145/3126594.3126613. 4
- [CFSS23] CALEPSO A. S., FLECK P., SCHMALSTIEG D., SEDLMAIR M.: Exploring Augmented Reality for Situated Analytics with Many Movable Physical Referents. In *ACM Symposium on Virtual Reality Software and Technology (VRST)* (2023), pp. 1–12. doi:10.1145/3611659.3615700. 3, 10
- [CKB09] COCKBURN A., KARLSON A., BEDERSON B. B.: A review of overview+detail, zooming, and focus+context interfaces. *ACM Comput. Surv.* 41, 1 (2009). doi:10.1145/1456650.1456652. 2, 3
- [Cum14] CUMMING G.: The new statistics: Why and how. *Psychological Science* 25, 1 (2014), 7–29. PMID: 24220629. doi:10.1177/0956797613504966. 7
- [DARS23] DOERR N., ANGERBAUER K., REINELT M., SEDLMAIR M.: Bees, birds and butterflies: Investigating the influence of distractors on visual attention guidance techniques. In *ACM Conference on Human Factors in Computing Systems Extended Abstracts (CHI-EA)* (2023). doi:10.1145/3544549.3585816. 3
- [DL24] DOERR N., LEE B.: Vishigh, 2024. URL: <https://github.com/doerrna/VisHigh.3>
- [DMT*18] DILLMAN K. R., MOK T. T. H., TANG A., OEHLBERG L., MITCHELL A.: A visual interaction cue framework from video game environments for augmented reality. In *ACM Conference on Human Factors in Computing Systems* (2018), p. 1–12. URL: <https://doi.org/10.1145/3173574.3173714>, doi:10.1145/3173574.3173714. 3
- [EBC*21] ENS B., BACH B., CORDEIL M., ENGELKE U., SERRANO M., WILLET W., PROUZEAU A., ANTHES C., BÜSCHEL W., DUNNE C., DWYER T., GRUBERT J., HAGA J. H., KIRSHENBAUM N., KOBAYASHI D., LIN T., OLAOSEBIKAN M., POINTECKER F., SAFFO D., SAQUIB N., SCHMALSTIEG D., SZAFIR D. A., WHITLOCK M., YANG Y.: Grand challenges in immersive analytics. In *ACM Conference on Human Factors in Computing Systems* (2021). doi:10.1145/3411764.3446866. 2
- [EST16] ELSAYED N. A. M., SMITH R. T., THOMAS B. H.: HORUS EYE: See the Invisible Bird and Snake Vision for Augmented Reality Information Visualization. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)* (2016), pp. 203–208. doi:10.1109/ISMAR-Adjunct.2016.0077. 3
- [ETM*15] ELSAYED N., THOMAS B., MARRIOTT K., PIANTADOSI J., SMITH R.: Situated Analytics. In *Big Data Visual Analytics (BDVA)* (2015), IEEE, pp. 1–8. doi:10.1109/BDVA.2015.7314302. 1, 2, 3
- [ETM*16] ELSAYED N. A. M., THOMAS B. H., MARRIOTT K., PIANTADOSI J., SMITH R. T.: Situated Analytics: Demonstrating immersive analytical tools with Augmented Reality. *J. Vis. Lang. Comput.* 36, C (2016), 13–23. doi:10.1016/j.jvlc.2016.07.006. 1, 2
- [FMS93] FEINER S., MACINTYRE B., SELIGMANN D.: Knowledge-based augmented reality. *Commun. ACM* 36, 7 (1993), 53–62. doi:10.1145/159544.159587. 10
- [FP21] FONNET A., PRIÉ Y.: Survey of Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 27, 3 (2021), 2101–2122. doi:10.1109/TVCG.2019.2929033. 2
- [GAM18] GROGORICK S., ALBUQUERQUE G., MAGNOR M.: Gaze guidance in immersive environments. In *IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2018), pp. 563–564. doi:10.1109/VR.2018.8446215. 3
- [GCC17] GUTWIN C., COCKBURN A., COVENEY A.: Peripheral Popout: The Influence of Visual Angle and Stimulus Intensity on Popout Effects. In *ACM Conference on Human Factors in Computing Systems* (2017), pp. 208–219. doi:10.1145/3025453.3025984. 4
- [GCOK21] GRANDI J. G., CAO Z., OGREN M., KOPPER R.: Design and simulation of next-generation augmented reality user interfaces in virtual reality. In *IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)* (2021), pp. 23–29. doi:10.1109/VRW52623.2021.00011. 3
- [GGL*14] GRATZL S., GEHLENBORG N., LEX A., PFISTER H., STREIT M.: Domino: Extracting, Comparing, and Manipulating Subsets Across Multiple Tabular Datasets. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 20, 12 (2014), 2023–2032. doi:10.1109/TVCG.2014.2346260. 3
- [GLH*18] GRUENEFELD U., LANGE D., HAMMER L., BOLL S., HEUTEN W.: FlyingARrow: Pointing Towards Out-of-View Objects on Augmented Reality Devices. In *ACM International Symposium on Pervasive Displays (PerDis)* (2018), pp. 1–6. doi:10.1145/3205873.3205881. 1, 3, 4, 9
- [GSL*14] GEYMAYER T., STEINBERGER M., LEX A., STREIT M., SCHMALSTIEG D.: Show me the invisible: Visualizing hidden content. In *ACM Conference on Human Factors in Computing Systems* (2014), pp. 3705–3714. doi:10.1145/2556288.2557032. 3
- [Har06] HART S. G.: Nasa-task load index (nasa-tlx); 20 years later. *Human Factors and Ergonomics Society Annual Meeting* 50, 9 (2006), 904–908. doi:10.1177/154193120605000909. 5

- [HRD07] HEER J., ROBERTSON G.: Animated Transitions in Statistical Data Graphics. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 13, 6 (2007), 1240–1247. doi:10.1109/TVCG.2007.70539.9
- [HRD*19] HURTER C., RICHE N. H., DRUCKER S. M., CORDEIL M., ALLIGIER R., VUILLEMOT R.: Fiberclay: Sculpting three dimensional trajectories to reveal structural insights. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 25, 1 (2019), 704–714. URL: <https://doi.org/10.1109/TVCG.2018.2865191>, doi:10.1109/TVCG.2018.2865191.10
- [JHPR11] JO H., HWANG S., PARK H., RYU J.-H.: Aroundplot: Focus+context interface for off-screen objects in 3d environments. *Computers & Graphics* 35, 4 (2011), 841–853. *Semantic 3D Media and Content*. doi:10.1016/j.cag.2011.04.005.3
- [JLB*00] JULIER S., LANZAGORTA M., BAILLOT Y., ROSENBLUM L., FEINER S., HOLLERER T., SESTITO S.: Information filtering for mobile augmented reality. In *IEEE and ACM International Symposium on Augmented Reality (ISAR)* (2000), pp. 3–11. doi:10.1109/ISAR.2000.880917.10
- [KCWK20] KREKHOV A., CMENTOWSKI S., WASCHK A., KRÜGER J.: Deadeye visualization revisited: Investigation of preattentiveness and applicability in virtual environments. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 26, 1 (2020), 547–557. doi:10.1109/TVCG.2019.2934370.3
- [Kei02] KEIM D. A.: Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 8, 1 (2002), 1–8. doi:10.1109/2945.981847.2,5
- [KFS*22] KRAUS M., FUCHS J., SOMMER B., KLEIN K., ENGELKE U., KEIM D., SCHREIBER F.: Immersive analytics with abstract 3d visualizations: A survey. *Computer Graphics Forum (CGF)* 41, 1 (2022), 201–229. doi:<https://doi.org/10.1111/cgf.14430>.2
- [KKO*14] KISHISHITA N., KIYOKAWA K., ORLOSKY J., MASHITA T., TAKEMURA H., KRUIJFF E.: Analysing the effects of a wide field of view augmented reality display on search performance in divided attention tasks. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (2014), pp. 177–186. doi:10.1109/ISMAR.2014.6948425.3
- [KMH01] KOSARA R., MIKSCH S., HAUSER H.: Semantic depth of field. In *IEEE Symposium on Information Visualization (INFOVIS)* (2001), pp. 97–104. doi:10.1109/INFVIS.2001.963286.3,4
- [KPV*18] KOYTEK P., PERIN C., VERMEULEN J., ANDRÉ E., CARPENDALE S.: MyBrush: Brushing and Linking with Personal Agency. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 24, 1 (2018), 605–615. doi:10.1109/TVCG.2017.2743859.2,3
- [KR14] KASAHARA S., REKIMOTO J.: JackIn: Integrating first-person view with out-of-body vision generation for human-human augmentation. In *Augmented Human International Conference (AH)* (2014), ACM, pp. 1–8. doi:10.1145/2582051.2582097.3
- [KSF10] KRUIJFF E., SWAN J. E., FEINER S.: Perceptual issues in augmented reality revisited. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (2010), pp. 3–12. doi:10.1109/ISMAR.2010.5643530.2
- [KVZ*13] KALKOFEN D., VEAS E., ZOLLMANN S., STEINBERGER M., SCHMALSTIEG D.: Adaptive ghosted views for Augmented Reality. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (2013), pp. 1–9. doi:10.1109/ISMAR.2013.6671758.3
- [LBBH10] LEE C., BONEBRAKE S., BOWMAN D. A., HÖLLERER T.: The role of latency in the validity of ar simulation. In *IEEE Virtual Reality Conference (VR)* (2010), IEEE Computer Society, pp. 11–18. doi:10.1109/VR.2010.5444820.3
- [LBBH09] LEE C., BONEBRAKE S., HOLLERER T., BOWMAN D. A.: A replication study testing the validity of ar simulation in vr for controlled experiments. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (2009), IEEE Computer Society, pp. 203–204. doi:10.1109/ISMAR.2009.5336464.3,10
- [LDG*21] LISLE L., DAVIDSON K., GITRE E. J., NORTH C., BOWMAN D. A.: Sensemaking Strategies with Immersive Space to Think. In *IEEE Virtual Reality and 3D User Interfaces (VR)* (2021), pp. 529–537. doi:10.1109/VR50410.2021.00077.2
- [LHC*21] LEE B., HU X., CORDEIL M., PROUZEAU A., JENNY B., DWYER T.: Shared Surfaces and Spaces: Collaborative Data Visualisation in a Co-located Immersive Environment. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 27, 2 (2021), 1171–1181. doi:10.1109/TVCG.2020.3030450.2,4
- [LKM*17] LAVIOLA J. J., KRUIJFF E., MCMAHAN R. P., BOWMAN D., POUPYREV I. P.: *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, 2017.2
- [LPED20] LIU J., PROUZEAU A., ENS B., DWYER T.: Design and Evaluation of Interactive Small Multiples Data Visualisation in Immersive Spaces. In *IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2020), pp. 588–597. doi:10.1109/VR46266.2020.00081.2
- [LRM*13] LEE C., RINCON G. A., MEYER G., HÖLLERER T., BOWMAN D. A.: The effects of visual realism on search tasks in mixed reality simulation. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 19, 4 (2013), 547–556. doi:10.1109/TVCG.2013.41.3,10
- [LSGB20] LANGE D., STRATMANN T. C., GRUENEFELD U., BOLL S.: Hivefive: Immersion preserving attention guidance in virtual reality. In *ACM Conference on Human Factors in Computing Systems* (2020), p. 1–13. doi:10.1145/3313831.3376803.3
- [LSS23] LEE B., SEDLMAIR M., SCHMALSTIEG D.: Design Patterns for Situated Visualization in Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 30, 1 (2023), 1–12. doi:10.1109/TVCG.2023.3327398.1,2
- [LYBP23] LIN T., YANG Y., BEYER J., PFISTER H.: Labeling Out-of-View Objects in Immersive Analytics to Support Situated Visual Searching. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 29, 3 (2023), 1831–1844. doi:10.1109/TVCG.2021.3133511.1,2,4
- [ML14] MCINTIRE J. P., LIGGETT K. K.: The (possible) utility of stereoscopic 3d displays for information visualization: The good, the bad, and the ugly. In *IEEE VIS International Workshop on 3DVis (3DVis)* (2014), pp. 1–9. doi:10.1109/3DVis.2014.7160093.2
- [MM21] MA Q., MILLET B.: Design guidelines for immersive dashboards. *Human Factors and Ergonomics Society Annual Meeting* 65, 1 (2021), 1524–1528. doi:10.1177/1071181321651177.1
- [MSD*18] MARRIOTT K., SCHREIBER F., DWYER T., KLEIN K., HENRY RICHE N., ITOH T., STUERZLINGER W., THOMAS B. H.: *Immersive Analytics*. Lecture Notes in Computer Science. Springer International Publishing, 2018. doi:10.1007/978-3-030-01388-2.2
- [MTE*20] MARQUARDT A., TREPKOWSKI C., EIBICH T. D., MAIERO J., KRUIJFF E., SCHÖNING J.: Comparing non-visual and visual guidance methods for narrow field of view augmented reality displays. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 26, 12 (2020), 3389–3401. doi:10.1109/TVCG.2020.3023605.3
- [PKB05] POLYS N. F., KIM S., BOWMAN D. A.: Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. *ACM Symposium on Virtual Reality Software and Technology (VRST)* (2005), 46–55. doi:10.1145/1101616.1101626.3
- [PLE*19] PROUZEAU A., LHUILLIER A., ENS B., WEISKOPF D., DWYER T.: Visual Link Routing in Immersive Visualisations. In *ACM International Conference on Interactive Surfaces and Spaces (ISS)* (2019), pp. 241–253. doi:10.1145/3343055.3359709.3

- [PZB*23] PIETSCHMANN L., ZÜRCHER P., BUBIK E., CHEN Z., PFISTER H., BOHNÉ T.: Quantifying the Impact of XR Visual Guidance on User Performance Using a Large-Scale Virtual Assembly Experiment. In *IEEE Visualization and Visual Analytics (VIS)* (2023), IEEE Computer Society, pp. 211–215. doi:10.1109/VIS54172.2023.00051. 2
- [RBR22] ROBERTS J. C., BUTCHER P. W., RITSOS P. D.: One view is not enough: Review of and encouragement for multiple and alternative representations in 3d and immersive visualisation. *Comput. 11* (2022), 20. doi:10.3390/COMPUTERS11020020. 2
- [Rob07] ROBERTS J. C.: State of the art: Coordinated & multiple views in exploratory visualization. In *International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV)* (2007), IEEE Computer Society, pp. 61–71. doi:10.1109/CMV.2007.20. 2, 3, 6
- [RPA*22] RICHER G., PISTER A., ABDELAAL M., FEKETE J.-D., SEDLMAIR M., WEISKOPF D.: Scalability in visualization. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2022), 1–15. doi:10.1109/TVCG.2022.3231230. 5
- [RWBH09] RAGAN E., WILKES C., BOWMAN D. A., HOLLERER T.: Simulation of augmented reality systems in purely virtual environments. In *IEEE Virtual Reality Conference (VR)* (2009), IEEE Computer Society, p. 287–288. doi:10.1109/VR.2009.4811058. 3
- [SBB*23] SHIN S., BATCH A., BUTCHER P. W. S., RITSOS P. D., ELMQVIST N.: The Reality of the Situation: A Survey of Situated Analytics. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2023), 1–19. doi:10.1109/TVCG.2023.3285546. 1, 2
- [SBDE23] SAFFO D., BATCH A., DUNNE C., ELMQVIST N.: Through Their Eyes and In Their Shoes: Providing Group Awareness During Collaboration Across Virtual Reality and Desktop Platforms. In *ACM Conference on Human Factors in Computing Systems* (2023), pp. 1–15. doi:10.1145/3544548.3581093. 2
- [SCZ*20] SIDENMARK L., CLARKE C., ZHANG X., PHU J., GELLERSEN H.: Outline Pursuits: Gaze-assisted Selection of Occluded Objects in Virtual Reality. In *ACM Conference on Human Factors in Computing Systems* (2020), pp. 1–13. doi:10.1145/3313831.3376438. 3, 10
- [SD21] SATKOWSKI M., DACHSELT R.: Investigating the impact of real-world environments on the perception of 2d visualizations in augmented reality. In *ACM Conference on Human Factors in Computing Systems* (2021). doi:10.1145/3411764.3445330. 2
- [SH16] SCHMALSTIEG D., HÖLLERER T.: *Augmented Reality - Principles and Practice*. Addison-Wesley Professional, 2016. 2
- [SL00] SUOMELA R., LEHIKONEN J.: Context Compass. In *Digest of Papers. International Symposium on Wearable Computers* (2000), IEEE Computer Society, pp. 147–147. doi:10.1109/ISWC.2000.888481. 3
- [SLC*19] SICAT R., LI J., CHOI J., CORDEIL M., JEONG W.-K., BACH B., PFISTER H.: DXR: A Toolkit for Building Immersive Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 25, 1 (2019), 715–725. doi:10.1109/TVCG.2018.2865152. 4
- [SLQS23] SAYARA A., LEE B., QUIJANO-CHAVEZ C., SEDLMAIR M.: Designing Situated Dashboards: Challenges and Opportunities. In *IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)* (2023), pp. 97–102. doi:10.1109/ISMAR-Adjunct60411.2023.00028. 1
- [SS22] SERAJI M. R., STUERZLINGER W.: XVCollab: An Immersive Analytics Tool for Asymmetric Collaboration across the Virtuality Spectrum. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)* (2022), pp. 146–154. doi:10.1109/ISMAR-Adjunct57072.2022.00035. 2
- [TEM*19] TREPKOWSKI C., EIBICH D., MAIERO J., MARQUARDT A., KRUIJFF E., FEINER S.: The effect of narrow field of view and information density on visual search performance in augmented reality. In *IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2019), pp. 575–584. doi:10.1109/VR.2019.8798312. 3
- [TK06] TONNIS M., KLINKER G.: Effective control of a car driver’s attention for visual and acoustic guidance towards the direction of imminent dangers. In *IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR)* (2006), IEEE Computer Society, pp. 13–22. doi:10.1109/ISMAR.2006.297789. 3
- [TMB02] TVERSKY B., MORRISON J. B., BETRANCOURT M.: Animation: Can it facilitate? *Int. J. Hum.-Comput. Stud.* 57, 4 (2002), 247–262. doi:10.1006/ijhc.2002.1017. 9
- [TOK*16] TATZGERN M., ORSO V., KALKOFEN D., JACUCCI G., GAMBERINI L., SCHMALSTIEG D.: Adaptive information density for augmented reality displays. In *IEEE Virtual Reality Conference (VR)* (2016), pp. 83–92. doi:10.1109/VR.2016.7504691. 3, 10
- [TWD*18] THOMAS B. H., WELCH G. F., DRAGICEVIC P., ELMQVIST N., IRANI P., JANSEN Y., SCHMALSTIEG D., TABARD A., EL-SAYED N. A. M., SMITH R. T., WILLETT W.: Situated analytics. In *Immersive Analytics*, Marriott K., Schreiber F., Dwyer T., Klein K., Riche N. H., Itoh T., Stuerzlinger W., Thomas B. H., (Eds.). Springer International Publishing, 2018, pp. 185–220. doi:10.1007/978-3-030-01388-2_7. 1
- [Ver02] VERTEGAAL R.: Designing attentive interfaces. In *Symposium on Eye Tracking Research & Applications* (2002), ACM, p. 23–30. doi:10.1145/507072.507077. 2
- [VMFS11] VEAS E. E., MENDEZ E., FEINER S. K., SCHMALSTIEG D.: Directing attention and influencing memory with visual saliency modulation. In *ACM Conference on Human Factors in Computing Systems* (2011), pp. 1471–1480. doi:10.1145/1978942.1979158. 3
- [WF09] WHITE S., FEINER S.: SiteLens: Situated visualization techniques for urban site visits. In *ACM Conference on Human Factors in Computing Systems* (2009), pp. 1117–1120. doi:10.1145/1518701.1518871. 1, 2
- [WGRF22] WIELAND J., GARCIA R. C. H., REITERER H., FEUCHTNER T.: Arrow, bézier curve, or halos? – comparing 3d out-of-view object visualization techniques for handheld augmented reality. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)* (2022), pp. 797–806. doi:10.1109/ISMAR55827.2022.00098. 3, 4, 9
- [WJD17] WILLETT W., JANSEN Y., DRAGICEVIC P.: Embedded Data Representations. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 23, 1 (2017), 461–470. doi:10.1109/TVCG.2016.2598608. 1, 2
- [WSS20] WHITLOCK M., SMART S., SZAFIR D. A.: Graphical perception for immersive analytics. In *IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2020), pp. 616–625. doi:10.1109/VR46266.2020.00084. 3
- [WVEG11] WOLFE J. M., VÔ M. L.-H., EVANS K. K., GREENE M. R.: Visual search in scenes involves selective and nonselective pathways. *Trends in Cognitive Sciences* 15, 2 (2011), 77–84. doi:https://doi.org/10.1016/j.tics.2010.12.001. 2
- [YLF*20] YU D., LIANG H.-N., FAN K., ZHANG H., FLEMING C., PAPANGELIS K.: Design and evaluation of visualization techniques of off-screen and occluded targets in virtual reality environments. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 26, 9 (2020), 2762–2774. doi:10.1109/TVCG.2019.2905580. 3, 4, 9
- [YLS24] YU X., LEE B., SEDLMAIR M.: Design Space of Visual Feedforward And Corrective Feedback in XR-Based Motion Guidance Systems, 2024. arXiv:2402.09182, doi:10.1145/3613904.3642143. 2
- [ZLG*21] ZOLLMANN S., LANGLLOTZ T., GRASSET R., LO W. H., MORI S., REGENBRECHT H.: Visualization techniques in augmented reality: A taxonomy, methods and patterns. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 27, 9 (2021), 3808–3825. doi:10.1109/TVCG.2020.2986247. 2