

Towards Understanding Diminished Reality

Yi Fei Cheng
Department of Computer Science,
Swarthmore College
United States

Hang Yin
School of Computer Science, Carnegie
Mellon University
United States

Yukang Yan
Department of Computer Science and
Technology, Tsinghua University
China

Jan Gugenheimer
Télécom Paris - LTCI, Institut
Polytechnique de Paris
France

David Lindlbauer
Human-Computer Interaction
Institute, Carnegie Mellon University
United States




Figure 1: We present two empirical studies on diminished reality (DR). This figure demonstrates two applications: (A) using DR to visually remove work-space clutter, (B) using DR to support search by distinguishing between targets and non-targets.

ABSTRACT

Diminished reality (DR) refers to the concept of removing content from a user's visual environment. While its implementation is becoming feasible, it is still unclear how users perceive and interact in DR-enabled environments and what applications it benefits. To address this challenge, we first conduct a formative study to compare user perceptions of DR and mediated reality effects (e. g., changing the color or size of target elements) in four example scenarios. Participants preferred removing objects through opacity reduction (i. e., the standard DR implementation) and appreciated mechanisms for maintaining a contextual understanding of diminished items (e. g., outlining). In a second study, we explore the user experience of performing tasks within DR-enabled environments. Participants selected which objects to diminish and the magnitude of the effects when performing two separate tasks (video viewing, assembly). Participants were comfortable with decreased contextual understanding, particularly for less mobile tasks. Based on the results, we define guidelines for creating general DR-enabled environments.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9157-3/22/04...\$15.00

<https://doi.org/10.1145/3491102.3517452>

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI; Mixed / augmented reality.**

KEYWORDS

Diminished Reality, Mediated Reality, Empirical study

ACM Reference Format:

Yi Fei Cheng, Hang Yin, Yukang Yan, Jan Gugenheimer, and David Lindlbauer. 2022. Towards Understanding Diminished Reality. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3491102.3517452>

1 INTRODUCTION

Augmented Reality (AR) enables users to personalize the visual appearance of their surroundings by directly introducing digital augmentations into their visual field. For example, users could embed virtual applications like web browsers and email clients into their environment (cf. Cheng et al. [6]). This promises a rich set of experiences involving a seamless blend between the virtual and real worlds, and continuous availability of digital information.

However, placing additional virtual elements into users' field of view has the potential to overload them with information (cf. Lindlbauer et al. [29]). This is particularly problematic considering that the physical world, on which virtual elements are overlaid, is already fraught with its own sources of information overload and distraction. One class of augmentations that aims to alleviate this challenge is referred to as *diminished reality* (DR). DR was first introduced by Steve Mann as part of his Mediated Reality framework,

which broadly refers to artificially modifying the human perception by way of devices [34]. DR addresses the problem of information overload by deliberately removing virtual and real-world content from users' perceived environment.

DR has been explored in application domains such as interior design [53], manufacturing [45], and gaming [50]. Most research, however, focused on its technical implementation (e. g., the realization of the technique on AR headsets). While DR becomes increasingly feasible to implement (see Mori et al. [41] for an overview), it is still unclear how users will perceive and interact in environments with DR augmentations enabled.

In this work, we aim towards building an in-depth understanding of DR, including users' perception of and considerations for its usage, and its potential applications. For this purpose, we conducted two studies.

In the first study, we investigate the question of how users perceive different augmentations which seek to remove real-world content from their visual environment. Participants ($N = 16$) were asked to alter an environment by using seven augmentation effects, like changing the color or scale of target objects, within four application scenarios. The effects and applications were inspired by prior AR research, theories and frameworks from perception, and image processing techniques. To avoid current technical limitations of AR HMDs, for both our experiments, we simulated the DR experiences using a VR HMD and 3D scanned environments. This is a frequently applied methodology in AR research to overcome issues like limited visual field and occlusion, which would potentially impact conceptual findings [6, 29]. This approach allowed us to ensure that our participants were always presented with the same environments, increasing the overall internal validity of our experiments. We also benefited from having full control over the environment. Results revealed a generally positive attitude towards DR, as well as a scenario-invariant preference for the opacity adjustment effect as a means of achieving DR. Furthermore, participants reported that the selection of target objects for DR effects should be personalized, and dependent on various contextual factors such as item proximity, relevance, and task mobility.

In the second experiment, we investigate how participants apply DR effects when performing different tasks to maximize their level of perceived comfort. Participants chose which objects to diminish with either the opacity adjustment or outline effect, and the magnitude of the effect. We compare this custom condition to performing the same tasks in an environment where all task-irrelevant objects are diminished, and an environment where no objects are diminished. Additionally, we analyzed their modifications in the custom condition to identify patterns in DR usage. Participants modified an average of 29.7% ($SD = 20.0\%$) of objects using DR in each environment, and applied DR to non-task relevant objects in order to support their task completion. Participants avoided diminishing items that could physically interfere with their movements (i. e., participants generally avoided DR which would result in unintentional collisions) and applied less DR when there was a social presence.

Based on the results, we distill six recommendations for future implementations of DR. We recommend the use of opacity adjustment as a general purpose DR augmentation effect. We also recommend for DR to always be complemented with mechanisms for

retaining contextual understanding where reductive augmentations are applied. Our results furthermore show that the acceptability of DR depends on the likelihood of physical interference from the diminished elements, their interaction requirements and behaviours, and the level of social presence. Lastly, we advise that future devices should allow users to retain a sense of agency over the DR augmentations.

2 RELATED WORK

2.1 Diminished Reality

Diminished Reality was introduced by Mann [34, 35] as a concept in the context of Mediated Reality, later embedded in the context of All Reality [36]. DR refers to the removal ("diminishing") of real-world physical objects from users' visual perception. In the context of AR, diminishing objects is oftentimes performed to enable users to see through or remove existing objects. Avery et al. [3] and Kalkofen et al. [24], for example, both proposed AR x-ray vision-like approaches to enable users to see through walls. Taylor et al. [55] used DR to support robotic tele-manipulation. Barnum et al. [5] enabled users to see moving objects through static ones, such as a moving pedestrian behind a corner through view-port synthesis. Rameau et al. [46] used DR to visually remove other cars from a drivers' visual field to improve the safety of advanced driving assistance systems when overtaking. Kim et al. [27] applied DR to remove content-irrelevant real objects by modifying the transparency of the object. Besides the context of x-ray vision and driver safety, DR has been used in the context of sports [51], for interior design [53], to visualize data from autonomous vehicles [17], to enhance remote conversations [15], and to augment the attention of individuals with autism [58]. Pearson [44] explored the concept of chameleon devices to increase privacy and security around mobile devices. Haas et al. [13] explored how the idea of diminishing can be used in auditory Mediated Reality. The aforementioned works serve as inspiration for the usage scenarios and DR effects employed in our work. Furthermore, most of these works focus on alterations of a target objects' opacity to achieve DR. We expand on this approach and address the question of which DR effects users prefer, and under what circumstances.

2.1.1 Implementing DR. Previous work has been mostly concerned with developing and refining implementations of DR, focusing on how to fully remove target objects from images or videos while leveraging a variety of cameras and sensors. Mori et al. [41] provide an overview of the different technologies. They concluded that DR technology is used to implement diminish, see-through, replace, and in-paint functions, where "diminish" means "degrade visual functions for a certain purpose." Herling et al. [16] removed objects from a video feed using explicit object tracking and patch-based image completion, closely related to the PatchMatch [4] algorithm. Recent work furthermore employed deep learning-based methods for video inpainting (e. g., Kim et al. [26], Guida and Sra [12]). Similar methods have been extended to 3D environments by Mori et al. [39, 40]. Hasegawa and Saito [14] proposed extracting and removing pedestrians from video sequences using the histogram of oriented gradients feature descriptor. Zokai et al. [60] used multi-view projections to visually remove objects from widely separated stereo images.

Meerits and Saito [37] leverage multiple RGB-D cameras to hide objects from a scene. This is comparable to Remixed Reality [32], which leverages a live 3D reconstruction of an environment for the manual removal of objects. Kari et al. [25] proposed a pipeline for real-time object substitution for tablet-based and head-mounted MR, and addressed technical challenges of diminishing moving objects with only a real-time stream of monocular RGB information.

All these works assume a video pass-through technology that can visually erase objects from users' environments. Projection mapping can be seen as an alternative to this approach, and has been explored as a viable option to implement DR. Seo et al. [52], for example, employ radiometric compensation to visually hide objects. Iwai et al. [20] enable users to set specific document on their desk to transparent for easier access. Inami et al. [18] expands uses retro-reflective material to diminish parts of what users see.

Besides visual augmentation, Lindlbauer et al. [30] built objects with controllable transparency from optically dynamic material. Their work requires target objects to be manufactured with the specific DR-implementation in mind, and is therefore less flexible compared to visual augmentation techniques.

The aforementioned methods and algorithms work towards making the implementation of DR feasible. In this work, we focus on the concept and usage of DR. We therefore employ fully immersive and controllable environments to avoid challenges of low frame rate or imperfections in the effects. Our insights, however, are intended to directly inform the design and implementation of current and future DR techniques.

2.2 Visual alterations beyond DR

The goal to alter users' perception of target objects is shared by research beyond DR. Rixen et al. [47], for instance, presented how visual alterations affected user's experience with interpersonal communications. They measured user's comfort, acceptance of altering and being altered, and how it is impacted by being able to perceive or not perceive the alteration. We have taken a similar approach in our experiment, but focus on understanding the concept of diminished reality rather than understanding the impact of visual alterations in a conversation. Jones et al. [22, 23] use projection mapping to visually alter the environment that users are in. They carried out a user study that involved performing tasks and rating the effects. Our second user study (Section 5) was done in a similar manner although the technologies involved are different. We refer readers to the work of Grundhöfer and Iwai [11] for an overview of projection mapping systems and how the technique is used to augment users' visual reality. Lindlbauer et al. [31] visually changed the surrounding of target objects to alter users' perception of the target. One of their applications changed the perceived visibility of objects, which embodies the idea of DR. In our work, we investigate the influence on different effects on users' perception and comfort, which can be used to inform the design of above systems.

Besides controlling users' visual environment, researchers have explored predicting users' gaze behavior and saliency (cf. Itti et al. [19]), including modulating the saliency in an AR setting [38, 56]. We adopt saliency modulation as one possible effect in the first experiment. Our scenarios were furthermore inspired by work on visual clutter (cf. Rosenholtz et al. [49]) and visual search (cf. Wolfe

and Horowitz [57]). The findings of the second experiment can inform the design of techniques to facilitate both.

We believe DR is potentially a valuable addition to the toolset of research on visual augmentation. Our work aims at providing the foundation to understanding when and to what extent DR should be applied to provide a beneficial experience for end users.

3 DIMINISHED REALITY

In the following, we provide a functional definition of DR to set our work in the context of previous approaches. We base our definition on Haas et al.'s framework for interactive auditory Mediated Reality [13], which is similar to visual DR, but for the auditory domain. We furthermore provide a set of application scenarios which we used to inform the use cases employed in our experiments. Note that we are not concerned with the effective implementation of DR, but assume complete scene understanding and an output device that is capable of arbitrarily altering the physical world. We chose these parameters to enable a less constrained exploration of DR.

3.1 Functional definition

With a wearable device situated in front of a user's field of view (i. e., an AR headset), incoming visual stimuli can be segmented by source. Based on the user's objectives and requirements, a target source or their respective features can be obfuscated to various extents. The result of the DR process is an altered and curated environment whereby targeted aspects of the real environment are rendered less perceptible. To remain consistent with past work, we will primarily use the term DR to refer to achieving perceptual reduction of entire objects via opacity adjustment. In past work, DR most often targeted individual objects rather than their features. Additionally, past implementations primarily focused on either decreasing the opacity of objects or removing them entirely. We note that DR can be applied on a more granular level and that there are a variety of ways beyond opacity adjustment for achieving this purpose. For instance, to reduce the saliency of a target object's color, a desaturation effect could be applied, whereas if the objective was to obfuscate a target object's texture, a blur effect may be more appropriate. We will more broadly classify such alternative approaches as Mediated Reality techniques. In our first experiment, we compare users' perception of the opacity adjustment to several Mediated Reality techniques.

With respect to the DR process, the user can adopt several roles. We distinguish between the moderator of a DR instance, who defines the DR target, effect, and extent, and the recipient, whose visual perception is augmented through DR. The user roles are not mutually exclusive. In a single user scenario, where a user defines a specific DR augmentation on themselves, they occupy both. We make this distinction to explore cases where users may apply DR on others.

We see this definition of DR as a functional starting point for our work. For instance, our current understanding of DR excludes implementations where objects are rendered less perceptible through introducing modifications to their surroundings [31]. Additionally, effects designed to obfuscate certain features (e. g., desaturation for color) may inadvertently make the target item more salient to a user overall (e. g., a desaturated object in a colorful environment).




Figure 2: Depictions of DR usage scenarios: (S1) clutter management, (S2) search, (S3) privacy protection, (S4) remote guidance. For each scenario, the unmodified (left) and DR-enabled (right) views are illustrated. Note that in (S2) and (S4), we use red boxes to emphasize objects that the user is focusing on, they are not a part of our effect. Descriptions of each scenario are provided in section 3.2.

We believe that our work grounds such future exploration of a more diverse set of combinations of DR effects and their combination with an existing environment.

3.2 Scenarios for DR

To explore user preferences for applying DR, we provide four sample DR usage scenarios (Figure 2). The applications were selected based on prior work on DR (e.g. [3, 27, 35, 44]) and to span different aspects of the *user role* and *target* factors defined in Section 3.1:

- **S1. Work-space clutter management.** An office clerk is typing up a report at their work station. They have not organized their desk in a while, and as such their space has

become rather cluttered with items unrelated to their current task. The DR function on their AR headset enables them to temporarily hide the task-irrelevant items in their periphery to produce an illusion of a cleaner and more comfortable working environment.

- **S2. Search.** A student is searching for a book in the library. Their AR headset automatically recognizes which books are relevant for their task. Through DR, the system assists the user in their search by filtering out irrelevant books from their field of view. Thus, with DR, they can perform their search in a more efficient manner, while previously they would have to scan through the shelves sequentially. A DR approach may be preferable to avoid information overload compared to introducing visual indicators (e.g., arrows or highlighting bounding boxes) if the environment is cluttered to begin with.
- **S3. Privacy protection.** A person is browsing through social media on their mobile phone in a coffee shop. They would prefer to keep their phone use private. Through DR, they obfuscate the contents of their device from the view of others.
- **S4. Remote guidance.** An instructor is guiding a student through a circuit assembly task. They are given remote access to their students' AR headsets. To guide them through the problem step-by-step, they filter out the irrelevant items in students' field of view. By doing so, they direct students towards the relevant circuit components for the current step, while avoiding information overload, similar to S2 (search).

The presented scenarios assume future AR headsets that are capable of recognizing users' context and perform DR perfectly in real-time. In S1 (clutter) and S2 (search), the user is both the moderator and recipient of the DR, whereas S3 (privacy) and S4 (guidance) explore use cases where users apply DR on others. S1 (clutter) and S3 (privacy) present scenarios where the objective of the DR is to diminish a selected target. The use of DR in S2 (search) and S4 (guidance) is to highlight unmodified targets by diminishing non-targets. In all scenarios, different approaches to removing aspects of the environment (i.e., visual effects) may be more appropriate depending on the context. Similarly, the ideal extent to which the target is diminished might vary. For instance, in S4 (guidance), hiding the irrelevant items entirely from the student may potentially be dangerous (e.g., rendering a sharp object as invisible), but in S3 (privacy), for protecting a user's privacy, just obfuscating the texture and color of the mobile screen may be sufficient.

4 EXPERIMENT I

We conducted a formative study that compares the standard DR implementation (opacity change) with six Mediated Reality alternatives. 16 participants were asked to report on their experiences in 3D scanned environments in VR.

4.1 Design

We used a within-subject design with two independent variables: *effect* with seven levels (Section 4.2) and *scenario* with four levels (Section 3.2). Each participant experienced every effect in every

scenario. The *environment* in which each scenario is instantiated is set as a control factor. The scenario order was counterbalanced using a Latin Square.

We instantiated the four scenarios in three environments each (Figure 4). The environments were sampled to represent typical settings for each usage scenarios, with manually defined DR targets according to the scenario description. The effect order within each scenario is set randomly. In each scenario, participants can toggle between each effect freely and adjust the *extent* to which it is applied to the pre-defined *targets* within the scene.

4.2 Effects

We compare the standard implementation of DR to a selection of Mediated Reality-oriented effects informed by prior research in DR and AR, perception, and image processing. The selection of effects diminish different features of a given object (e. g., color, texture, geometry), and are illustrated in Figure 3.

- **Reduce opacity (E1).** The standard DR implementation approach involves recovering the occluded background of a target object [60], which we consider equivalent to rendering the target object as fully transparent. A natural extension is semi-transparent rendering for partial diminishing of target object.
- **Outline (E2).** E2 extends E1 by rendering an outline around a target to preserve some of the geometric context. The technique is explored by Taylor et al. [55].
- **Blur (E3).** We introduce blur as a means of making regions less perceptible. Prior research has shown that modulating the sharpness of portions of an image is an effective mechanism for guiding users' attention [28, 59].
- **Reduce saliency (E4).** Mendez et al. [38] proposed modulating the saliency of image regions to direct user attention. We adapt their approach by applying only conspicuity-reducing modulations onto the DR targets.
- **Desaturate (E5).** The effect reduces users' understanding of the target's coloring. Previously, the effect was compared with E4 in the work of Mendez et al. [38] as a baseline method for diminishing.
- **Reduce contrast (E6).** The effect reduces users' understanding of both the target's coloring and texture, and was used as baseline in the work of Mendez et al. [38].
- **Reduce scale (E7).** We decrease the apparent size of targets as a means of decreasing its perceived presence, since the resulting object will literally occupy less of users' field of view. Additionally, this reduces users' understanding of the object's true geometry. This effect is inspired by prior works that changed objects' apparent size [31, 42].

We consider the opacity effect (E1) as equivalent to the standard DR implementation. We classify E2-E7 as Mediated Reality techniques that aim at reducing the saliency and conspicuity of target with respect to the environment. The outline effect (E2) augments the opacity effect (E1) with additional visual feedback of the object geometry. Rather than simply removing objects from a users' visual perception, E3-E6 modifies their texture and E7 changes their scale.




Figure 3: Illustrations of sample augmentation effects evaluated in our first experiment: (E1) reduce opacity, (E2) reduce opacity + outline, (E3) blur, (E4) reduce saliency, (E5) desaturate, (E6) reduce contrast, (E7) reduce scale. The top-left image illustrates the unmodified view of the environment. The remaining images demonstrate each effect applied to the painting.

4.3 Apparatus

The DR experiences were developed in Unity 2021 for the Oculus Quest headset. The environments were shown in VR as 3D scans (Canvas [43]). We additionally populated our environments with scanned objects (Qlone [33]) and people from the Microsoft Rocketchbox repository [9]. All sessions were video recorded (SideQuest [10], Oculus Casting [8]).

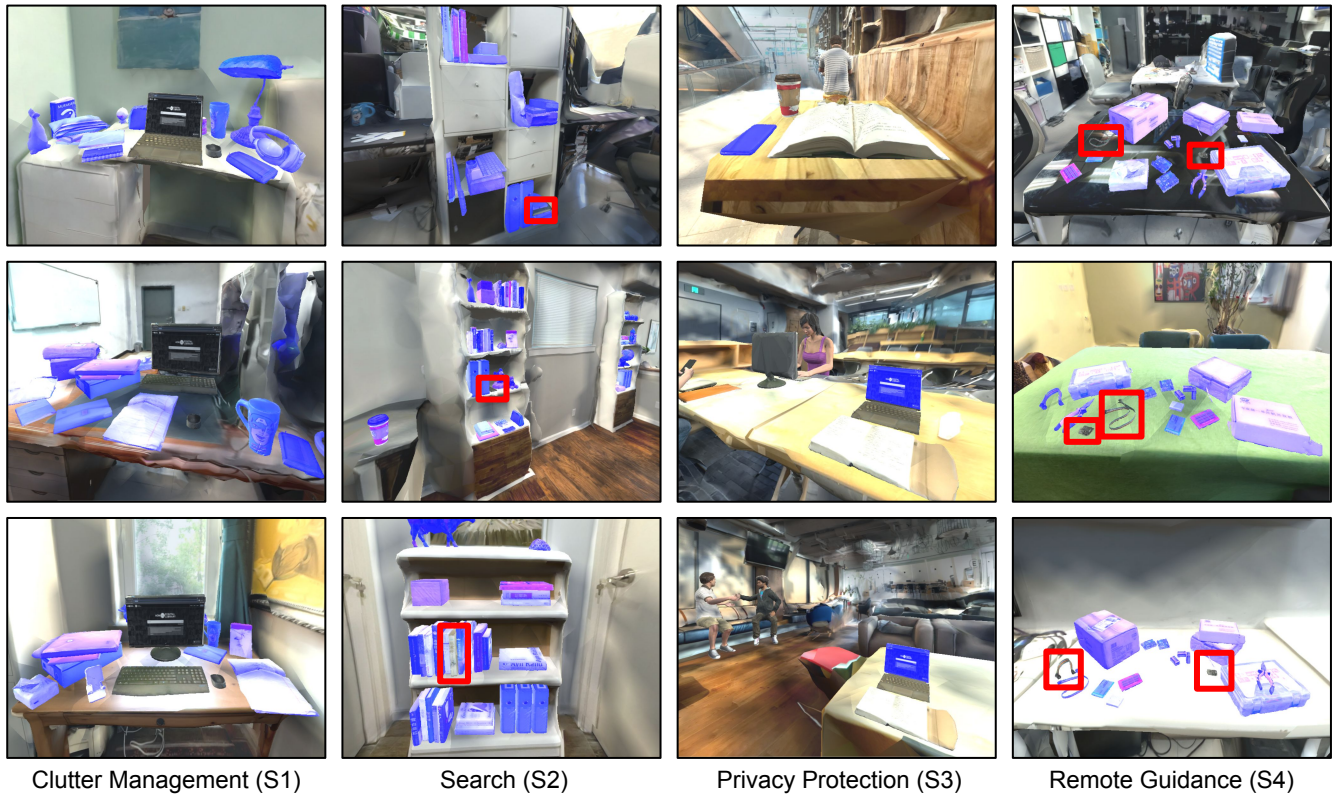


Figure 4: Illustrations of the environments tested in experiment I: (S1) clutter management, (S2) search, (S3) privacy protection, (S4) remote guidance. We highlight the objects which were pre-selected to be diminished (blue). In (S2) and (S4), we also draw red bounding boxes around the items which we asked the participant to find during the study, they are not part of our effects.

4.4 Participants

We recruited 16 participants (10 male, 6 female; age: $M = 21.94$, $SD = 2.77$) via snowball sampling starting from university mailing lists and social networks. The study was conducted remotely through Zoom due to COVID-19. Participants needed to have access to an Oculus Quest. We asked participants to install our application before the study. A study session took around 60 minutes. Participants had varying amounts of experience with AR ($M = 3.00$, $SD = 1.37$) and VR ($M = 3.24$, $SD = 1.30$) on a Likert-type scale from 1 (low) to 5 (high). Participants were compensated with \$10.

4.5 Procedure

After completing the consent form and demographic questionnaire, participants performed a training session. The session prompted the participant to toggle between different DR effects in a test environment, and apply the effects to various extents. The experimenter subsequently introduced the scenarios and environments, and participants reported on their experience with each DR effect via a think-aloud format. For each scenario, participants were also asked to complete a questionnaire with Likert ratings on preference, comfort, and perceived utility. The session ended with a semi-structured interview, including general impressions of DR and opinions on the various effects and usage scenarios.

4.6 Results

Participants generally had a positive attitude towards the concept of DR and preferred the opacity (E1) and outline (E2) effects. We coded experiment session transcriptions to identify key findings, and cross referenced our analysis with the questionnaire results (Figure 5). We performed our statistical analysis using JASP 0.14.1 [21].

4.6.1 Overall Perception of DR. Participants were generally positive towards the idea of using DR across our four example scenarios (S1 (clutter): $M = 4.25$, $SD = 0.77$; S2 (search): $M = 4.44$, $SD = 0.63$; S3 (privacy): $M = 4.25$, $SD = 0.58$; S4 (guidance): $M = 4.50$, $SD = 0.63$). They particularly appreciated the use of DR to distinguish between relevant and irrelevant environment objects, seeing benefits of this feature to support search and sustain attention (e.g., "I like how it will make me focus on a single thing.", P4). Several participants also mentioned the benefits of DR in protecting their privacy (e.g., "it's helpful outdoors because I'm super self-conscious about people noticing what videos I'm watching", P11).

Participants were skeptical about using DR on others, since they saw increased potential for abuse ($N=14$). They were specifically concerned with malicious parties hiding dangerous objects or nudging users to make decisions against their self-interest. When applying DR on others, participants stressed the importance of consent ($N = 9$), and that the recipient must be aware of ($N = 8$) and able

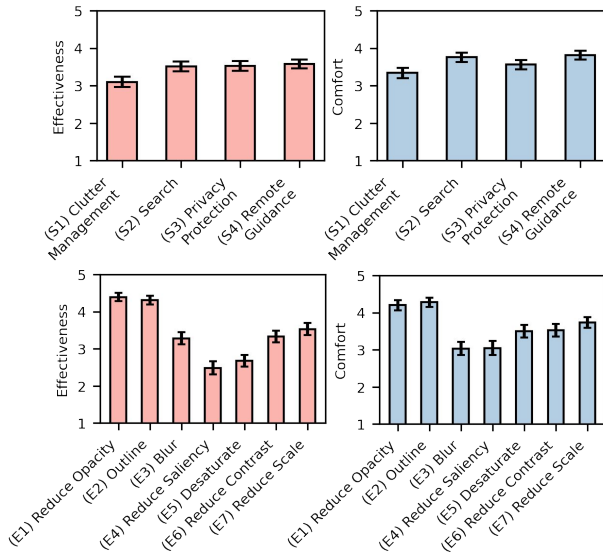


Figure 5: Summary of the study I results. Top: Participant ratings of comfort and effectiveness of DR usage by scenario. Bottom: Participant ratings of comfort and effectiveness of DR usage by effect. The error bars represent the standard error for each value.

to retain agency over its application ($N = 3$). Participants were also concerned with DR causing them to unintentionally collide with objects ($N = 14$). Lastly, they preferred the DR effects to be kept to a minimum for task involving movement ($N = 7$) and that objects that are tangentially or soon-to-be relevant to stay visually accessible ($N = 13$).

4.6.2 Effect preferences. Participants evaluated each effect from the perspectives of effectiveness and comfort. A Friedman’s test showed a main effect of the DR effect on participant comfort ($\chi^2(6) = 60.055, p < 0.001$) and perceived effectiveness ($\chi^2(6) = 110.239, p < 0.001$). We performed a series of pairwise Wilcoxon signed rank tests with Bonferroni adjustments for post hoc analysis.

Participants considered both opacity adjustment (E1) and outline (E2) to be more effective than the remaining effects (E3-E7) (all $p < 0.05$). In terms of comfort, they preferred opacity adjustment over blur (E3) ($W = 882.5, p < 0.001$), saliency modulation (E4) ($W = 1040, p = 0.002$), and scale (E7) ($W = 248, p = 0.014$). They preferred outline (E2) over every other effect (all $p < 0.05$) aside from opacity adjustment (E1). They appreciated outline (E2) for providing them with an awareness of what was diminished. Overall, the results indicate a general preference for the opacity adjustment (E1) and outline (E2) effects, indicating the relevance of the standard implementation of DR.

A Friedman’s test also indicated that participant perceptions of DR effects depended on the scenario (comfort: $\chi^2(6) = 12.356, p = 0.006$; effectiveness: $\chi^2(6) = 8.489, p < 0.037$). A series of pairwise Wilcoxon signed rank tests with Bonferroni adjustments revealed that participants found the use of DR least comfortable and effective for clutter management (S1) (comfort: $W = 1711,$

$p = 0.042$ compared to search (S2), $W = 1705.5, p = 0.005$ compared to remote guidance (S4); effectiveness: $p < 0.05$ compared to S2-S4). We attribute this comparatively negative perception of DR use for clutter management (S1) to a misalignment of what participants considered clutter and what was pre-set to be diminished. For DR to be effective, it must remove what participants actually want it to remove ($N = 8$). Overall, participants agreed that how DR should be applied depends on the scenario, environment, and appearance of objects ($N = 16$). For instance, for privacy preservation (S3), participants were just concerned with obfuscating the contents of their device, and therefore the blur (E3) and contrast reduction effects (E6) was most appropriate in that context ($N = 11$). Similarly, applying the de-saturation effect on a monochromatic object would be redundant ($N = 13$).

5 EXPERIMENT II

In our second experiment, we explore the optimal level of DR for various contexts. To achieve this, we gave users full control of what objects in their environment should be diminished and to what extent. 12 participants performed two tasks in two environments, each time with three DR conditions: (1) no DR applied, (2) DR applied to all task-irrelevant objects, and (3) user-driven custom DR. The evaluated conditions are shown in Figure 6. As a motivating scenario, participants were instructed to manage visual clutter.

5.1 Design

We used a within-subject design with three independent variables: *task* (blocks, video), *environment* (apartment, office), and *DR condition* (no DR, full DR, custom DR). Each participant performed every task three times, once under each DR condition, in every environment. Based on the results of experiment I, we use the transparency effect (E1) for DR and enable participants to toggle an outline (E2) around targets. The order of the task and environment were fully permuted between participants. We additionally counterbalanced the order of the DR condition using a Latin Square.

5.2 Tasks

We evaluate user preferences for how DR should be applied in the context of two tasks (blocks, video). We decided to evaluate the use of DR in the context of a selected primary task as this would better reflect the typical usage where applying DR is not the focus of attention.

For *blocks*, participants were asked to construct a block structure based on a template photo. The task was limited to one minute, and participants were instructed to replicate the block to the best of their ability. Grabbing and assembling blocks was implemented using standard VR interaction techniques. We chose this task as an instance of a moderate complexity task that required participants’ full attention. Additionally, the task is interactive and required participants to physically engage with their environment (e.g., avoid collisions with surrounding objects while building).

As second task (*video*), participants watched a 1-minute video clip while monitoring two peripheral application windows for notifications. Participants were asked to report notification they observed in a timely manner. After the video, participants had to provide a two sentence summary of what they saw. The task was designed




Figure 6: An overview of the study II tasks, environments, and DR conditions.

to resemble a typical multi-display AR interface that is passive and of lower cognitive load.

5.3 Environment

Participants performed tasks in two different environments, *apartment* and *office*, selected to represent a private and public setting, respectively. Environments were set up so that each individual item within the room could be selected and diminished, except the room structure (i. e., walls, ceiling, floor). The office included other people, which participants were also able to diminish. We hypothesized that participants would treat people differently from other scene elements, and take their presence into consideration when deciding on appropriate usage of DR. We additionally included dynamic elements (screen with videos playing) that were likely to draw participants' attention within both environments.

5.4 DR Condition

Participants performed the tasks in three DR conditions: (1) with no DR applied (no DR), (2) with DR applied in full (i. e., fully transparent with no outline) to all immediately task irrelevant objects (full DR), and (3) where the use of DR is user-driven (custom DR). In the *custom DR* condition, participants are given the option to apply DR to any object within their environment. Participants controlled the extent of the DR applied to each object according to their preferences.

5.5 Apparatus

We simulated the experience of performing tasks in different environments augmented with different DR conditions in VR, developed in Unity 2021 for the Oculus Quest headset. We use environments from the Replica dataset [54] as a starting point instead of scanned environments to enable greater customization within our environments and since they exhibited fewer visual artifacts. We first processed the environments using MeshLab [7] and MeshMixer [2] to clear the environments of all objects. We then populated the environments with items obtained from different sources: 3D scanned (i. e., using Qlone [33]), the Microsoft Rocketbox Repository [9], and the Amazon Berkeley Object dataset [1]. In order to ensure that participants registered collisions with (diminished) objects, the VR controllers vibrated whenever their hand collided with an object in VR. All sessions were video recorded.

5.6 Participants

We recruited 12 participants (8 male, 4 female; age: $M = 22.67$, $SD = 3.23$) from a university. The study was conducted remotely through Zoom. Participants needed to have access to an Oculus Quest and had to download and install our application before the study. A study session took around 60 minutes in total. Participants were generally familiar with both AR ($M = 3.50$, $SD = 1.00$) and VR ($M = 4.17$, $SD = 0.72$) on a Likert-type scale from 1 (low) to 5 (high). Participants were compensated with \$10 for their efforts.

5.7 Procedure

After completing the consent form and demographic questionnaire, participants performed a training session to get familiar with the experiment. The training session demonstrated the three DR conditions, and required participants to make adjustments in the custom DR condition. Participants subsequently performed all tasks in all conditions employing a think-aloud protocol. For conditions with *custom DR*, participants were asked to maximize their perceived comfort. The session ended with a questionnaire on comfort and perceive performance, and a structured interview.

5.8 Results

We analyzed the questionnaire ratings with JASP 0.14.1 [21] and performed a thematic analysis of the think-aloud and interview recording transcriptions. We additionally analyzed participants' custom DR usage in terms of usage distribution by task and environment, spatial relationships, and diminished objects. In summary, we found that participants preferred the custom DR condition over no DR and full DR for the given tasks. On average, they applied DR to 29.7% of objects in the target environment (of 83 modifiable objects in the private environment; 77 in the public environment). The environment and task did not significantly influence the number of modified objects.

5.8.1 Perceptions of DR conditions. Participants' ratings are shown in Figure 7. A Friedman test showed a main effect of the DR condition on participant comfort ($\chi^2(2) = 12.181, p = 0.002$) and perceived performance ($\chi^2(2) = 20.074, p < 0.001$). We performed a series of pairwise Wilcoxon signed rank tests with Bonferroni adjustments for post hoc analysis. Results showed that the custom DR condition was preferred in terms of perceived performance and comfort compared to both the no DR and full DR conditions ($p < 0.05$). This result further supports the idea that users are open to applying DR to create a more comfortable, personalized visual environment, if given the capabilities of doing so easily.

We performed a series of Wilcoxon signed rank tests to analyse differences between the tasks and scenarios across the DR conditions. Within the full DR condition, participants reported to perform better and feel more comfortable completing the video task compared to the blocks task (performance: $W = 72, p = 0.01$; comfort: $W = 136, p < 0.001$). We primarily attribute participants' preference for the full DR condition in context of the video task compared to the blocks task to their concern relating to unintentionally colliding with objects. The blocks task generally required more movement and direct interaction with the environment than the video task, thus a higher level of contextual understanding. It follows that diminishing the surrounding objects is less acceptable in the blocks task than within the video task, where no movement was required whatsoever. No additional significant differences in performance and comfort ratings were identified.

5.8.2 Custom DR usage. Across the environments and tasks, on average, participants applied diminished reality to 29.7% ($SD = 20\%$) of the modifiable items in their surroundings. Mann-Whitney U tests indicated that neither the task ($W = 207.5, p = 0.099$) nor the environment ($W = 350.5, p = 0.201$) yielded a main effect on the number of modified objects. As evident in the standard deviation,




Figure 7: Summary of the study II questionnaire results. Top: Overall participant comfort and perceived performance ratings of the three DR conditions. Bottom: Participant comfort and perceived performance ratings by DR condition and task. * indicates a statistically significant difference. Error bars indicate standard error.

participants varied in how they customized their environments with DR. Figure 8 offers a break down of all modifications, including the amount of objects that participants set as fully transparent, fully transparent with outline enabled, semi-transparent, and semi-transparent with outline enabled.

Participants set objects as fully transparent as a primary method of DR (Figures 8 and 9). Participant use of DR was driven by four main considerations: the spatial arrangement of objects ($N = 12$), movement ($N = 8$), object interaction requirements and functionality ($N = 10$), and social presence ($N = 8$).




Figure 8: Average distribution of DR usage by task and environment. The sum of each bar represents the total rate of modified items for each condition (i. e., the sum of all effects). Error bars represent the standard error of the sum of all modified items per condition.

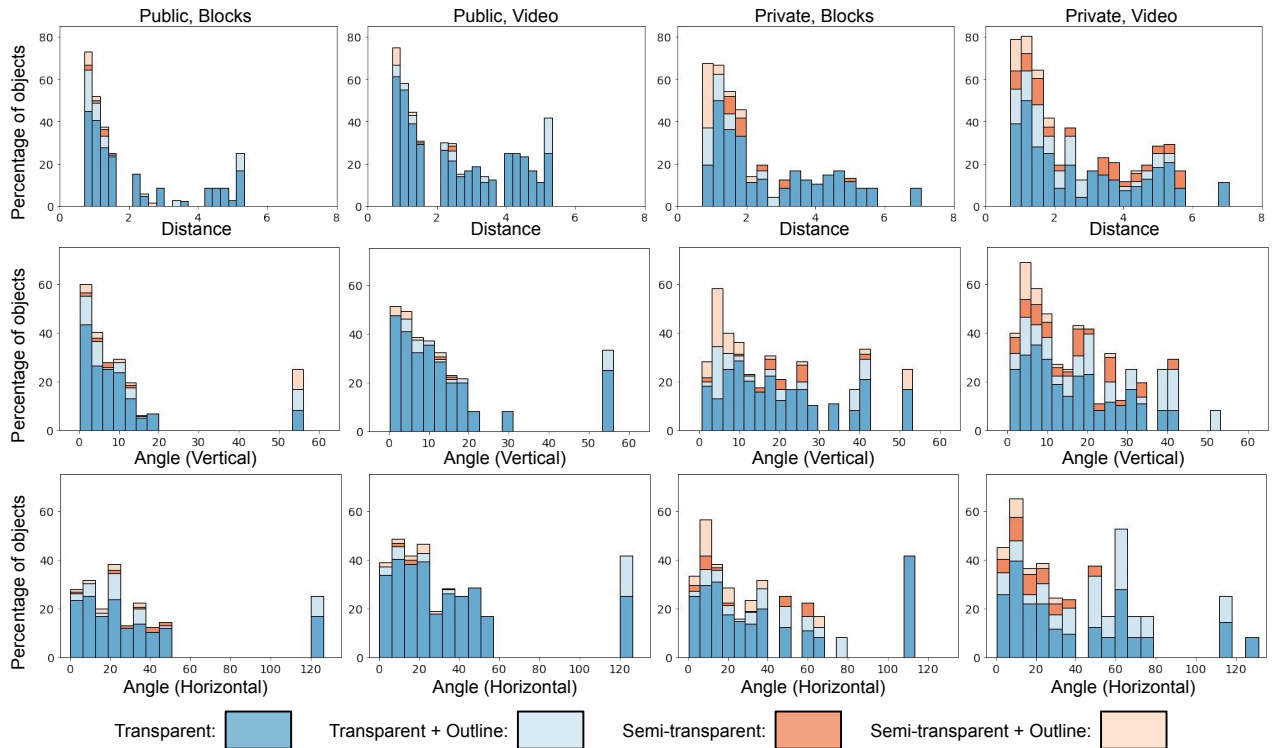


Figure 9: Distribution of DR usage by object distance and angle from the participant. We normalize our values by the total number of modifiable objects in each bin (e. g., the first dark-blue bar in the top-left graph indicates that around 45% of the modifiable objects $[0.71, 0.94]$ units away from the user were set as transparent). From left to right: DR usage in our four conditions.

Spatial arrangement. Figure 9 illustrates how participants' decisions were often driven by the spatial arrangement of objects within their surroundings. Participants visibly focused their augmentations on objects in close proximity (i. e., the percentage of objects they modify appears to decrease with distance) and within their field of view (i. e., modified objects are generally within 60°

from participants' line of sight). Several participants noted that they did not feel inclined to alter the environment objects at a distance or outside their peripheral vision since the objects did not serve as a significant distraction to their completion of the assigned tasks ("I left the other objects unmodified because I feel like my focus will just be fixated on the objects immediately in front of me", P2).

Participants also indicated that they were weary of diminishing all background objects as doing so would erase all visual indicators of their original context. Participants found that the erasure of visual indicators yielded an uncomfortable otherworldly space (*"It makes me feel like I'm in a liminal space"*, P6).

While participants generally saw greater utility in applying DR on objects close by, they were also cautious about using it on objects within their reach ($N = 12$). Echoing our findings from experiment I, participants were primarily concerned about how applying DR to objects within their reach could result in unintended collisions or interfere with their task completion ($N = 12$). Figure 10 illustrates how this concern manifested in practice. Evaluating the usage of DR for the blocks task in tandem with its use for the video task, observably fewer objects directly in front of the participant were set as fully transparent for the former (i. e., notebook, wallet, phone, coffee cup, scone, and napkins in the private environment; mouse, cup, book, phone in the public environment). Participants were generally more concerned with avoiding collisions with objects within their reach in the blocks task compared to the video task because it required them to move around and interact more with their environment (*"In the blocks tasks, I think it's easier if I know if there are actual objects on the table because of collisions. For the video tasks, I won't have to really interact with objects. Therefore, I don't have to know if they're actually there"*, P10). As a result, participants were less open to setting objects within their reach as fully transparent for the blocks task.

If participants felt the need to modify the appearance of close-by objects nonetheless, they saw the semi-transparent and outline settings as desirable options for achieving the effect of DR while provide themselves with more context of the objects they were concerned about colliding with (Figures 10 and 11). Participants largely regarded hiding-and-outlining and setting objects as semi-transparent as serving the same purpose - they help reduce distraction while retaining an operational level of contextual understanding such that collisions can be avoided (*"applying diminished reality helps me keep in mind that objects are there, so I can avoid them, while making them less distracting"*, P1). The decision to hide and outline as opposed to setting objects as semi-transparent appeared to primarily be based on personal preferences. For instance, several participants regarded the outlines themselves as a source of distraction (P3, P6, P9). One participant made decisions about whether a given object should be left unaltered or partially diminished based on the ramifications of collisions with the object. For instance, several participants reasoned that causing a cup to spill may be more problematic than dropping a bag of tissues (*"I think the tissues and looks like wallets and books can be fully transparent because I think there are some things that like if knocked onto the floor, it doesn't matter"*, P13).

In summary, participants considered the following aspects with respect to spatial arrangement. First, they considered the objects' distances from themselves ($N = 12$). They generally found fully diminishing objects within reach uncomfortable and disorienting; diminishing objects out of their reach was considered more acceptable. Second, they considered each object's visibility ($N = 10$). Participants were generally indifferent to the state of objects (i. e., either diminished or unmodified) which were either outside their point of view or physically occluded.

Movement. In augmenting their environment using DR, participants additionally took into account whether a given object was moving. We included a television with moving imagery in both environments and additional active monitors in the private environment. Eight participants made note of how dynamic objects are particularly distracting. However, of the eight participants, not all consistently applied DR to either hide or reduce the presence of the screens within the room (Figures 10 and 11). We attribute this in part to the fact that the screens were not necessarily within their task-focused vision. Additionally, several participants reasoned that since the objects are dynamic, it may be beneficial to have a peripheral awareness of their state in spite of how that requires them to divert their attention from their primary task (*"I have a cat. I might not want to see like them completely, but I kind of want to track their movement"*, P3). For dynamic objects, especially moving screen-based content, participants suggested introducing a freezing effect rather than hiding as an alternative.

Object interaction requirements and functionality. The majority of our participants ($N = 10$) repeatedly considered the functionality and relevance of individual objects in making their decisions about DR usage (i. e., which objects are expected to be used during or shortly after the task?). In the video task for instance, several participants left a subset of the items on the table unaltered because they imagined accessing them while accomplishing the task (*"Maybe this time I will keep these three objects [referring to phone, coffee, and pastry] as they are because I want to drink a coffee, look at my phone, or eat the pastry while watching the video"*, P16). Several participants regarded the phone in particular as something they are already used to having present next to them in their daily lives, therefore leaving it unmodified while diminishing the objects in its surroundings (*"The phone is already something I'm used to having in front of me"*, P3). Select objects are also perceived to have certain connotations, and depending on the context, it may be beneficial to have them obfuscated. For instance, P6 noted that it may be beneficial to hide the guitar placed in the private environment as he associates it with something he can play. Sharing this sentiment, P4 preferred for the books on his table to be diminished so it doesn't "reminds [him] of stuff that would distract him."

Social presence. A class of objects which our participants paid special attention to was people. Six participants left the visual appearance of the people in their surroundings entirely unmodified across both tasks, four set some as either hidden but outlined or semi-transparent, and two set them as fully transparent. Participants indicated that they were generally uncomfortable with applying DR on others ($N = 8$). Several participants felt that it would be beneficial to be aware of the people around them as well as their surroundings when others are present ($N = 5$). One scenario that participants frequently described as an explanation for this need involves a diminished colleague asking them a question (e. g., *"what happens if one of them has a question for me or wants to talk to me? then I'm going to hear a disembodied voice coming out of the void!"*, P6). A subset of the participants furthermore explicitly stated that applying visual alterations to others would be disrespectful ($N = 3$) (*"I feel like it's there's something about making someone disappear, they have like a identity attached to them"*, P1). Participants were additionally concerned with potential awkward situations that arise out of a disjointed understanding of reality (e. g., *"I can imagine my*




Figure 10: Distribution of transparency effect usage by object. The color and size of the circles indicates the percentage of participants who applied the specified augmentation on the object it overlays. At most 83% of participants chose to augment any given object. We highlight objects which $\geq 25\%$ of participants chose to apply the same effect. Top: objects entirely hidden. Bottom: objects set as transparent but outlined.




Figure 11: Distribution of semi-transparent effect usage by object. The color and size of the circles indicates the percentage of participants who applied the specified augmentation on the object it overlays. At most 83% of participants chose to augment any given object. We highlight objects which $\geq 25\%$ of participants chose to apply the same effect. In the public environment, the semi-transparency effect was not consistently used on the same objects by our participants ($< 25\%$ agreement). Hence, we only visualize the usage of the effect in the private environment. Left: objects set as semi-transparent. Right: objects set as semi-transparent and outlined.

friends were telling me like, ‘check out what’s happening with TV.’ And I’m like, ‘where’s the TV?’ And I’ll look stupid.”, P9). In certain contexts, the presence of people may also be desirable, such as to boost productivity (“When I go to co-working spaces, I usually just go there and go on my laptop. Something about other people working helps me focus”, P1). Several participants furthermore noted that the presence of others introduces a consideration of ownership ($N = 3$). They would be cautious towards diminishing objects that they perceived as belonging to others. In the words of P1, “similar to when you go to a library, or you’re sitting in an airport, ... and

you see someone save a seat with a jacket, you probably won’t sit there.”

Additional considerations. In addition to the four main considerations of spatial arrangement, movement, interaction requirements, and social presence, from our qualitative data and observations, we identified two minor considerations: object saliency (e. g., bright colors) and the physical plausibility of the result (e. g., diminishing a chair underneath a person was considered strange by one participant because it is not physically possible, P3).

6 DISCUSSION

Most prior work was focused on either making DR technically feasible, or completely erasing target objects from users' visual field for applications such as improved vehicle safety, interior design, or privacy. In our work, we aimed to take a step back and asked if this type of augmentation is something that users actually want, what type of effects are preferred, and to what extent DR should be applied. Through two empirical studies, we found that users' attitude towards DR is generally positive, although its application is highly dependent on context, including task and environment. They preferred simple yet effective effects, while maintaining contextual awareness. Changing the opacity of a target while adding an outline achieved this goal for most participants. In the following, we distill guidelines from both evaluations, and discuss limitations of our approach, ethical and safety implications, and further work.

6.1 Implications for Design

Since our goal was to better understand how we should apply DR to optimize for user comfort, we will present our discussion of both studies in the form of recommendations:

- **Opacity as primary DR effect.** While opacity adjustment may not always be the optimal effect, participants generally agreed that it is suitable and effective, and considered straightforward and aesthetically acceptable.
- **Maintain contextual understanding.** In many situations, users may desire a reduced understanding of their surrounding environment, but not necessarily to the extent where they are wholly unaware of everything around them. Ideally, they would prefer maintaining a vague understanding of their context (i. e., a less distracting or information-heavy representation). This can be achieved with an outline effect, for instance. Techniques for providing contextual understanding at lower levels of distraction, like outline extraction, may therefore be valuable for DR implementations in the future.
- **Retain agency.** Users were concerned about losing control of the augmentations applied to their vision, particularly due to safety implications. As a result, they generally felt uncomfortable without agency over the mechanism. Users desired, at the minimum, an awareness that their perception of the environment was augmented with DR.
- **Avoid unintended collisions.** Users were concerned with the potential to physically collide with diminishing object. Spatial context of the environment within a user's reach should always be retained.
- **Consider object interaction requirements and behaviour.** Objects which are immediately or soon-to-be relevant to users should generally be kept unmodified or at least visually accessible. Moving objects or objects with dynamic content (e. g., screens), while distracting, might be relevant to users' situational awareness.
- **Social presence increases contextual awareness.** Users want to retain awareness of other people within their surroundings, as well as the surrounding context when others are present. Several considerations driving this preference are expectations of potential interactions (e. g., if a colleague

asks a question, an awareness of their position in the environment would be beneficial) and a respect for others (e. g., not intruding upon others' personal space or belongings).

We believe that above guidelines serve as stepping stones towards the implementation of context-aware DR systems.

6.2 Limitations

We chose to simulate the experience of DR using VR instead of employing AR headsets. This was motivated primarily by enhanced control over the environment and because a satisfying implementation of DR on see-through head-mounted display is still beyond the capabilities of commercially-available devices. To the best of our ability, we repeatedly tried to articulate that the DR effects applied to the environment are visual augmentations rather than physical adjustments. This was also reinforced through the inclusion of haptic feedback in the second experiment. The degree of perceived realism, however, varies across participants and environments.

For instance, the 3D environments employed within the experiment, though arguably comparable in fidelity to those typically used for HCI user studies (e. g., [6, 29]), are not entirely photo-realistic and contain artefacts due to the limitations of current commercial scanning devices. Though the artifacts are typically concentrated in areas outside the focus region of our experiment tasks, seeing the artifacts may still result in a lapse in participant immersion. Likewise, since we can only approximate the experience of DR, participants may occasionally forget the key caveat that the augmentations are visual rather than physical.

Future work leveraging haptic re-targeting may be helpful in verifying the results of our study. We nonetheless believe the trade-off was warranted, since it allowed us to create a better representation of DR that's not limited by current technology.

In the first experiment, we evaluated user preferences for DR effects in four sample scenarios. While the scenarios we curated are meant to span several dimensions, they strike a balance between covering a large space and limited experimental duration. The same holds true for the target objects that could be modified in each environment. Lastly, while participants were instructed to imagine different tasks, they did not actually perform them. All those factors might influence the results. The consistency across the scenarios and tasks makes us believe, though, that the results on effect types are valid and generalize beyond the presented scenarios. We plan to verify this in future experiments.

In the second experiment, the selection of tasks and the employed time limit might have influenced participants' perception of DR. Furthermore, all environments were unfamiliar to participants. Results indicate that participants were comfortable with diminishing around 30% of objects. It will be interesting in the future to repeat a similar experiment in participants' own environments, and with longer tasks, to see if this changes the usage of DR.

6.3 Ethical & safety considerations

The goal of this work is to explore the feasibility and user preference in a scenario where DR is technically feasible, rather than advocate for its deployment in all AR systems. We wanted to find out if DR is conceptually something that users could see value in, and it seems they do. Diminished Reality, like most augmentation techniques,

has the potential to lead to ethical and safety challenges, which must be addressed in any real life implementation of the technology.

There are clear safety implications relating to a mismatch between a user's visual perception of the environment and the actual physical state of the environment. A lack of understanding of one's spatial context can lead to unintended collisions. Diminishing moving objects may also lead to unexpected safety issues. Future DR systems must therefore implement guidance mechanisms to complement its usage to preventing such situations, similar to chaperoning techniques in VR.

In addition to potential physical safety hazards, DR could further increase the issue of information filter bubbles. While this is prevalent in all AR interfaces, DR's unique ability to hide information, objects or people could lead to further exacerbation of the problem. While further work could address this by limiting the target objects to which DR can be applied to, addressing issues of communications and the implications of personalized reality-modifying technologies remains an open problem. This challenge is even more apparent when considering bad actors, who might manipulate users towards making decisions against their self interest, e. g., decreasing the saliency of safety hazards. Ensuring privacy, data integrity and security on a device level and application level is therefore a premier challenge that needs to be addressed (cf. Roesner et al. [48]).

Lastly, the possibility of diminishing people has a large set of deeply problematic and dystopian implications, such as social exclusion. While both our study and the study of Kari et al. [25] discussed users' attitudes towards diminishing and substituting people, questions of what the boundaries should be remains open and important to contemplate.

6.4 Future Work

The goal of our work is to be starting point for building a more human-centered understanding of DR and its potential applications. In addition to addressing the limitations articulated in Section 6.2, there is a large range of potential for future research. Firstly, after determining a workable DR effect, we focused on investigating its use for augmenting cluttered environments to support task performance. The utility of DR as filtering mechanism in support of search remains an interesting open question. Future work may want to consider comparing the mechanism with other attention direction approaches. Secondly, building on the insights from this study, an interesting direction is to explore how the considerations articulated in Section 6.3 may be quantified in service of automating the process. Lastly, we currently only explore DR from a visual perspective. An interesting direction may be to investigate multi-sensory diminishing via combining our insights with the work of Haas et al. [13] on auditory mediated reality.

7 CONCLUSION

We present insights from two empirical studies exploring potential usage scenarios of diminished reality, as well as user perceptions of and interactions in DR-enabled environments. In the first study, we investigate what augmentation effects users prefer within different scenarios. We found that across the different scenarios we presented, users strongly preferred adjusting the object opacity to diminish their appearance. Users additionally required mechanisms

for revealing the hidden context in a less distracting manner (e. g., outlining). In the second study, we explored the user experience of performing tasks within DR-enabled environments. Participants customized their surroundings via adjusting the opacity of objects and enabling outline. They compared their experience in their custom usage of DR with their experience in the same environment with no DR applied, and with DR applied to all task-irrelevant items. Participants generally saw benefits in decreasing their contextual understanding, particularly for performing less mobile tasks. We provide a set of design guidelines, meant to support content creators in integrating DR into visual augmentation devices. We hope our insights will inform further work on understanding DR, and visual augmentations in general.

REFERENCES

- [1] Amazon.com. 2021. Amazon Berkeley Object Dataset. <https://amazon-berkeley-objects.s3.amazonaws.com/index.html#home>
- [2] Autodesk. 2021. Meshmixer. <https://www.meshmixer.com/>
- [3] Benjamin Avery, Christian Sandor, and Bruce H Thomas. 2009. Improving Spatial Perception for Augmented Reality X-Ray Vision. In *2009 IEEE Virtual Reality Conference*. IEEE Computer Society, USA, 79–82. <https://doi.org/10.1109/VR.2009.4811002>
- [4] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. 2009. PatchMatch: a randomized correspondence algorithm for structural image editing. *ACM transactions on graphics* 28, 3 (July 2009), 1–11. <https://doi.org/10.1145/1531326.1531330>
- [5] Peter Barnum, Yaser Sheikh, Ankur Datta, and Takeo Kanade. 2009. Dynamic seethroughs: Synthesizing hidden views of moving objects. In *2009 8th IEEE International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, USA, 111–114. <https://doi.org/10.1109/ISMAR.2009.5336483>
- [6] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. 2021. SemanticAdapt: Optimization-Based Adaptation of Mixed Reality Layouts Leveraging Virtual-Physical Semantic Connections. In *The 34th Annual ACM Symposium on User Interface Software and Technology (Virtual Event, USA) (UIST '21)*. Association for Computing Machinery, New York, NY, USA, 282–297. <https://doi.org/10.1145/3472749.3474750>
- [7] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. 2008. MeshLab: an Open-Source Mesh Processing Tool. In *Eurographics Italian Chapter Conference*, Vittorio Scarano, Rosario De Chiara, and Ugo Erra (Eds.). The Eurographics Association, Salerno, Italy, 129–136. <https://doi.org/10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136>
- [8] Facebook. 2021. Oculus. <https://www.oculus.com/>
- [9] Mar Gonzalez-Franco, Eyal Ofek, Ye Pan, Angus Antley, Anthony Steed, Bernhard Spanlang, Antonella Maselli, Domna Banakou, Nuria Pelechano, Sergio Orts-Escolano, Veronica Orvalho, Laura Trutoiu, Markus Wojcik, Maria V. Sanchez-Vives, Jeremy Bailenson, Mel Slater, and Jaron Lanier. 2020. The Rocketbox Library and the Utility of Freely Available Rigged Avatars. *Frontiers in Virtual Reality* 1 (2020), 20. <https://doi.org/10.3389/frvir.2020.561558>
- [10] The Khronos Group. 2021. SideQuest. <https://sidequestvr.com/>
- [11] Anselm Grundhöfer and Daisuke Iwai. 2018. Recent advances in projection mapping algorithms, hardware and applications. *Computer graphics forum: journal of the European Association for Computer Graphics* 37, 2 (May 2018), 653–675. <https://doi.org/10.1111/cgf.13387>
- [12] Jake Guida and Misha Sra. 2020. Augmented Reality World Editor. In *26th ACM Symposium on Virtual Reality Software and Technology (Virtual Event, Canada) (VRST '20)*. Association for Computing Machinery, New York, NY, USA, Article 47, 2 pages. <https://doi.org/10.1145/3385956.3422125>
- [13] Gabriel Haas, Evgeny Stemasov, Michael Rietzler, and Enrico Rukzio. 2020. Interactive Auditory Mediated Reality: Towards User-Defined Personal Soundscapes. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. Association for Computing Machinery, New York, NY, USA, 2035–2050. <https://doi.org/10.1145/3357236.3395493>
- [14] Kunihiko Hasegawa and Hideo Saito. 2015. Diminished Reality for Hiding a Pedestrian using Hand-Held Camera. In *2015 IEEE International Symposium on Mixed and Augmented Reality Workshops*. IEEE Computer Society, USA, 47–52. <https://doi.org/10.1109/ISMARW.2015.18>
- [15] Sjoerd Hendriks, Simon Mare, Mafalda Gamboa, and Mehmet Aydın Baytaş. 2021. Azalea: Co-Experience in Remote Dialog through Diminished Reality and Somaesthetic Interaction Design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 261, 11 pages. <https://doi.org/10.1145/3411461.3411721>

- //doi.org/10.1145/3411764.3445052
- [16] Jan Herling and Wolfgang Broll. 2010. Advanced self-contained object removal for realizing real-time Diminished Reality in unconstrained environments. In *2010 IEEE International Symposium on Mixed and Augmented Reality*. IEEE Computer Society, USA, 207–212. <https://doi.org/10.1109/ISMAR.2010.5643572>
 - [17] Sei Ikeda, Iwao Takemura, Asako Kimura, and Fumihisa Shibata. 2018. Diminished Reality System Based on Open-Source Software for Self-Driving Mobility. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE Computer Society, USA, 354–357. <https://doi.org/10.1109/ISMAR-Adjunct.2018.00103>
 - [18] Masahiko Inami, Naoki Kawakami, and Susumu Tachi. 2003. Optical camouflage using retro-reflective projection technology. In *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings*. IEEE Computer Society, USA, 348–349. <https://doi.org/10.1109/ISMAR.2003.1240754>
 - [19] Laurent Itti, Christof Koch, and Ernst Niebur. 1998. A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE transactions on pattern analysis and machine intelligence* 20, 11 (Nov. 1998), 1254–1259. <https://doi.org/10.1109/34.730558>
 - [20] Daisuke Iwai, Sawako Hanatani, Chinatsu Horii, and Kosuke Sato. 2006. Limpid Desk: Transparentizing Documents on Real Desk in Projection-Based Mixed Reality. In *IEEE Virtual Reality Conference (VR 2006)*. IEEE Computer Society, USA, 319–319. <https://doi.org/10.1109/VR.2006.95>
 - [21] JASP Team. 2020. JASP (Version 0.14.1)[Computer software]. <https://jasp-stats.org/>
 - [22] Brett Jones, Hrvoje Benko, Eyal Ofek, and Andrew D Wilson. 2013. IllumiRoom: peripheral projected illusions for interactive experiences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 869–878. <https://doi.org/10.1145/2470654.2466112>
 - [23] Brett Jones, Rajinder Sodhi, Michael Murdock, Ravish Mehra, Hrvoje Benko, Andrew Wilson, Eyal Ofek, Blair MacIntyre, Nikunj Raghuvanshi, and Lior Shapira. 2014. RoomAlive: magical experiences enabled by scalable, adaptive projector-camera units. In *Proceedings of the 27th annual ACM symposium on User interface software and technology (Honolulu, Hawaii, USA) (UIST '14)*. Association for Computing Machinery, New York, NY, USA, 637–644. <https://doi.org/10.1145/2642918.2647383>
 - [24] Denis Kalkofen, Eduardo Veas, Stefanie Zollmann, Markus Steinberger, and Dieter Schmalstieg. 2013. Adaptive ghosted views for Augmented Reality. In *2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE Computer Society, USA, 1–9. <https://doi.org/10.1109/ISMAR.2013.6671758>
 - [25] Mohamed Kari, Tobias Grosse-Puppenthal, Luis Falconeri Coelho, Andreas Rene Fender, David Bethge, Reinhard Schütte, and Christian Holz. 2021. TransformMR: Pose-Aware Object Substitution for Composing Alternate Mixed Realities. In *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE Computer Society, USA, 69–79. <https://doi.org/10.1109/ISMAR52148.2021.00021>
 - [26] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. 2019. Deep Video Inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, USA, 5792–5801.
 - [27] Hanseob Kim, TaeHyung Kim, Myungho Lee, Gerard Jounghyun Kim, and Jae-In Hwang. 2020. Don't Bother Me: How to Handle Content-Irrelevant Objects in Handheld Augmented Reality. In *26th ACM Symposium on Virtual Reality Software and Technology (Virtual Event, Canada) (VRST '20)*. Association for Computing Machinery, New York, NY, USA, Article 32, 5 pages. <https://doi.org/10.1145/3385956.3418948>
 - [28] Robert Kosara, Silvia Miksch, and Helwig Hauser. 2001. Semantic Depth of Field. In *Proceedings of the IEEE Symposium on Information Visualization 2001 (INFOVIS '01) (INFOVIS '01)*. IEEE Computer Society, USA, 97.
 - [29] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. 2019. Context-Aware Online Adaptation of Mixed Reality Interfaces. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (New Orleans, LA, USA) (UIST '19)*. Association for Computing Machinery, New York, NY, USA, 147–160. <https://doi.org/10.1145/3332165.3347945>
 - [30] David Lindlbauer, Jörg Müller, and Marc Alexa. 2016. Changing the Appearance of Physical Interfaces Through Controlled Transparency. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (Tokyo, Japan) (UIST '16)*. Association for Computing Machinery, New York, NY, USA, 425–435. <https://doi.org/10.1145/2984511.2984556>
 - [31] David Lindlbauer, Jörg Müller, and Marc Alexa. 2017. Changing the Appearance of Real-World Objects By Modifying Their Surroundings. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (Denver, Colorado, USA) (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3954–3965. <https://doi.org/10.1145/3025453.3025795>
 - [32] David Lindlbauer and Andrew D. Wilson. 2018. *Remixed Reality: Manipulating Space and Time in Augmented Reality*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173703>
 - [33] EyeCue Vision Technologies LTD. 2021. Qlone. <https://www.qlone.pro/>
 - [34] Steve Mann. 1999. Mediated reality. *Linux J.* 1999, 59es (March 1999), 5–es.
 - [35] Steve Mann. 2002. Mediated reality with implementations for everyday life. *Presence Connect* 1 (2002).
 - [36] Steve Mann, Tom Furness, Yu Yuan, Jay Iorio, and Zixin Wang. 2018. All Reality: Virtual, Augmented, Mixed (X), Mediated (X,Y), and Multimeditated Reality. arXiv:1804.08386 [cs.HC]
 - [37] Siim Meerits and Hideo Saito. 2015. Real-Time Diminished Reality for Dynamic Scenes. In *2015 IEEE International Symposium on Mixed and Augmented Reality Workshops*. IEEE Computer Society, USA, 53–59. <https://doi.org/10.1109/ISMARW.2015.19>
 - [38] Erick Mendez, Steven Feiner, and Dieter Schmalstieg. 2010. Focus and Context in Mixed Reality by Modulating First Order Salient Features. In *Proceedings of the 10th International Conference on Smart Graphics (Banff, Canada) (SG '10)*. Springer-Verlag, Berlin, Heidelberg, 232–243.
 - [39] Shohei Mori, Okan Erat, Wolfgang Broll, Hideo Saito, Dieter Schmalstieg, and Denis Kalkofen. 2020. InpaintFusion: Incremental RGB-D Inpainting for 3D Scenes. *IEEE Transactions on Visualization and Computer Graphics* 26, 10 (2020), 2994–3007. <https://doi.org/10.1109/TVCG.2020.3003768>
 - [40] Shohei Mori, Jan Herling, Wolfgang Broll, Norihiko Kawai, Hideo Saito, Dieter Schmalstieg, and Denis Kalkofen. 2018. 3D PixMix: Image Inpainting in 3D Environments. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE Computer Society, USA, 1–2. <https://doi.org/10.1109/ISMAR-Adjunct.2018.00020>
 - [41] Shohei Mori, Sei Ikeda, and Hideo Saito. 2017. A survey of diminished reality: Techniques for visually concealing, eliminating, and seeing through real objects. *IPSP Transactions on Computer Vision and Applications* 9, 1 (2017), 1–14.
 - [42] Takuji Narumi, Yuki Ban, Takashi Kajinami, Tomohiro Tanikawa, and Michitaka Hirose. 2012. Augmented Perception of Satiety: Controlling Food Consumption by Changing Apparent Size of Food with Augmented Reality. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Austin, Texas, USA) (CHI '12)*. Association for Computing Machinery, New York, NY, USA, 109–118. <https://doi.org/10.1145/2207676.2207693>
 - [43] Occipital. 2021. Canvas. <https://canvas.io/>
 - [44] Jennifer Pearson, Simon Robinson, Matt Jones, Anirudha Joshi, Shashank Ahire, Deepak Sahoo, and Sriram Subramanian. 2017. *Chameleon Devices: Investigating More Secure and Discreet Mobile Interactions via Active Camouflaging*. Association for Computing Machinery, New York, NY, USA, 5184–5196. <https://doi.org/10.1145/3025453.3025482>
 - [45] Alexander Plopski, Ada Virginia Taylor, Elizabeth Jeanne Carter, and Henny Admoni. 2019. InvisibleRobot: Facilitating Robot Manipulation Through Diminished Reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE Computer Society, USA, 165–166. <https://doi.org/10.1109/ISMAR-Adjunct.2019.00-55>
 - [46] Francois Rameau, Hyowon Ha, Kyungdon Joo, Jinsoo Choi, Kibaek Park, and In So Kweon. 2016. A Real-Time Augmented Reality System to See-Through Cars. *IEEE transactions on visualization and computer graphics* 22, 11 (Nov. 2016), 2395–2404. <https://doi.org/10.1109/TVCG.2016.2593768>
 - [47] Jan Ole Rixen, Teresa Hirzle, Mark Colley, Yannick Etzel, Enrico Rukzio, and Jan Gugenheimer. 2021. Exploring Augmented Visual Alterations in Interpersonal Communication. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 730, 11 pages. <https://doi.org/10.1145/3411764.3445597>
 - [48] Franziska Roesner, Tadayoshi Kohno, and David Molnar. 2014. Security and privacy for augmented reality systems. *Commun. ACM* 57, 4 (April 2014), 88–96. <https://doi.org/10.1145/2580723.2580730>
 - [49] Ruth Rosenholtz, Yuanzhen Li, and Lisa Nakano. 2007. Measuring visual clutter. *Journal of vision* 7, 2 (Jan. 2007), 17–17. <https://doi.org/10.1167/7.2.17>
 - [50] Shunsuke Sakai, Yohei Yanase, Yasutsuna Matayoshi, and Masahiko Inami. 2018. D-Ball: Virtualized Sports in Diminished Reality. In *Proceedings of the First Superhuman Sports Design Challenge: First International Symposium on Amplifying Capabilities and Competing in Mixed Realities (Delft, Netherlands) (SHS '18)*. Association for Computing Machinery, New York, NY, USA, Article 6, 6 pages. <https://doi.org/10.1145/3210299.3210305>
 - [51] Shunsuke Sakai, Yohei Yanase, Yasutsuna Matayoshi, and Masahiko Inami. 2018. D-Ball: Virtualized Sports in Diminished Reality. In *Proceedings of the First Superhuman Sports Design Challenge: First International Symposium on Amplifying Capabilities and Competing in Mixed Realities (Delft, Netherlands) (SHS '18)*. Association for Computing Machinery, New York, NY, USA, Article 6, 6 pages. <https://doi.org/10.1145/3210299.3210305>
 - [52] Byung-Kuk Seo, Moon-Hyun Lee, Hanhoon Park, and Jong-Il Park. 2008. Projection-Based Diminished Reality System. In *2008 International Symposium on Ubiquitous Virtual Reality*. IEEE Computer Society, USA, 25–28. <https://doi.org/10.1109/ISUVR.2008.21>
 - [53] Sanni Siltanen. 2017. Diminished reality for augmented reality interior design. *The Visual Computer* 33, 2 (2017), 193–208.
 - [54] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler,

- Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. 2019. The Replica Dataset: A Digital Replica of Indoor Spaces. arXiv:1906.05797 [cs.CV]
- [55] Ada V. Taylor, Ayaka Matsumoto, Elizabeth J. Carter, Alexander Plopski, and Henny Admoni. 2020. Diminished Reality for Close Quarters Robotic Telemanipulation. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE Computer Society, USA, 11531–11538. <https://doi.org/10.1109/IROS45743.2020.9341536>
- [56] Eduardo E. Veas, Erick Mendez, Steven K. Feiner, and Dieter Schmalstieg. 2011. *Directing Attention and Influencing Memory with Visual Saliency Modulation*. Association for Computing Machinery, New York, NY, USA, 1471–1480. <https://doi.org/10.1145/1978942.1979158>
- [57] J. Wolfe and T. S. Horowitz. 2008. Visual search. *Scholarpedia* 3, 7 (2008), 3325. <https://doi.org/10.4249/scholarpedia.3325> revision #145401.
- [58] Asim Evren Yantaç, Doga Çorlu, Morten Fjeld, and Andreas Kunz. 2015. Exploring Diminished Reality (DR) Spaces to Augment the Attention of Individuals with Autism. In *2015 IEEE International Symposium on Mixed and Augmented Reality Workshops*. IEEE Computer Society, USA, 68–73. <https://doi.org/10.1109/ISMARW.2015.21>
- [59] Lining Yao, Anthony DeVincenzi, Anna Pereira, and Hiroshi Ishii. 2013. FocalSpace: Multimodal Activity Tracking, Synthetic Blur and Adaptive Presentation for Video Conferencing. In *Proceedings of the 1st Symposium on Spatial User Interaction* (Los Angeles, California, USA) (*SUI '13*). Association for Computing Machinery, New York, NY, USA, 73–76. <https://doi.org/10.1145/2491367.2491377>
- [60] Siavash Zokai, Julien Esteve, Yakup Genc, and Nassir Navab. 2003. Multiview Paraperspective Projection Model for Diminished Reality. In *Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR '03)*. IEEE Computer Society, USA, 217.