Thinking Outside the Lab: VR Size & Depth Perception in the Wild

RAHUL ARORA, University of Toronto, Canada

JIANNAN LI, University of Toronto, Canada

GONGYI SHI, University of Toronto, Canada
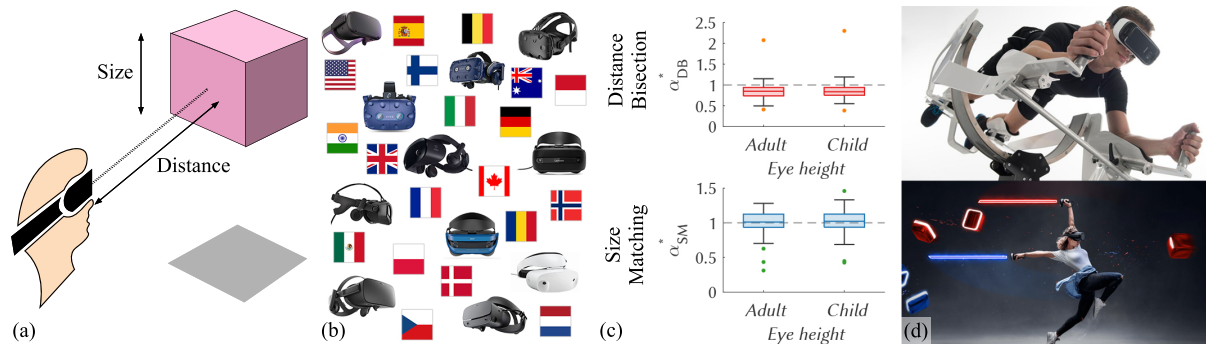
KARAN SINGH, University of Toronto, Canada

Fig. 1. We conducted a remote unsupervised "in-the-wild" study to understand size and egocentric distance perception in VR using a gamified methodology (a). Our 60-participant study spanned 19 countries and 11 different HMDs (b). The summary of the two perceptual matching tasks is shown here (c): distance bisection judgements show a trend of nonlinear distance foreshortening (top) while size matching judgements suggest neither foreshortening nor anti-foreshortening (bottom). These measurements can inform "distance correction" functions that help improve user performance in applications such as VR flight simulators and games such as Beat Saber (d). © ICAROS (top) and Beat Games (bottom).

Size and distance perception in Virtual Reality (VR) have been widely studied, albeit in a controlled laboratory setting with a small number of participants. We describe a fully remote perceptual study with a gamified protocol to encourage participant engagement, which allowed us to quickly collect high-quality data from a large, diverse participant pool (N=60). Our study aims to understand medium-field size and egocentric distance perception in real-world usage of consumer VR devices. We utilized two perceptual matching tasks—distance bisection and size matching—at the same target distances of 1–9 metres. While the bisection protocol indicated a near-universal trend of nonlinear distance compression, the size matching estimates were more equivocal. Varying eye-height from the floor plane showed no significant effect on the judgements. We also discuss the pros and cons of a fully remote perceptual study in VR, the impact of hardware variation, and measures needed to ensure high-quality data.

CCS Concepts: • **Human-centered computing** → **Virtual reality**; • **Computing methodologies** → **Perception**.

Additional Key Words and Phrases: Depth perception, size perception, virtual reality, remote studies, gamification.

1 INTRODUCTION

Understanding human perception in Virtual Reality (VR) is a fundamental question in VR research. From casual applications such as games and movies to performance-critical technical training, modelling biases in people's perception of their surroundings in VR is an important endeavour. Unsurprisingly then, ever since the early days

Authors' addresses: Rahul Arora, University of Toronto, 40 St. George Street, Toronto, Ontario, Canada, arorar@dgp.toronto.edu; Jiannan Li, University of Toronto, 40 St. George Street, Toronto, Ontario, Canada, jiannanli@dgp.toronto.edu; Gongyi Shi, University of Toronto, 40 St. George Street, Toronto, Ontario, Canada, gongyi.shi@mail.utoronto.ca; Karan Singh, University of Toronto, 40 St. George Street, Toronto, Ontario, Canada, karan@dgp.toronto.edu.

of VR, visual perception has received considerable attention in psychology and cognitive science research [1, 2, 12]. Of particular interest is the perception of the geometric attributes of visual space, namely distance and size. Real-world size and distance perception experiments have a rich century-old history in cognitive science [9, 30, 77], and similar experiments in VR serve a dual purpose. Not only do such studies advance our understanding of visual perception in VR, but a virtual world closely simulating reality provides a unique opportunity to construct environments that are difficult or impossible in the real world. Such fantastical constructs can be used to manipulate perceptual cues at will and gain insights into the human visual system [5]. VR size and distance perception studies can thus improve our understanding of real-world perception as well. At the same time, due to the shortcomings of current VR devices, such as the well-known vergence-accommodation conflict [27], it is essential to conduct perceptual studies in VR to inform the design of VR applications where perceptual accuracy is crucial.

Most geometric perception studies, whether real-world or in VR, utilize strictly controlled laboratory conditions (e.g. [11, 20, 21, 54]). While maintaining strict control over the experiment is invaluable for gaining theoretical insights into the human visual system as noted above, the practical utility of the results to real-world applications is limited. For example, a typical perceptual experiment in VR would utilize the same hardware for all users, ensure strict calibration of the interpupillary distance (IPD), and control the position of the head-mounted display (HMD) on the user's head. In real-world usage, users utilize diverse VR setups, and are typically untrained in proper IPD adjustment—adjusting IPD and HMD position for comfort, rather than physical accuracy. Recently, Hornsey et al. [29] reported encouraging results on a size and shape constancy task where a) two different devices, Oculus Rift and HTC Vive, were utilized and b) IPD was not adjusted or controlled for. However, the results were still conducted in a laboratory setting, requiring significant investment of experimenter and participant time and effort. Inspired by their results, we present the first remote study on geometric perception in VR (Fig. 1a). Our size and distance perception experiment aims at achieving results that are *ecologically valid*, that is, we want to model real-world HMD usage to inform the development of VR applications. In addition to removing strict controls on the experiment conditions, a remote study enables us to access a large and diverse pool of participants (Fig. 1b). In this work, we report on a medium-scale study with N=60. Studies such as ours are especially relevant today owing to the COVID-19 physical distancing requirements which make remote studies an attractive option for experimental research involving human subjects.

Study Summary. In our experiment, we study the perception of size and egocentric distance in VR using two different measurement protocols. The first protocol involves a classical size matching task [28, 66], where participants resize a nearby *comparison* object to match the physical size of a relatively distant *reference* object. Judgements of perceived egocentric distance can be achieved by assuming the size-distance invariance hypothesis (SDIH) [13] (also see Howard [30, Sec. 29.3.2], Wagner [77]) and inverting the perceived size judgements. The classical SDIH posits that given a visual (retinal) size, the perceived size of an object is proportional to the perceived distance to it. Our second protocol—distance bisection—elicits distance perception judgements directly: participants move the *comparison* object so that it lies half as far from them as a fixed *reference* object. Thus, the *reference* object is placed at similar positions for both the protocols, while the *comparison* is farther in the distance bisection protocol as compared to the size matching protocol. While our protocol choices are similar to Gruber [23], who applied both protocols and compared their results to evaluate SDIH in real environments, we do not formally evaluate the applicability of SDIH in our scenario due to the differences in the egocentric position of the *comparison*. In addition, our protocols elicit perceived distance *relative* to the physical distance, but do not help estimate the perceived *absolute* distance. This relative measure indicates whether the visual space exhibits *foreshortening*, that is, fixed depth intervals being increasingly compressed with further distance from the observer, or *anti-foreshortening*, where the visual space is expanded. Accurate judgements indicate that the perceived distance varies linearly with the physical distance and the visual space is *invariant*.

A second factor, explored with both the measurement protocols, is eye height. We compare distance and size perception when the scene was rendered assuming an eye height similar to an average adult (170 cm) vs. a much shorter height of 50 cm. This investigation is motivated by two disparate applications—seated or short characters (child, dwarf, or quadruped) in VR games, and the use of VR for tele-operating robots that are often short [7].

A common challenge with remote studies is control over data quality. Malicious participants can develop strategies to game the system, and the collected data may need thorough analysis to detect cheating [17, 26, 67]. An additional challenge is the lack of participant engagement [17], which can reduce data validity as the experiment goes on. We designed a gamified experiment to tackle these challenges: participants were awarded points for their perceptual judgements, with more accurate judgements getting a higher score. An online leaderboard prompted participants of their rank vis-à-vis others, and top participants received a bonus compensation. Further, we broke down the tasks into easily-fulfilled "chunks" to prevent VR cybersickness or boredom due to the monotonous tasks, thereby ensuring high quality of the collected data.

Contributions. Following are our main contributions.

— The first fully remote study on geometric perception in VR.
— Understanding how perceived distance $\mathcal{R}$ varies with actual egocentric distance R using two different perceptual matching protocols. With size matching, the overall trend did not point clearly to either non-linear foreshortening or anti-foreshortening: $\mathcal{R} = kR^{1.001}$. In contrast, with distance bisection, a clear trend of foreshortening emerged: $\mathcal{R} = kR^{0.863}$.
— Testing size and depth perception with two different eye heights. Perceptual judgements were similar across the two eye heights.
— A gamified protocol for our remote study and a discussion on adapting it for other perceptual studies in VR.

2 RELATED WORK

Our work straddles three themes: understanding biases in the size and depth perception, modelling human perception in virtual environments, and remote experiments for VR and perceptual studies.

2.1 Bias in Human Depth and Size Perception

Fundamental to human vision, size and depth perception have received sustained interest in psychology and cognitive science. Howard's book [30] provides a comprehensive review of this broad area. Size and depth perception biases reveal basic mechanisms of the vision system [45] and are critically relevant to efficiency and safety in tasks such as aircraft piloting [42]. Prior research has presented rich and sometimes contrasting results about biases in egocentric depth perception. While some studies found that people can perceive depth accurately and reliably [46, 57, 62], others showed that they are prone to systematic errors [16, 19, 42]. A large body of research noted the effect of foreshortening [19, 45, 53, 72, 75]. The effect was found to be stronger with depth cues further removed [56]. However, several studies using a bisection protocol presented the opposite evidence in support of an invariant or expanded visual space. Rieser et al. [62] reported that participants were largely accurate in judging the mid-points of 4–24 metre self-to-target distances in an open field. Lappin et al. [42] found an environment-sensitive antiforeshortening effect; participants perceived the mid-points of 15 and 30 metre distances as further away than the true mid-points. These contradictory results remain unexplained yet.

Size perception is closely related to depth perception and the two are commonly measured together in prior studies [21, 24]. The size-distance invariance hypothesis (SDIH) [13, 15] directly links the two, positing that with a fixed visual angle, the perceived size is proportional to the perceived distance. Researchers have applied SDIH to measure perceived distance through perceived size as a proxy [37, 44]. However, the validity of SDIH is controversial [15]. Through size-matching tasks, a number of experiments [33, 34, 66] have found that the sizes of

further placed objects are increasingly overestimated, inconsistent with the commonly observed foreshortening effect if SDIH holds true. More directly, Gruber [23] showed that participants underestimated the sizes of objects in size matching but overestimated the distances to the same objects in distance bisection. Various attempts have been made to resolve the size-distance paradox [63]. Some researchers argue that size and depth perception are indeed independent processes [10, 24].

Our study revisits biases in these two important aspects of spatial perception in an immersive virtual environment, using a size matching and a distance bisection protocol employed in classical work (e.g. Holway and Boring [28] and Gruber [23]). The virtuality of the computer-generated world allows for conducting studies at a much larger scale (N=60) in comparison to traditional perception studies which typically involved fewer than 20 participants. While the virtual and physical worlds are perceived notably differently, we hope that data collected from large and diverse sample pools can also contribute to deepening our understanding of human vision.

2.2 Perception Studies in Immersive VR

Presenting a virtual space that can be accurately perceived by human eyes is important for VR applications to be more immersive and useful. In contrast to the somewhat contradictory results about depth and size perception in the real world, research on egocentric depth perception in VR has, in general, found a tendency of depth underestimation [61] and size underconstancy [29, 37]. In a study replicating the real-world experiment of Lappin et al. [42] in a virtual setting [3], participants placed the perceived mid-points closer than the true mid-points, in contrast to the antiforeshortening effect observed by Lappin et al. in the real-world setting. Corujeira and Oakley [7] also found depth underestimation in VR using a blind walking protocol. They further pointed out the effect of eye height on depth perception accuracy; a lying-on-the-floor eye height of 20 cm led to more accurate judgements than when seated (110 cm). Size-matching studies in VR have found underestimation of distant object sizes [29, 37, 38], largely compatible with results from studies on depth perception.

In addition to the differences in the general tendency of biases from the real world, depth perception in VR has also been found to be less precise (higher depth discrimination threshold) than in the real world [52]. This discrepancy between real and virtual environments has been of much research interest. We refer the reader to the survey by Renner et al. [61] for an in-depth review on the topic. The significant variability of available pictorial depth cues—such as texture [71], lighting [70], number of objects [38]—in different virtual environments could cause varying levels of distortion in perceived depth. The technical properties of current VR systems, such as FoV [35], IPD [4, 78], and vergence-accommodation conflict [27] may also affect depth perception accuracy.

The clear influence of hardware configurations on perception in VR may suggest that only tightly controlled lab studies offer a reliable avenue for studying VR perception. However, recent work has reported encouraging results with uncalibrated consumer VR devices [29], consistent with controlled studies with lengthy calibration processes. Our work further extends VR perception studies outside the laboratory to gather data from a large, diverse, sample pool and evaluate the feasibility of this promising methodology.

2.3 Unsupervised Remote Studies and Gamification

Online studies have recently gained strong traction for a variety of research topics, including decision making [49], social behaviour [18, 48], perception [64, 79], design optimization [40, 41], and user interface performance [39]. Researchers have noted their potential benefits of reaching the "internet scale" [49, 60]: enabling running parallel sessions to drastically reduce study turnover time and accessing large participants pools beyond the researchers' local communities. Despite the usually unsupervised and uncontrolled setting of these studies, past research has shown that proper measures can ensure high-quality data collection [49, 79].

Remote VR studies [31, 48, 50, 58, 69] are quickly becoming feasible with the increasing dissemination of consumer VR devices. Recent efforts have shown that online VR studies can efficiently gather reliable data [50, 69]

and replicate established results [31, 48]. Mottelson and Hornbæk [50] distributed VR devices to participants, who were asked to test two input techniques and a VR illusion study in an *in-the-wild* setting. They found effects comparable to traditional laboratory settings , albeit with larger variation. Ma et al. [48] were able to successfully induce behavioural manipulation in three VR studies conducted through Amazon Mechanical Turk. Huber and Gajos [31] replicated past results obtained from conventional lab settings for a spatial navigation task and a negotiation task with uncompensated VR users.

While results from prior works are encouraging, researchers have also noted uncertainties and difficulties in remote VR data collection, including variations in hardware [48], unknown instruction effectiveness [69], and ethical concerns [26, 50, 67, 69]. In particular, crowdsourced studies tend to face diminishing engagement from participants, who may optimize for return on effort rather than quality [17, 22, 25]. Gamification has been identified as a method to sustain engagement and increase data quality [47, 74]. Introducing *intrinsic* motivations, such as helping participants know themselves better, can have similar positive effects on engagement [31].

Experiments on low-level human perception have typically been conducted in supervised and strictly controlled environments to ensure internal validity. Our study explores the potential of an unsupervised, fully-remote protocol for collecting high-quality low-level perception data in VR. While building on recommended practices suggested by prior work, we introduce additional measures to combat data quality challenges unique to our study tasks, such as limiting head motion. We also add gamification mechanisms and intrinsic motivations to our study design to improve participant engagement.

3 SIZE AND DEPTH PERCEPTION STUDY

We will first describe the perceptual experiment, before detailing our gamification strategy.

3.1 Participants

Sixty participants were recruited via two main channels: a) VR interest groups on social media (Facebook/Reddit) and b) HCI/VR research e-mail lists. Almost all were male (56M, 2F, 2 unspecified), reflecting current estimates of VR ownership [59, 73]. Most were young (25: 18–25, 24: 26–35), but some were middle-aged (9: 36–45, 2: 46–55). Participants hailed from 19 countries across 5 continents; however, over half (36) were from the USA. See Fig. 2 and supplemental **S1** (Sec. 1.1) for details. Participants received US $5 (or equivalent), and the top ten on the gamified leaderboard (§ 4) were paid an additional $15.
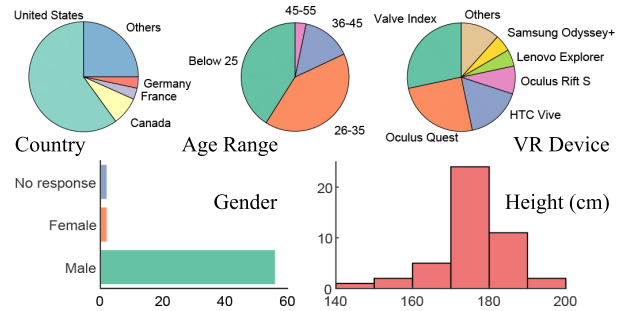

Fig. 2. Demographics information of the study participants.

3.2 Apparatus

All participants had access to a six–degree of freedom (6-DoF) VR device with a controller. They used 11 different devices, with Valve Index (17) and Oculus Quest (15) being the most popular. Device resolutions ranged between 1280×1440 and 1440×1600 per eye, while the Field of View (FoV) lay between $78° \times 88°$ and $106° \times 112°$.

3.3 Stimuli

In both the protocols, a minimal scene containing only the fixed *reference* object, the subject-controlled *comparison* object and a ground plane were rendered over a solid dark grey background. Both the objects were dull-pink coloured cubes and the ground plane was textured with grey Perlin noise [55]. Objects more complex than cubes

could have been chosen, but pilot testing showed no significant difference between cubes and complex objects such as tori and Stanford bunnies. Further, the colour choice for the cubes was arbitrary, but the colour saturation was used to clearly contrast the foreground objects against the background and ground plane consisting of greys only. In addition to ambient lighting, a directional light pointing downwards lit the scene, causing the cubes to cast shadows on the plane. The shadows were meant to emphasize the difference between the two *eye heights* to the participants. No additional context or environmental cues were provided (Fig. 3b). The off-the-ground object positioning is motivated by mid-air objects common in popular VR games such as Beat Saber and Superhot, and creation applications such as Tilt Brush and Gravity Sketch.

The stimuli were positioned relative to the participant in a spherical coordinate system (Fig. 3) centred at the participant's cyclopean eye (point midway between the eyes), with a vertical zenith ($\theta = 0°$) and horizontal reference plane ($\theta = 90°$). The initial cyclopean line of sight projected onto the horizontal served as the azimuth reference direction ($\phi = 0°$). Note that in this coordinate system, a point is represented by the vector (r, θ, ϕ), where r is the egocentric distance, or interchangeably, the depth from the participant.

In all the conditions, both the cubes were positioned near the participant's fovea. The *reference* cube was positioned so that the centre of its front face was slightly to the right of the line of sight ($\phi_{\mathrm{ref}} = -15°$) and slightly above the horizontal ($\theta_{\mathrm{ref}} = 82.5°$) and the *comparison* positioned at the same inclination, but to the left ($\phi_{\mathrm{comp}} = 15°$). For different trials, the distance r to the *reference* cube was varied between 1 and 9 metres, with 2m increments, that is, $r_{\mathrm{ref}} \in \{1m, 3m, 5m, 7m, 9m\}$. Note again, that similar to the ϕ and θ coordinates, the distance is to the centre of the front face (the one facing the subject) of the cube. The *reference* cube's size s_{ref} was chosen uniformly randomly between 25 and 35 cm; the randomization intended to reduce the chance of inadvertently turning the *reference* cube into a familiar-sized object.

Task-relevant distance cues thus available to the participants were: stereopsis, convergence, motion parallax, perspective, lighting & shading, shadows, and to a degree, texture gradient on the ground plane. We now describe protocol-dependent stimuli characteristics.

Distance Bisection Protocol. For each trial, the distance r_{comp} as well as size s_{comp} of the *comparison* were chosen so that the position was not close to the accurate response ($r_{\mathrm{comp}} \not\approx r_{\mathrm{ref}}/2$), and the size was not close to that of the *reference*: $s_{\mathrm{comp}} \not\approx s_{\mathrm{ref}}$. These choices were made to avoid participants simply choosing the initial position as their response and to avoid the use of size matching as an effective cue, respectively. Specifically, for setting r_{comp}, we first chose uniformly randomly whether to make it larger (75% of r_{ref}) or smaller (25% of r_{ref}) than the correct response, and then added a uniform random perturbation of 10%. Thus,

$$r_{\mathrm{comp}} = (f \pm \epsilon)r_{\mathrm{ref}}, \text{ where } f \sim \mathcal{U}\{0.25, 0.75\}, \ \epsilon \sim \mathcal{U}(0, 0.1). \tag{1}$$

s_{comp} was similarly chosen w.r.t. s_{ref} but with $f \sim \mathcal{U}\{0.7, 1.3\}$.

Size Matching Protocol. The size and distance parameters of the *comparison* were similarly initialized, with the only exception being the range of the random perturbation ϵ for s_{comp}, which was sampled from $\mathcal{U}(0, 0.15)$. This was motivated by the higher inter-participant dissimilarity observed in our *size matching* pilots.

3.4 Tasks and Controls

For the *bisection task*, participants were able to control the distance of the *comparison* cube by using the thumbstick on their controller. Note that θ_{comp} and ϕ_{comp} remained fixed. Participants were instructed to position the *comparison* so that the center of its front face was half the distance from them as that of the *reference*.

For the *size matching* task, participants used the thumbstick to increase/decrease the size of the *comparison*. The instruction was to match the physical size, that is, the 3D size, of the *comparison* to the *reference*. Note that the precise instruction is important here, as the size matching instruction can substantially impact the participants'
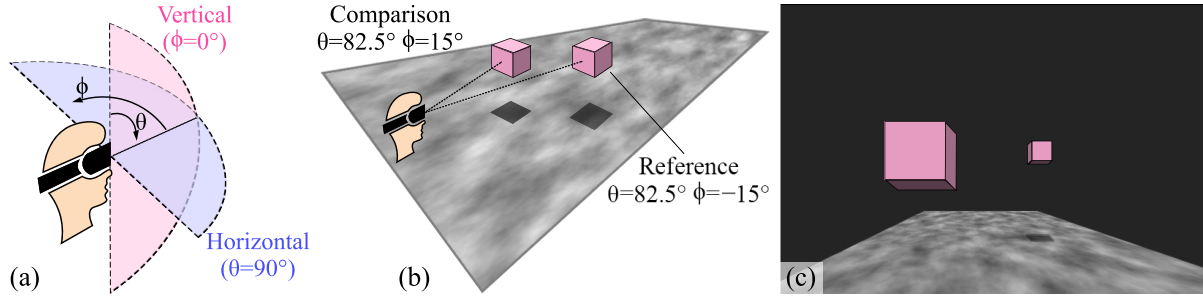
Fig. 3. The spherical coordinate system used in setting up the stimuli (a). An illustration of the stimuli, showing the *reference* and *comparison* cubes, the Perlin-noise textured ground plane, and the cast shadows (b). The actual distance to the *comparison* depends on the *protocol*: for *size matching*, it is positioned very close (75 cm) to the user, while for *distance bisection*, it is placed farther. An illustration of the scene as seen from the subject's perspective (c).

responses (see Wagner [77, p.64]). In Wagner's terms, we used *objective instructions* for the task. Readers are encouraged to scrutinize the instructions video in supplemental archive **S3** for details.

In either protocol, participants had 20 seconds to complete each trial. Participants pressed a button on their controller to confirm their response. The remaining time was always displayed, and if the response was not confirmed in the given time limit, the trial was rejected and excluded from the analysis. Pilot tests without time restrictions showed that ~10–12 s was enough to finish a trial and we therefore chose a conservative time limit of 20 s per trial.

3.5 Procedure

The experiment utilized a 2×2×5 full-factorial within-subject design, with 3 repetitions for each factor combination. For an experiment session, the order of the *protocols* was randomized. Within each protocol, the order of the *eye height* variable was randomized between *adult* and *child*. The ground plane was rendered 170 cm below eye-height for the former, and 50 cm for the latter. Trials for a (*protocol, eye height*) pair composed an experiment "block". Each block consisted of 15 trials in a random order: the five possible *reference distances*—1m, 3m, 5m, 7m, 9m—repeated thrice each. Thus, each experiment session contained 60 trials.

For a short delay of 0.5 s between trials, no stimuli was displayed. Before each block, a text prompt informed the participants of the upcoming condition. The first block for both *protocols* started with 2 additional untimed practice trials. An experiment session could take up to 22 mins in theory, but none took over 15 mins in practice.

Freedom of Movement. As soon as participants pressed a button to start the first block, the position and orientation of their head was stored and used to construct the coordinate system described in § 3.3. During the experiment session, participants could move their head up to 15 cm from this initial position and rotate it up to 30° from the initial orientation. Upon exceeding either limit, the stimuli were hidden and a warning message was displayed, until the participants went back to the allowed range of motion.

3.6 Hypotheses

Based on existing work and pilot tests, we formulated four hypotheses. Prior work has indicated that distance to the bisection point is typically underestimated in VR [3], and we expected the same result (**H1**). For the *size matching* task, our cue-limited situation suggests a mix of objective size matching (perfect matching) and visual size matching leading to size underconstancy [12]. Thus, we expected slight size underconstancy (**H2**). Finally, we expected participants to be more accurate for both *size matching* and *distance bisection* in the *adult eye height*

condition (**H3**). Note that this is in contrast to existing work, where subjects were typically more accurate in judging distance when positioned close to the ground [7]. However, unlike most existing work, we position objects off-the-ground. Our hypothesis **H3** is mainly driven by the expectation that participants will be more accurate in the *adult* condition, which is closer to their real-world experience than *child eye height*.

4 GAMIFICATION METHODOLOGY

Realizing that the experimental tasks for gauging size and depth perception can be rather dry and unengaging, we decided to use game elements in our study. Gamification has been shown to be an effective tool for enhancing user engagement [47] and improving task performance in quantitative studies [74]. It can also help maintain high data-quality in a remote study such as ours [47].

For each trial, participants were awarded scores for their responses, using a strictly concave function over the domain of response (*comparison* size in size matching and distance in bisection), with the maximum at the accurate response (**S1**, Sec. 1.3). Participants' scores over all successful trials in each experiment block were tallied up and displayed at the very end of the session. Note that participants were only shown the aggregate scores for each block, so no feedback was given on individual trials or on whether over- or underestimation happened. Participants were encouraged to complete multiple experiment sessions, up to three. These additional sessions allowed participants to improve their score, and also get bonus points for simply completing the 2nd and 3rd sessions. Thus, instead of 3 judgements for each combination of experiment factors (§ 3.5), we could get up to 9 judgements for each participant. Without appropriate safeguards, repeated sessions could lead to biased judgements as participants could try to reverse-engineer the score. Our methodology eliminated this possibility by a) only providing aggregated feedback in the form of total score for each block of 15 trials and b) limiting each participant to three sessions. Thus, participants' ability to learn the mapping from the score to their responses was severely hindered and no learning effects were expected.

An online leaderboard was maintained where participants could see their (and others') current scores. To keep participants engaged with the study tasks, receive honest judgements, and to encourage the completion of additional experiment sessions, the top ten participants on the leaderboard were provided with an additional US $15 on top of the base compensation of US $5 for participation. Further, advertisements for the study highlighted that it allowed participants to test their size and distance perception in VR and to compete against strangers in perception tasks, thus adding intrinsic motivation as well as a competitive spin to the task.

Study Flow. Participants started by providing their informed consent. They then watched the short instructions video (2.5 min), and downloaded an executable. After starting the application and completing an experiment session, participants filled a demographics survey and provided an introspective report of the cues they utilized for the tasks. To complete additional sessions, participants could simply restart the application. We maintained participant privacy by a) assigning a randomly-generated playful username which was then displayed on the leaderboard, and b) using a hash of unique IDs associated with their device as an authentication mechanism, so no personal information was required (details in **S1**, Sec. 1.4). Participants only provided an email address for gift card delivery. Finally, participants could opt-out of the leaderboard; only one did.

Success of the Methodology. The gamification methodology was an immediate success. The study was able to reach our budgeted target of 60 participants within 5 days. A number of participants completed multiple sessions (21: 3, 10: 2, 29: 1, mean: 1.9), contributing **a total of 113 experimental sessions encompassing 3360 judgements each for size matching and for distance bisection**. Participants commented that they did the study purely out of interest in VR and/or scientific research; for example, they participated to help in "building a foundation of understanding for future technologies", because they thought it was "definitely an interesting thing
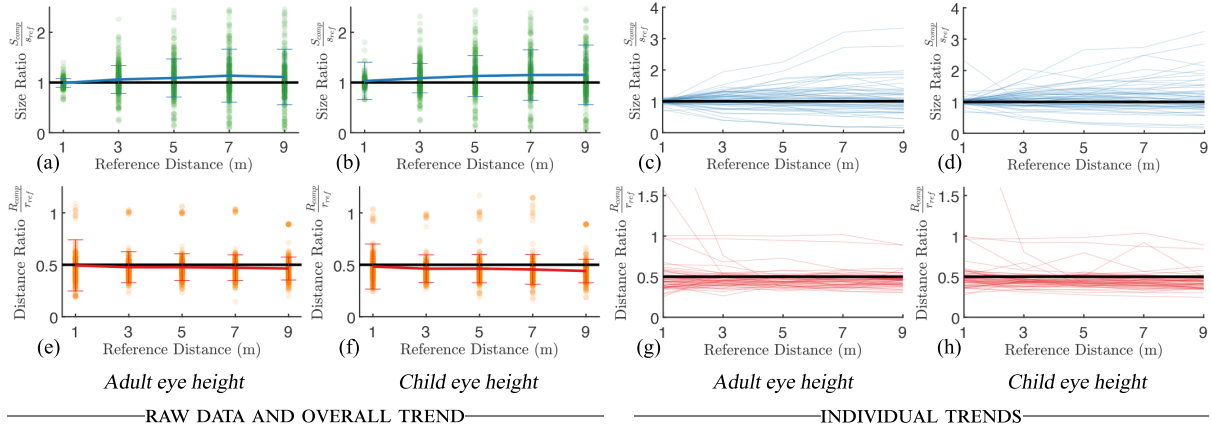
Fig. 4. Results for the *size matching* (top) and *distance bisection* (bottom). Left and center-left columns: all recorded judgements along with the mean and standard deviation over all participants for *adult* and *child* height respectively. Center-right and right: mean trends for individual participants. Accurate judgements are shown in **black**, while mean judgements are shown in **blue**/**red**. Notice that the distance to the bisection point is underestimated near-consistently, but for size matching, trends of both underconstancy and overconstancy are common.

to study", or that it was "more about the research than the money". However, the opportunity to earn additional money was appreciated by some as "a nice bonus".

5 RESULTS

Raw data for all participants has been plotted in Fig. 4a–b for *size matching* and for *distance bisection* tasks in Fig. 4e–f. Note that we denote the participants' responses in uppercase: *comparison* size in *size matching* with S_{comp} and *comparison* distance in *distance bisection* with R_{comp}, to contrast with the initial values s_{comp} and r_{comp} provided to them. We also show the mean (M) and standard deviation (σ) values over all recorded data. The means here have not been reweighted to account for the variation in the number of sessions participants completed. However, discounting judgements from additional sessions reveals very similar results (see **S1**, Sec. 2). Fig. 4c,d,g,h show individual trends for each participant. A very clear trend of underestimating the distance to the bisection point quickly emerges (**H1** validated). For *size matching*, however, the trend is ambiguous (**H2** partially invalidated) but with a slight tendency towards overconstancy. The overall trends were similar for the two *eye heights* for both the tasks (**H3** invalidated).

Data Cleaning and Outlier Removal. Outliers were detected for each of the four *protocol*, *eye height* pairs by comparing the best-fit exponent α^* for each participant with those of other participants. After performing the fitting procedure (see § 5.1), we computed the mean and standard deviation of the natural logarithm of the exponents across all participants. Participants deviating from the mean by a 3σ threshold or higher were considered to be outliers.

Participants whose *distance bisection* exponent was deemed to be an outlier for either of the two *eye heights* were not considered in any of the *distance bisection* analyses. Outliers for *size matching* were analogously removed. This procedure led to the removal of 2 participants each for *distance bisection* and *size matching* analyses. Consequently, for analyses involving both the *protocols*, a total of 4 participants were not considered. In addition, one participant reported hardware issues and their data was completely discounted. Note that Fig. 4 plots have been drawn after this cleaning.

From the remaining data, a very small number of trials were excluded because the participants did not press the confirm button within the time limit: a total of 12 trials (5 for *size matching* and 7 for *distance bisection*) over all the conditions. Thus, 3205 *size matching* trials and 3203 *distance bisection* trials were used for analysis.

Size Matching. The main dependent variable is the size ratio $S_{\text{comp}}/s_{\text{ref}}$ for each trial. Across all distances and eye-heights, the size ratio ($M \pm \sigma$) was 1.094 ± 0.430. To understand the impact of the two main factors of *eye height* and *reference distance*, we ran a repeated measures (RM-) ANOVA on each participant's mean response for every given pair of factor values. *Reference distance* was found to be a significant factor ($F_{4,224} = 4.99, p < .001$), while *eye height* had a noticeable but insignificant effect ($F_{1,56} = 3.76, p = .057$). No interaction effects were found. While *reference distance* was a significant factor, post-hoc analysis using Tukey's Honest Significant Difference (HSD) did not reveal any significant ($p < .05$) pairwise differences. In order to assess if participants' answers were more self-consistent for certain factor levels than others, we also performed an RM-ANOVA on the standard deviation of the size ratio, which found only *reference distance* to be a significant factor ($F_{4,224} = 17.61, p < .001$). Tukey's HSD indicated that the distance pairs (1m,3m), (1m,5m), (1m,7m), (1m,9m), (3m,7m), (3m,9m), and (5m,9m) were significantly different.

Another effect we analyzed was "inertia"—the tendency of initial size s_{comp} to impact the participant response S_{comp}. A 1-way ANOVA found that it had a significant effect indeed: in the *adult eye height*, the final size ratios were significantly larger when the initial *comparison* size was larger than the *reference*; $s_{\text{comp}} > s_{\text{ref}}$: 1.132 ± 0.439, $s_{\text{comp}} < s_{\text{ref}}$: 1.021 ± 0.361 ($F_{1,1601} = 30.6, p < .001$). The same effect was noticed for *child eye height* trials: 1.146 ± 0.481 vs. 1.073 ± 0.410 ($F_{1,1600} = 10.6, p = .001$).

Distance Bisection. The main independent factor for *distance bisection* was the distance ratio $R_{\text{comp}}/r_{\text{ref}}$ whose value over all trials was 0.469 ± 0.157. An RM-ANOVA revealed that no factor significantly affected the mean distance ratio ($p > .05$). However, a significant effect of *reference distance* on the standard deviation was observed ($F_{4,224} = 3.72, p = .006$). Post-hoc comparisons using Tukey's HSD did not show any significant differences. As compared to the *size matching* trials, the impact of "inertia" here was small. Specifically, the initial comparison distance d_{comp} did not have a significant effect on participants' responses in the *adult eye height* condition, it had a small but significant effect in the *child eye height* condition—distance ratio was 0.468 ± 0.142 when $d_{\text{comp}} > d_{\text{ref}}/2$ and 0.453 ± 0.163 when not ($F_{1,1598} = 3.91, p = .048$).

Learning Effects. To confirm that our minimal aggregated feedback strategy prevented bias, we tested if participants with multiple sessions improved over time. Using paired *t*-tests for participants with two or three sessions, we observed no significant difference ($p > .10$) between the mean size ratios (*size matching*) or mean distance ratios (*distance bisection*) across sessions. That is, **no learning effects** were observed, as expected.

5.1 Fitting Curves of Perceived Distance

The size-distance invariance hypothesis (SDIH) has been used in prior work to derive perceived depth judgements from judgements of perceived size [29, 30]. In our *size matching* task, participants use perceptual size matching to indicate the perceived size of the *reference* by setting the *comparison* size S_{comp}. The SDIH states that the ratio of perceived size $\mathcal{S}$ to perceived egocentric distance $\mathcal{R}$ is fixed: $\mathcal{S}/\mathcal{R} = c^{\text{p}}$.

Given the actual physical size S and egocentric distance R for an object then, we can use the fact that their ratio is also fixed ($S/R = c^{\text{a}}$) to invert a given judgement of perceived size, to get the perceived distance as $\mathcal{R} = cR\frac{\mathcal{S}}{S}$, where $c = c^{\text{a}}/c^{\text{p}}$ is a constant modelling the ratio of an object's angular size to its perceived angular size. We assume that angular size is perceived linearly, and thus c is an stimulus-independent constant.

In order to compare the perceived distance implied by the *distance bisection* and the *size matching* judgements, we can fit curves with the same representational degrees of freedom to trials from both and test if the parameters

are similar. Similar to prior work [76], we model the perceived distance $\mathcal{R}$ as an exponential function of actual distance R, along with a linear scaling term. That is,

$$\mathcal{R} = kR^{\alpha} \text{ with } k, \alpha \geq 0. \tag{2}$$

In *size matching* judgements, a participant indicates that the perceived *reference* size $\mathcal{S}_{\text{ref}}$ and comparison size $\mathcal{S}_{\text{comp}}$ are equal, implying that

$$c^{\text{p}}_{\text{ref}}\mathcal{R}_{\text{ref}} = c^{\text{p}}_{\text{comp}}\mathcal{R}_{\text{comp}}.$$

Multiplying both sides by c, we get $c^{\text{a}}_{\text{ref}}\mathcal{R}_{\text{ref}} = c^{\text{a}}_{\text{comp}}\mathcal{R}_{\text{comp}}$, which further expands to

$$\mathcal{R}_{\text{ref}}\frac{s_{\text{ref}}}{r_{\text{ref}}} = \mathcal{R}_{\text{comp}}\frac{S_{\text{comp}}}{R_{\text{comp}}}.$$

Finally, substituting the perceived distances from Eq. 2 and rearranging, we get the ratio of perceived distance of the *comparison* object to that of the *reference*,

$$\frac{k(R_{\text{comp}})^{\alpha}}{k(r_{\text{ref}})^{\alpha}} = \frac{s_{\text{ref}}}{r_{\text{ref}}}\frac{R_{\text{comp}}}{S_{\text{comp}}}. \tag{3}$$

For accurate judgements implying an invariant visual space with no foreshortening or anti-foreshortening, the above equation should hold. Therefore, the best-fit value for α for *size matching*, α^*_{SM}, is obtained by minimizing the deviation from Eq. 3.

$$\alpha^*_{SM} = \underset{\alpha \in (0,\infty)}{\arg\min} \sum_i \left\| \frac{s^i_{\text{ref}}}{r^i_{\text{ref}}}\frac{R^i_{\text{comp}}}{S^i_{\text{comp}}} - \frac{(R^i_{\text{comp}})^{\alpha}}{(r^i_{\text{ref}})^{\alpha}} \right\|^2, \tag{4}$$

where the summation is over a participant's *size matching* trials. Fitting a similar curve for the *distance bisection* trials is relatively straightforward.

$$\alpha^*_{DB} = \underset{\alpha \in (0,\infty)}{\arg\min} \sum_i \left\| 1 - \frac{\mathcal{R}^i}{k(R^i)^{\alpha}} \right\|^2 = \sum_i \left\| 1 - \frac{2(R^i_{\text{comp}})^{\alpha}}{(r^i_{\text{ref}})^{\alpha}} \right\|^2. \tag{5}$$

Intuitively, note that if the curve models a participant's responses with full accuracy, then $\mathcal{R}^i = k(R^i)^{\alpha}$ for all i, and the residual is zero. The minimization in Eq. (4) penalizes relative deviation of the participant's response from the perceived distance modelled by the curve. Eq. (5) is similar, except that we want the modelled perceived distance to the *reference* to be twice that of the *comparison*. Both the minimizations are non-linear least squares problems and are efficiently solved using a trust-region based approach [6] available in MATLAB. We initialize with the trivial value $\alpha^* = 1$ for both.

Fig. 1c shows the curve-fitting summary results for all conditions. The α^* found via *size matching* judgements (α^*_{SM}) for *adult eye height* was 0.991 ± 0.185 ($M \pm \sigma$), and for *child eye height*, it was 1.012 ± 0.188. For *distance bisection* (α^*_{DB}), it was 0.862 ± 0.219 for *adult* and 0.863 ± 0.238 for *child* eye heights. Note that exponents are fit for each participant before taking the mean and standard deviation over the participants, thus weighing single-session and multiple-session participants equally. As already suggested by analysis of the raw data, *distance bisection* judgements showed a clear trend towards compression of perceptual distance space, while *size matching* data indicated a mix of compression and expansion. An RM-ANOVA with dependent variable α^* and independent variables *protocol* and *eye height* found that while *protocol* had a significant effect ($F_{1,54} = 15.22, p < .001$), *eye height* did not ($F_{1,54} = 0.007, p = .965$).

5.2 Random Factors Effects

As noted in § 3.1, we collected data on a number of random factors: gender, age range, height[1], device used, 3D games experience, and 3D modelling experience. We also extracted resolution and FoV factors from device specifications. For device-independent online studies to be successful and useful for perceptual experiments, the effect of device specifications should be minimal. Therefore, we analyze the effect of each factor in isolation for the maximum statistical power, allowing us to detect and report even small deviations from the null hypotheses. For nominals (device, gender, age range), this is a 1-way ANOVA with the mean values of α_{DB}^* and α_{SM}^* for each participant as the dependents and with the random factor in consideration the only independent factor. First considering device usage, while 11 different devices were used, we only consider the ones used by ≥ 5 participants: Valve Index (17), Oculus Quest (15), and HTC Vive (10). 1-way ANOVAs showed that neither device used nor age range were significant factors for α_{DB}^*, but age had a significant effect on α_{SM}^* ($F_{3,53} = 3.46, p = .023$). Tukey's HSD revealed a pairwise difference between participants aged below 25 and those between 36 to 45 years old. Impact of gender was not analyzed due to our heavily-skewed sample.

Similarly, for numeric variables (height, resolution, and FoV) and Likert scale ordinals treated as numeric for analysis (3D games and modelling experience), we computed the correlation coefficient between each variable and the exponents α_{SM}^* and α_{DB}^*, but no significant correlations were found. This is an encouraging result, indicating that an *in-the-wild* strategy is viable for spatial perception experiments such as ours.

5.3 Quantitative Analysis of Participant Strategies

The use of 6-DoF HMDs allowed us to track participants' head position and orientation. This data can help us understand participants' strategies for judging size/distance and answer questions such as "Do they move and rotate their head a lot to use motion parallax cues?", "Do they look towards the ground, potentially utilizing shadows as cues?", and "Do they respond quickly, or not?". To understand whether participants changed their strategies to account with changes in experiment factors, we perform RM-ANOVA with head-tracking data as dependent and experiment factors as independent variables. For simplicity, interaction effects were not included.

Data Collection. Participants' head position and orientation were recorded every 200ms. Over each trial, the difference between consecutive position samples (in $\mathbb{R}^3$) and consecutive orientation samples (in terms of angle between them) were added up to estimate the total head translation and rotation. For each participant and factor values, these two datapoints were then averaged over all associated trials to get the *head translation* and *head rotation* variables. In order to understand the role of shadows, we also compute *looking down percentage*: the percentage of orientation samples where the cyclopean line of sight was at least $5°$ below horizontal. We also show how the *time taken* by the participant to record a response was impacted by experiment factors.
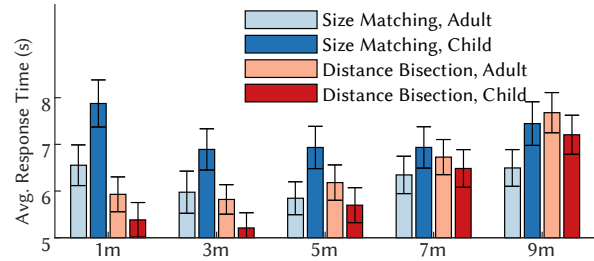

Fig. 5. Average response time (with std. error) across the experiment factors. Note that the **y-axis does not start at zero**. Notice the large differences across the *protocols* and *eye heights* at 1m and 3m *reference* distances, which shrink when the *reference* is farther.

---

[1]Height was solicited via email after the study: 45 participants responded.

Change in Strategy with Experiment Factors. RM-ANOVA indicated that the *head translation* was significantly affected by the *eye height* ($F_{1,53} = 4.56, p = .037$)[2], with participants moving their heads significantly more in the *adult eye height* condition (15.5 ± 22.5 cm) as compared to *child* (13.8 ± 20.9 cm).

Head rotation was similarly impacted by *eye height* ($F_{1,53} = 6.33, p = .015$), with *adult* (19.8° ± 24.6°) showing significantly more rotation as compared to *child* (16.6° ± 20.8°). Interestingly, *looking down percentage* was significantly affected by the *protocol* ($F_{1,53} = 23.28, p < .001$) as well as *eye height* ($F_{1,53} = 4.82, p = .033$). Post-hoc comparisons showed that participants looked down much more in the *distance bisection protocol* (29.7 ± 35.9%) as compared to *size matching* (9.0 ± 18.1%), and in the *adult eye height* condition (21.8 ± 31.8%) as compared to *child* (17.0 ± 28.4%). Note that the effect sizes are larger for the difference between *protocols*. The *time taken* measure saw a significant effect of *reference distance* ($F_{4,216} = 20.9, p < .001$). Participants took the most time of 7.21 s for the 9m distance, as expected, but the least time was taken for the 3m distance: 5.97 s. *Time taken* for the 9m distance significantly differed from other distance values, and a significant difference was also observed between (1m, 3m), (3m, 5m), and (5m, 7m) pairs. Fig. 5 summarizes the results for the *time taken* variable.

Impact of Strategy on Perceptual Judgements. In order to understand the impact of participants' strategy (encoded in the variables described above) on their perceptual judgements, we compute correlations of these variables with the distance perception exponent α^* for all four combinations of *protocol* and *eye height*. Multiple significant but weak correlations were found for α^*_{DB} for *adult eye height*—with *head translation* ($r(54) = .29, p = .030$), *head rotation* ($r(54) = .29, p = .030$), and *looking down percentage* ($r(54) = .27, p = .044$).

Lastly, we had asked participants for introspective reports of how they judged size/distance in each condition.

Participants' Introspective Reports. We sought introspective reports from participants in two forms: freeform and structured. In order to not bias participants with our own presumptions, we first asked them for freeform feedback on the cues they used in both the *protocols*, and if the *eye height* impacted the cues they used. In a later section of the survey form (see supplemental **S3**), participants answered structured (multiple-choice) questions on the cues utilized in the two *protocols*.

Participants were surprisingly accurate in certain aspects. For example, those who mentioned the use of motion parallax in the structured questions exhibited much larger *head translation*: mean value was 21.6 cm vs. 11.0 cm for *size matching* and 17.5 cm vs. 10.7 cm for *distance bisection*. An ANOVA found the former distinction to be insignificant ($F_{1,55} = 3.00, p = .089$), but the latter significant ($F_{1,55} = 6.26, p = .015$). *Head rotation* showed a similar trend—*size matching*: 21.7° vs. 14.1°, *distance bisection*: 24.2° vs. 14.7°. The former was again found to be statistically insignificant ($F_{1,55} = 2.79, p = .100$) but the latter was significant ($F_{1,55} = 5.69, p = .021$). Similarly, participants who reported using shadows as a cue were more likely to spend time *looking down*: 13.3% vs. 6.6% ($F_{1,55} = 2.66, p = .108$) for *size matching* and 35.9% vs. 21.5% ($F_{1,55} = 2.30, p = .135$) for *distance bisection*.

The structured reports helped identify a few other interesting cue usage patterns. For instance, many participants tried to measure absolute egocentric size/distance in standard units (*SM*: 25, *DB*: 25), and a smaller but still significant number used relative size/distance cues such as their own body parts (*SM*: 16, *DB*: 6). A number of participants (*SM*: 12, *DB*: 16) felt that the speed of the continuous change in size/distance of the *comparison* was a useful cue. This is in spite of our efforts to minimize the usefulness of this cue: we used an acceleration term so that the speed was a function of the amount of time the participant held the thumbstick (used to increase or decrease the size/distance) in place.

Finally, subjective analysis of participants' freeform reports revealed some interesting insights. 34 out of 55 participants (whose data was used for both *protocols*) reported that the *eye height* did not influence their judgements at all, while another 4 felt that it had minimal impact. Most (11) of the participants who felt that it affected their judgements noted that the shadows were harder to see or otherwise different in one of the *eye*

---

[2]One participant's head-tracking data was lost to a data conversion bug.

heights. Surprisingly, very few (4) noted any additional difficulty in the *child height* condition they had little experience with in real life. One participant commented that they were "very used to not looking at the floor for the sake of avoiding motion sickness", so the *eye height* did not influence them.

Some participants relied on prior experience with VR games to estimate physical sizes. One said, "I would think back to all the times I see [sic] an object in the distance and come close from other games I have played.", while another "channeled a couple of years of VR Beat Saber and Synth Riders as those involve basically constant estimation of distance." Other participants mentioned that the ability to continuously move the *comparison* back and forth was helpful for *distance bisection*: "I moved the controlled cube back and forth to try and understand how far away the reference cube was…". One participant talked about using the two monocular views separately to get a better estimate of depth: "I just tried to compare the sizes while alternating which eye was open."

6 DISCUSSION

We discuss the implications of our work on visual perception in VR, and on remote perceptual experiments.

6.1 Implications on VR Perception Understanding

Our size perception experiment used a perceptual matching protocol inspired by Holway and Boring's classic experiment [28]. While our results were similar to the overall trend of slight size overconstancy they noted, we observed more variation and ambiguity—with individuals exhibiting both underconstancy and overconstancy. Other classic experiments [14] and more recent VR experiments [54] have studied restricted viewing conditions. Aiming for ecological validity, we employed scenarios closer to real-world usage: binocular viewing, relatively unrestricted head movement, lighting, and a textured ground plane. In this aspect, the experiments closest to ours might be Kenyon et al. [38], Hornsey et al. [29], and Kelly et al. [37].

Kenyon et al. [38] looked at size constancy in a CAVE environment, asking participants to resize a far-away familiar object while the real object was also placed near them, and observed that participants made the distant comparison larger than the reference across the range of tested distances (0.6–2.4m). The other two studies [29, 37] used modern HMDs and a familiar-sized object that the participants could touch by hand for reference. They also observed a general trend of size underconstancy. Note that our protocol is reversed: the reference is far away while the participant-controlled comparison is nearby. Therefore, while the previous studies observed size-underconstancy in VR, we found a slight overconstancy trend. One possible explanation to this contrary effect may be the shape of the stimulus. Wagner's survey on real-world size constancy [77] points out that size perception may be related to object orientation: frontally-oriented (perpendicular to the floor and facing the participant) stimuli were commonly associated with overconstancy and flat (parallel to the floor) stimuli with underconstancy. In comparison to the spherical [29, 37] and cylindrical [38] objects used in prior VR studies, our cube-shaped stimuli have a more noticeable frontal face.

Wagner [77] notes that an important factor in size perception experiments is the instruction provided to the subjects. While we aimed for *objective instructions*—asking participants to match the physical (3D) size, some participants may still have been influenced by the apparent, or retinal, size, leading to size underconstancy. On the contrary, some may have overestimated the perspective foreshortening effect, leading to overconstancy.

Our bisection protocol is similar to bisection studies in the real world [3, 42] as well as those in virtual environments [3]. While Lappin et al. [42] noted an anomalous report of overestimation of the distance to the bisection point, most studies (especially in VR) report underestimation, similar to ours. Moreover, egocentric distance judgements made using very different protocols have typically suggested distance underconstancy as well [61, survey]. The latter lends further cadence to the values of α_{DB}^* our curve-fitting suggested: in either *eye height*, bisection judgements suggested a shrinkage of perceived distance towards the observer.

Similar to our exponential curve fitting, Murgia and Sharkey [51] reported mean exponent values of 0.830 in an environment with poor cues, and 0.975 with one with rich cues, albeit in a CAVE system. Our environment rendering falls somewhere between those two: we render a floor, but no walls. Thus, the exponent 0.896 we observed in our bisection protocol is more similar to their results, than the value 1.083 suggested in the size matching protocol. This dichotomy is especially interesting since their results were obtained via a perceptual size matching protocol.

While our experiments do not intend to test the validity of SDIH, we observed a weak correlation between α^*_{DB} and α^*_{SM}, which is worthwhile to put in context of Ross's investigation of SDIH alternatives [63]. Of particular interest is his suggestion of a "perceptual SDIH", which breaks the classical SDIH assumption of accurately perceived angular size. In a future experiment, Wagner's *angular size instruction* can be applied to gauge the accuracy of angular size perception, by using a perceptual matching protocol similar to Kaneko and Uchikawa [36].

Contrary to our initial hypothesis, comparisons across the two eye heights did not reveal any noticable differences. This is in contrast to Corujeira and Ian's study [7] where observers were *better* at judging egocentric distances in the lower eye-height (20 cm from the floor) condition. However, this comparison should be taken with a grain of salt: they position objects on the ground and define distance as the walking distance to the percept. Thus, lower the eye-height, better the correspondence between the 3D egocentric distance and the walking distance. Leyrer et al. [43], who studied virtually altered eye heights in VR similar to us, also found that the foreshortening effect was smaller when eye height was closer to the ground plane, but they also position stimuli on the ground and define distance as the 2D distance projected to the floor.

Finally, our quantitative observation of the minimal impact of head motion on size/distance judgements in VR confirms Hornsey et al.'s informal observation [29]. We further discuss Hornsey et al. and other relatively-uncontrolled perceptual experiments below.

6.2 Conducting Remote Perceptual Studies

Our remote unsupervised methodology allowed us to quickly recruit 60 participants within 5 days. Judging by comments and feedback on our advertisements, the most important source of participants was Reddit. VR-focussed communities on Reddit showed great enthusiasm for the study, reconfirming previous work hailing the platform as a useful community for recruiting participants [65]. To engage participants, we used a combination of gamification and intrinsic as well as extrinsic motivation. Specifically, participants were told that the study allowed them to test their size and depth perception abilities in VR and contribute to science (*intrinsic motivation*), while at the same time receiving a monetary compensation (*extrinsic motivation*), and getting to see their score on a leaderboard with a chance to get a higher compensation (*gamification* and *extrinsic motivation*).

We believe that the gamification elements ensured increased user engagement, resulting in high-quality data in our experiment. Furthermore, we did not see any evidence for malicious behaviour. This is in contrast to microtask-based crowdsourced studies on platforms such as Mechanical Turk, where testing for malicious users is commonplace [67]. Prior work has also shown that dropping monetary compensation altogether still allows successful unsupervised studies with a large number of participants [31]. Outside of perception research, many well-known scientific projects have also used gamification and intrinsic motivation to engage the target audience [8, 32]. This combination is an exciting exploration avenue for future VR perception research.

Our study included participants from 19 different countries and a large variety of HMDs. While this diversity is encouraging, access to 6-DoF VR devices continues to be dominated by men in rich countries. This demographics bias may have affected some of our results. We hope that this bias in VR ownership will go down with lowering hardware costs democratizing access in low- and medium-income countries, and that research on the unique challenges faced by female VR users [68] reduces the gender bias in the future. Some bias towards participants in the US and Canada may also have occurred due to better availability of many VR devices in the region and the

fact that we posted our recruitment calls in English-language forums. However, even with the currently skewed VR ownership demographics, remote VR studies are still an invaluable tool for perception research and offer researchers access to a more diverse audience than their immediate local community.

Finally, the variety of headsets used by our participants allowed us to test the impact of the device on perceptual judgements. Our results echo the partially-uncontrolled experiment of Hornsey et al. [29], who tested size and shape judgements with two headsets and did not observe any significant differences between the two. However, it is unclear if this generalizes to perceptual judgements unexplored by our or their work. Previous work has noted a significant effect of hardware factors such as FoV [35] and IPD [4]. Further exploration is needed to understand which perceptual studies may not be amenable to uncontrolled methodologies such as ours.

7 CONCLUSION AND FUTURE WORK

We conducted a remote, fully unsupervised, perceptual study to understand size and distance perception in Virtual Reality. To recruit study participants and to keep them engaged through the monotonous tasks, we combined an extrinsic monetary reward with gamification and an appeal to their intrinsic motivation. Ours is the first fully-uncontrolled VR distance/size perception study, aimed at ecological validity by collecting data from a diverse set of participants and VR hardware. Unfortunately, current VR ownership trends reduced participant diversity in some aspects, most notably gender.

We compare the results obtained via a distance bisection protocol with distance perception implied by a size matching protocol via the size-distance invariance hypothesis (SDIH). While the former indicated a clear trend of distance underestimation, the latter showed a very weak overestimation trend. We also investigated the influence of a rendered floor, simulating different eye heights, but noticed no significant effect. We distilled participants' judgements into exponential curves of perceived egocentric distance expressed as a function of actual distance. Such curves can be used to manipulate virtual worlds to improve task performance in performance-critical applications such as VR flight simulators and industrial training as well as in 3D judgement–heavy games (Fig. 1d).

To aid future research, we will **release the source code for our experiment and analysis, as well as anonymized data**. The study executable is included (see **S3**) for reviewers to test and analyze.

Potential Applications and Future Work. In the future, we would like to systematically explore the gamification aspects of our work, modifying leaderboard size and monetary rewards parameters, and exploring a deeper coupling of gamification with the perceptual judgement tasks. Directly comparing intrinsic motivations with microtasks is also an interesting avenue. On the other hand, we are also inspired by recent work by Huber and Gajos [31] that advocates for conducting large scale studies without any monetary rewards.

Finally, we are very interested in applying the perceptual distance curves to real-world applications. Perceptual judgement tasks, such as distance bisection, could be used as a "pre-calibration" step for modifying the rendered world in performance-critical applications, similar to gamma calibration on monitors. In games, such a calibration can be utilized for improving performance as well, or a sneaky way to make the game more difficult for users by applying the inverse transformation to the rendered distance.

REFERENCES

[1] A. C. Beall, J. M. Loomis, J. W. Philbeck, and T. G. Fikes. Absolute motion parallax weakly determines visual scale in real and virtual environments. In B. E. Rogowitz and J. P. Allebach, eds., *Human Vision, Visual Processing, and Digital Display VI*, vol. 2411, pp. 288 – 297. International Society for Optics and Photonics, SPIE, 1995. doi: 10.1117/12.207547

[2] G. P. Bingham, A. Bradley, M. Bailey, and R. Vinner. Accommodation, occlusion, and disparity matching are used to guide reaching: A comparison of actual versus virtual environments. *Journal of experimental psychology: human perception and performance*, 27(6):1314, 2001.

[3] B. Bodenheimer, J. Meng, H. Wu, G. Narasimham, B. Rump, T. P. McNamara, T. H. Carr, and J. J. Rieser. Distance estimation in virtual and real environments using bisection. In *Proceedings of the 4th Symposium on Applied Perception in Graphics and Visualization*, pp. 35–40, 2007.

[4] G. Bruder, A. Pusch, and F. Steinicke. Analyzing effects of geometric rendering parameters on size and distance estimation in on-axis stereographics. In *Proceedings of the ACM Symposium on Applied Perception*, pp. 111–118, 2012.

[5] G. Bruder, F. Steinicke, P. Wieland, and M. Lappe. Tuning self-motion perception in virtual reality with visual illusions. *IEEE Transactions on Visualization and Computer Graphics*, 18(7):1068–1078, 2011.

[6] T. F. Coleman and Y. Li. An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on optimization*, 6(2):418–445, 1996.

[7] J. G. P. Corujeira and I. Oakley. Stereoscopic egocentric distance perception: The impact of eye height and display devices. In *Proceedings of the ACM Symposium on Applied Perception*, SAP '13, p. 23–30. Association for Computing Machinery, New York, NY, USA, 2013. doi: 10.1145/2492494.2492509

[8] V. Curtis. Motivation to participate in an online citizen science game: A study of foldit. *Science Communication*, 37(6):723–746, 2015.

[9] J. Cutting. *Reconceiving perceptual space*, pp. 215–238. MIT Press, Cambridge, MA, USA, 01 2003.

[10] R. Day and T. E. Parks. To exorcize a ghost from the perceptual machine. *The moon illusion*, pp. 343–350, 1989.

[11] D. W. Eby and M. L. Braunstein. The perceptual flattening of three-dimensional scenes enclosed by a frame. *Perception*, 24(9):981–993, 1995.

[12] R. G. Eggleston, W. P. Janson, and K. A. Aldrich. Virtual reality system effects on size-distance judgements in a virtual environment. In *Proceedings of the IEEE 1996 Virtual Reality Annual International Symposium*, pp. 139–146. IEEE, 1996.

[13] W. Epstein. Attitudes of judgment and the size-distance invariance hypothesis. *Journal of experimental Psychology*, 66(1):78, 1963.

[14] W. Epstein and A. A. Landauer. Size and distance judgments under reduced conditions of viewing. *Perception & Psychophysics*, 6(5):269–272, 1969.

[15] W. Epstein, J. Park, and A. Casey. The current status of the size-distance hypotheses. *Psychological Bulletin*, 58(6):491, 1961.

[16] J. M. Foley, N. P. Ribeiro-Filho, and J. A. Da Silva. Visual perception of extent and the geometry of visual space. *Vision Research*, 44(2):147–156, 2004.

[17] C. Galais and E. Anduiza. "You cheated on me!" Causes and consequences of cheating in online surveys. In *Visions in Methodology (VIM)*. McMaster University, Hamilton, Canada, 2014.

[18] H. Gehlbach, G. Marietta, A. M. King, C. Karutz, J. N. Bailenson, and C. Dede. Many ways to walk a mile in another's moccasins: Type of social perspective taking and its effect on negotiation outcomes. *Computers in Human Behavior*, 52:523–532, 2015.

[19] A. S. Gilinsky. Perceived size and distance in visual space. *Psychological review*, 58(6):460, 1951.

[20] A. Glennerster, L. Tcheang, S. J. Gilson, A. W. Fitzgibbon, and A. J. Parker. Humans ignore motion and stereo cues in favor of a fictional stable world. *Current Biology*, 16(4):428 – 432, 2006. doi: 10.1016/j.cub.2006.01.019

[21] W. Gogel and J. A. Da Silva. Familiar size and the theory of off-sized perceptions. *Perception & psychophysics*, 41:318–28, 05 1987. doi: 10.3758/BF03208233

[22] S. J. Gould, A. L. Cox, and D. P. Brumby. Diminished control in crowdsourcing: an investigation of crowdworker multitasking behavior. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 23(3):1–29, 2016.

[23] H. E. Gruber. The relation of perceived size to perceived distance. *The American Journal of Psychology*, 67(3):411–426, 1954.

[24] R. N. Haber and C. A. Levin. The independence of size perception and distance perception. *Perception & psychophysics*, 63(7):1140–1152, 2001.

[25] K. Hata, R. Krishna, L. Fei-Fei, and M. S. Bernstein. A glimpse far into the future: Understanding long-term crowd worker quality. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pp. 889–901, 2017.

[26] M. Hirth, T. Hoßfeld, and P. Tran-Gia. Cost-optimal validation mechanisms and cheat-detection for crowdsourcing platforms. In *2011 fifth international conference on innovative mobile and internet services in ubiquitous computing*, pp. 316–321. IEEE, 2011.

[27] D. M. Hoffman, A. R. Girshick, K. Akeley, and M. S. Banks. Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of vision*, 8(3):33–33, 2008.

[28] A. H. Holway and E. G. Boring. Determinants of apparent visual size with distance variant. *The American Journal of Psychology*, 54(1):21–37, 1941.

[29] R. L. Hornsey, P. B. Hibbard, and P. Scarfe. Size and shape constancy in consumer virtual reality. *Behavior research methods*, 52(4):1587—1598, August 2020. doi: 10.3758/s13428-019-01336-9

[30] I. Howard. *Perceiving in Depth. Volume 3. Other mechanisms of depth perception.* Oxford University Press, 03 2012.

[31] B. Huber and K. Z. Gajos. Conducting online virtual environment experiments with uncompensated, unsupervised samples. *Plos one*, 15(1):e0227629, 2020.

[32] D. Huynh, L. Zuo, and H. Iida. Analyzing gamification of "duolingo" with focus on its course structure. In *International Conference on Games and Learning Alliance*, pp. 268–277. Springer, 2016.

[33] N. Jenkin. Effects of varied distance on short-range size judgments. *Journal of experimental psychology*, 54(5):327, 1957.

[34] N. Jenkin. A relationship between increments of distance and estimates of objective size. *The American Journal of Psychology*, 72(3):345–363, 1959.

[35] J. A. Jones, E. A. Suma, D. M. Krum, and M. Bolas. Comparability of narrow and wide field-of-view head-mounted displays for medium-field distance judgments. In *Proceedings of the ACM Symposium on Applied Perception*, pp. 119–119, 2012.

[36] H. Kaneko and K. Uchikawa. Perceived angular and linear size: the role of binocular disparity and visual surround. *Perception*, 26(1):17–27, 1997.

[37] J. W. Kelly, L. A. Cherep, and Z. D. Siegel. Perceived space in the HTC Vive. *ACM Transactions on Applied Perception (TAP)*, 15(1):1–16, 2017.

[38] R. V. Kenyon, D. Sandin, R. C. Smith, R. Pawlicki, and T. Defanti. Size-constancy in the cave. *Presence: Teleoperators and Virtual Environments*, 16(2):172–187, 2007.

[39] S. Komarov, K. Reinecke, and K. Z. Gajos. Crowdsourcing performance evaluations of user interfaces. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 207–216, 2013.

[40] Y. Koyama, D. Sakamoto, and T. Igarashi. Crowd-powered parameter analysis for visual design exploration. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, UIST '14, p. 65–74. Association for Computing Machinery, New York, NY, USA, 2014. doi: 10.1145/2642918.2647386

[41] Y. Koyama, I. Sato, D. Sakamoto, and T. Igarashi. Sequential line search for efficient visual design optimization by crowds. *ACM Trans. Graph.*, 36(4), July 2017. doi: 10.1145/3072959.3073598

[42] J. Lappin, A. Shelton, and J. Rieser. Environmental context influences visually perceived distance. *Perception & psychophysics*, 68:571–81, 06 2006. doi: 10.3758/BF03208759

[43] M. Leyrer, S. A. Linkenauger, H. H. Bülthoff, U. Kloos, and B. Mohler. The influence of eye height and avatars on egocentric distance estimates in immersive virtual environments. In *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization*, APGV '11, p. 67–74. Association for Computing Machinery, New York, NY, USA, 2011. doi: 10.1145/2077451.2077464

[44] J. Loomis and J. Knapp. Visual perception of egocentric distance in real and virtual environments. In L. Hettinger and M. Hass, eds., *Virtual and adaptive environments: Applications, implications, and human performance issues*, pp. 21–46. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, USA, 2003.

[45] J. M. Loomis, J. A. Da Silva, N. Fujita, and S. S. Fukusima. Visual space perception and visually directed action. *Journal of Experimental Psychology: Human Perception and Performance*, 18(4):906, 1992.

[46] J. M. Loomis, J. A. Da Silva, J. W. Philbeck, and S. S. Fukusima. Visual perception of location and distance. *Current Directions in Psychological Science*, 5(3):72–77, 1996.

[47] J. Looyestyn, J. Kernot, K. Boshoff, J. Ryan, S. Edney, and C. Maher. Does gamification increase engagement with online programs? A systematic review. *PloS one*, 12(3):e0173403, 2017.

[48] X. Ma, M. Cackett, L. Park, E. Chien, and M. Naaman. Web-Based VR Experiments Powered by the Crowd. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, p. 33–43. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2018. doi: 10.1145/3178876.3186034

[49] W. Mason and S. Suri. Conducting behavioral research on amazon's mechanical turk. *Behavior research methods*, 44(1):1–23, 2012.

[50] A. Mottelson and K. Hornbæk. Virtual reality studies outside the laboratory. In *Proceedings of the 23rd acm symposium on virtual reality software and technology*, pp. 1–10, 2017.

[51] A. Murgia and P. M. Sharkey. Estimation of distances in virtual environments using size constancy. *International Journal of Virtual Reality*, 8(1):67–74, 2009.

[52] A. Naceri, A. Moscatelli, and R. Chellali. Depth discrimination of constant angular size stimuli in action space: role of accommodation and convergence cues. *Frontiers in Human Neuroscience*, 9:511, 2015. doi: 10.3389/fnhum.2015.00511

[53] J. F. Norman, J. T. Todd, V. J. Perotti, and J. S. Tittle. The visual perception of three-dimensional length. *Journal of Experimental Psychology: Human Perception and Performance*, 22(1):173, 1996.

[54] E. Peillard, T. Thebaud, J. Normand, F. Argelaguet, G. Moreau, and A. Lécuyer. Virtual objects look farther on the sides: The anisotropy of distance perception in virtual reality. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 227–236, 2019.

[55] K. Perlin. Improving noise. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pp. 681–682. ACM, New York, NY, USA, 2002.

[56] J. W. Philbeck and J. M. Loomis. Comparison of two indicators of perceived egocentric distance under full-cue and reduced-cue conditions. *Journal of Experimental Psychology: Human Perception and Performance*, 23(1):72, 1997.

[57] J. Purdy and E. J. Gibson. Distance judgment by the method of fractionation. *Journal of Experimental Psychology*, 50(6):374, 1955.

[58] K. Ragozin, K. Kunze, K. Marky, and Y. S. Pai. MazeRunVR: An Open Benchmark for VR Locomotion Performance, Preference and Sickness in the Wild. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–8, 2020.

[59] Rakuten Intelligence. The brief blog: Mobile dominates the online reality of virtual reality sales. https://www.rakutenintelligence.com/blog/2016/virtual-reality-mostly-mobile, 2019.

[60] K. Reinecke and K. Z. Gajos. LabintheWild: Conducting large-scale online experiments with uncompensated samples. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, pp. 1364–1378, 2015.

[61] R. S. Renner, B. M. Velichkovsky, and J. R. Helmert. The perception of egocentric distances in virtual environments-a review. *ACM Computing Surveys (CSUR)*, 46(2):1–40, 2013.

[62] J. J. Rieser, D. H. Ashmead, C. R. Talor, and G. A. Youngquist. Visual perception and the guidance of locomotion without vision to previously seen targets. *Perception*, 19(5):675–689, 1990.

[63] H. E. Ross. Levels of processing in the size-distance paradox. In *Levels of perception*, pp. 149–168. Springer, 2003.

[64] K. Sasaki and Y. Yamada. Crowdsourcing visual perception experiments: a case of contrast threshold. *PeerJ*, 7, 2019.

[65] I. Shatz. Fast, free, and targeted: Reddit as a source for recruiting participants online. *Social Science Computer Review*, 35(4):537–549, 2017. doi: 10.1177/0894439316650163

[66] W. M. Smith. A methodological study of size-distance perception. *The Journal of Psychology*, 35(1):143–153, 1953.

[67] R. Snow, B. O'Connor, D. Jurafsky, and A. Y. Ng. Cheap and fast—but is it good? evaluating non-expert annotations for natural language tasks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '08, p. 254–263. Association for Computational Linguistics, USA, 2008.

[68] K. Stanney, C. Fidopiastis, and L. Foster. Virtual reality is sexist: But it does not have to be. *Frontiers in Robotics and AI*, 7:4, 2020. doi: 10.3389/frobt.2020.00004

[69] A. Steed, S. Frlston, M. M. Lopez, J. Drummond, Y. Pan, and D. Swapp. An 'in the wild' experiment on presence and embodiment using consumer virtual reality equipment. *IEEE transactions on visualization and computer graphics*, 22(4):1406–1414, 2016.

[70] N.-C. Tai. Daylighting and its impact on depth perception in a daylit space. *Journal of Light & Visual Environment*, 36(1):16–22, 2012.

[71] G. Thomas, J. H. Goldberg, D. J. Cannon, and S. L. Hillis. Surface textures improve the robustness of stereoscopic depth cues. *Human factors*, 44(1):157–170, 2002.

[72] R. C. Toye. The effect of viewing position on the perceived layout of space. *Perception & Psychophysics*, 40(2):85–92, 1986.

[73] UploadVR. Report: Vive Users Are 95 Percent Male And Spend 5 Hours Per Week in VR. https://uploadvr.com/vive-users-94-9-percent-male-spend-5-hours-week-vr-average/, 2017.

[74] N. van Berkel, J. Goncalves, S. Hosio, and V. Kostakos. Gamification of mobile experience sampling improves data quality and quantity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):1–21, 2017.

[75] M. Wagner. The metric of visual space. *Perception & psychophysics*, 38(6):483–495, 1985.

[76] M. Wagner. *The geometries of visual space.* Psychology Press, 2006.

[77] M. Wagner. Sensory and cognitive explanations for a century of size constancy research. *Visual Experience: Sensation, Cognition, and Constancy*, 07 2012. doi: 10.1093/acprof:oso/9780199597277.003.0004

[78] P. Willemsen, A. A. Gooch, W. B. Thompson, and S. H. Creem-Regehr. Effects of stereo viewing conditions on distance perception in virtual environments. *Presence: Teleoperators and Virtual Environments*, 17(1):91–101, 2008.

[79] A. T. Woods, C. Velasco, C. A. Levitan, X. Wan, and C. Spence. Conducting perception research over the internet: a tutorial review. *PeerJ*, 3:e1058, 2015.