Estimating the Curvature of Visual Space in Virtual Reality Using an Exocentric Pointing Task

Alexis P Chambers
Centre for Applied Computing (CAC)
University of Oulu
Oulu, Finland
alexis.chambers@oulu.fi

Alizée Pinet
Laboratoire Mathématiques, Image et Applications (MIA)
La Rochelle Université
La Rochelle, France
alizee.pinet@univ-lr.fr

Evan Center
Center for Machine Vision and Signal Analysis (CMVS)
University of Oulu
Oulu, Finland
evan.center@oulu.fi

Timo Ojala
Centre for Applied Computing (CAC)
University of Oulu
Oulu, Finland
timo.ojala@oulu.fi

Steven LaValle
Centre for Applied Computing (CAC)
University of Oulu
Oulu, Finland
Steven.LaValle@oulu.fi

Nicoletta Prencipe
Institut de Recherche Mathématique de Rennes
Université de Rennes
Rennes, France
nicoletta.prencipe@univ-rennes.fr

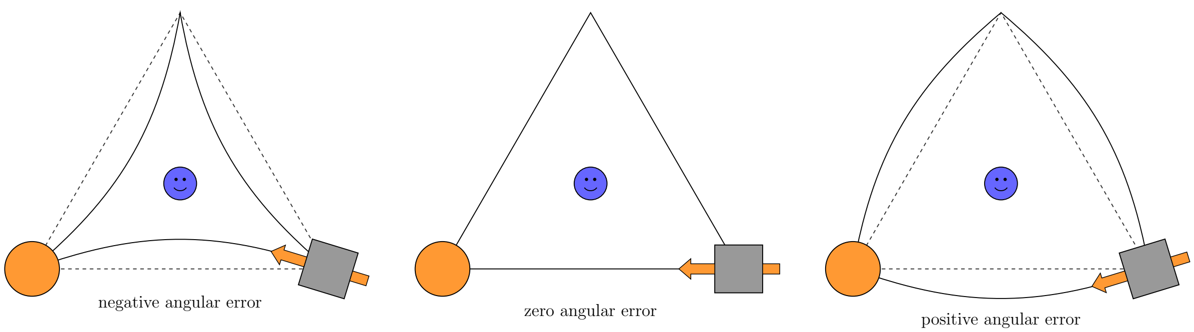

Figure 1: Three triangles illustrating possible experimental outcomes. From left to right: a negative angular error yielding a hyperbolic triangle, no angular error corresponding to a Euclidean triangle, and a positive angular error yielding an elliptic triangle. Triangle edges are drawn using Bézier curves for visualization purposes.

Abstract

Human perception of space deviates from a simple Euclidean representation of physical reality. In this study, we adapted a classic exocentric pointing paradigm to a Virtual Reality (VR) environment to investigate how the perceived curvature of space at gaze height varies with distance. Thirty participants performed a task using a VR headset, which consisted of aligning a pointer with targets arranged in equilateral triangular configurations with side lengths ranging from 2 to 20 m, while the participant was positioned at the barycenter of the triangle. The angular pointing error at each distance was interpreted as an indicator of perceived visual curvature. The results revealed a highly significant elliptic effect (positive angular error) in near space, whereas middle and far distances showed no evidence of hyperbolicity (negative angular error) and were perceived as closer to Euclidean (null angular error).

CCS Concepts

• **Computing methodologies → Perception**; **Virtual reality**.

Keywords

exocentric pointing, depth perception, virtual environments

ACM Reference Format:
Alexis P Chambers, Alizée Pinet, Evan Center, Timo Ojala, Steven LaValle, and Nicoletta Prencipe. 2026. Estimating the Curvature of Visual Space in Virtual Reality Using an Exocentric Pointing Task. In *ACM Symposium on Applied Perception 2026 (SAP '26), August 26–27, 2026, Rennes, France.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3821409.3821467

ACM ISBN 979-8-4007-2799-3/26/08
https://doi.org/10.1145/3821409.3821467

1 Introduction

While Euclidean geometry provides a reasonable model of the near physical space surrounding us, a number of studies in perceptual psychology suggest the non-Euclidean nature of the so-called *visual* *space*, understood as a mathematical model representing how humans perceive the physical space around them [Angell 1974; Blank 1953; Cuijpers et al. 2000a,b, 2001; Indow 1974; Koenderink and van Doorn 1998; Koenderink et al. 2008, 2002, 2000, 2003; Luneburg 1947, 1950; Schoumans and van der Gon 1999; Turski 2023; van Doorn et al. 2013; Wagner 1985].

If one seeks to provide a mathematical definition of visual space, it is fundamentally an object of three dimensions. From a modeling perspective, visual space can be hypothesized as an unknown three-dimensional Riemannian manifold representing physical Euclidean $\mathbb{R}^3$. If restricted to gaze height, it is treated as a two-dimensional Riemannian manifold. Because this manifold's exact properties are unknown, researchers use experimental techniques to characterize its geometry, in particular, its metric, geodesics or curvature.

Early models of visual space proposed that it may have non-Euclidean structure. Luneburg described binocular visual space under reduced-cue conditions to be hyperbolic and with constant negative curvature [Luneburg 1947, 1950]. Related accounts were developed by Blank [1953, 1958, 1959] and Indow [1974]. Other researchers argued for positive curvature or affine transformations of physical space, especially under fuller cue conditions [Angell 1974; Wagner 1985].

In the early 2000s, a substantial body of work revisited earlier models by questioning some of their assumptions. First, it emphasized perception under full-cue, rather than reduced-cue conditions. Although natural settings introduce confounding variables, strictly controlled laboratory conditions provide only a partial understanding of the phenomenon [Koenderink et al. 2000]. These studies developed and tested non-Euclidean models, undermining the idea of constant curvature. Different types of experimental tasks (including parallelity [Cuijpers et al. 2000b, 2001], bisection [Koenderink et al. 2002], and exocentric pointing [Cuijpers et al. 2000a; Koenderink et al. 2008, 2000, 2003; Schoumans and van der Gon 1999; van Doorn et al. 2013]) have been adopted, yielding inconsistent results between tasks.

Cuijpers et al. [Cuijpers et al. 2000b, 2001] used a parallelity judgment task with small samples ($n = 4$, $n = 8$), where participants adjusted a test bar to appear parallel to a reference bar in the horizontal gaze plane. They found non-constant curvature, revealing discrepancies between apparent and veridical parallelism under reduced visual context (a room covered with wrinkled black plastic). They also derived a Riemannian metric from the data, showing that curvature decreases with distance [Cuijpers et al. 2003]. When the bisection task was adopted [Koenderink et al. 2002], participants bisected a segment defined by a pair of stakes by controlling a third stake mounted on a radio-controlled vehicle. In a small sample ($n = 4$) under natural outdoor conditions, this experiment yielded apparent fronto-parallel lines that were curved and concave toward the observer, indicating positive curvature.

Koenderink and collaborators proposed the exocentric pointing paradigm, in which an observer uses a pointer at one location to indicate the direction toward a separate target, both at some distance from the observer. They compare this scenario to trying to guess where another person is looking at a cocktail party [Koenderink et al. 2000]. In contrast, the term egocentric pointing refers to aiming along the direct line of sight (e.g., with a weapon or telescope). Because exocentric pointing involves comparing directions within the scene rather than along a direct line of sight, angular pointing errors can be used, e.g., to infer the perceived curvature of the plane at gaze height or of the fronto-parallel plane. The experiment performed by Koenderink et al. took place in a natural outdoor environment, a meadow [Koenderink et al. 2000], and investigated the perceived curvature of the plane at gaze height. The pointer and target were placed at the same distance from the observer with a fixed angle of 120°. Their findings suggest that the curvature of visual space at gaze height is elliptic at close distances and hyperbolic at greater distances, in contradiction with results obtained with parallelity judgments [Koenderink et al. 2002]. A more recent theoretical work by Turski [Turski 2023] appears to be consistent with the findings in [Koenderink et al. 2000] regarding the curvature.

Other variants of the exocentric pointing task were tested: on the horizontal plane, but with the target and pointer placed at variable distances from the observer, with full-cue [Koenderink et al. 2008, 2003] and limited-cue [Cuijpers et al. 2000a] conditions; both in the fronto-parallel plane and in the horizontal one in a simple virtual environment (VE) [Schoumans and van der Gon 1999]; only in the fronto-parallel plane using a monitor with no head movements allowed [van Doorn et al. 2013]. Although the sample size across these studies was limited ($n = 3$ for [Cuijpers et al. 2000a; Koenderink et al. 2000, 2003; van Doorn et al. 2013] and $n = 5$ for [Koenderink et al. 2008; Schoumans and van der Gon 1999]), and some included the authors as participants [Koenderink et al. 2000, 2003; van Doorn et al. 2013], the overall conclusion was that there was a systematic deviation in the pointing task with respect to the veridical direction, which was considered an indicator of an intrinsic non-Euclidean geometry of both the fronto-parallel plane and the horizontal gaze-height plane.

In two early studies [Schoumans et al. 2000; Schoumans and van der Gon 1999], a simple virtual setup (LCD shutter goggles) was already used to investigate exocentric pointing. Since then, immersive technologies have developed considerably with Virtual Reality (VR) which has made it possible to revisit classical experiments with the precise control of experimental stimuli required in psychophysics (e.g., for spatial cognition), as noted in Bülthoff and van Veen [2001]. Unlike conventional laboratory displays, VR allows participants to perceive and act within a spatially structured environment while the experimenter fine-tunes the viewing conditions, environmental layout, object features, and availability of depth cues.

VR provides a useful intermediate approach between field studies and traditional laboratory displays: it allows for this precise control while still presenting participants with an immersive spatial scene. At the same time, extended reality (XR) can influence perceived space. Distance judgments in XR can differ from those in natural environments and are affected by display characteristics, contextual cues, observer height, and object location in the visual field [Bruder et al. 2015; Interrante et al. 2006; Leyrer et al. 2011; Peillard et al. 2019b; Renner et al. 2013]. Although these effects are most often studied with egocentric distance tasks, they are relevant here because exocentric pointing depends on perceived spatial relationships between the observer, pointer, and target.

There already exists a large body of work investigating egocentric distance perception in XR [Adams et al. 2022; Bruder et al. 2015; Creem-Regehr et al. 2023; Hmaiti et al. 2023; Interrante et al. 2006;

Peillard et al. 2019b; Renner et al. 2013; Vaziri et al. 2021, 2017]. A commonly adopted task is blind walking, in which a participant is shown a distant object and asked to close their eyes and walk to where they believe it is, to record whether people tend to over- or underestimate perceived distances.

Peillard et al. [2019b] reported lack of uniformity in perception across the visual field. In particular they investigated egocentric distance perception for objects presented in front of and to the sides of an observer in VR. They found that objects in the periphery were consistently perceived to be farther away than objects at the same distance in front of them. This result could imply anisotropy of visual space, consistent with Cuijpers et al. [2000a].

Exocentric distance estimation was analyzed using AR in Peillard et al. [2019a]. The task consisted of reproducing the distance between two objects placed 2.1 m away from the observer, in various combinations of real and virtual objects. The authors found a significant overestimation of the veridical distances, which was significantly greater when both objects were real rather than both virtual.

Despite extensive work on distance perception in VR, to our knowledge there is no published study that uses an exocentric pointing task to measure the curvature of visual space at gaze height in VR. To address this gap, we replicated and extended the paradigm used by Koenderink et al. [Koenderink et al. 2000] in VR, in particular to assess whether the reported curvature effects remain valid with a larger sample size. We conducted the study in a reduced-cue setting, and future work will involve running the task in a full-cue VE, similarly to the original study [Koenderink et al. 2000].

We first conducted a pilot experiment (n = 20) to test the paradigm and estimate effect sizes to be used for an *a priori* power analysis, because the original experiment included only the three authors as participants. The pilot examined whether participants in VR would show *positive* angular excess at near distances, *negative* angular excess at far distances, and *near-zero* angular excess at intermediate distances. It also explored whether individual factors, such as vision type and VR experience, modulated angular pointing errors. Bayes factors were calculated for angular errors in the pilot experiment, which indicated very strong evidence for an effect of elliptic triangles in the very close visual space (2 m range; BF_{10} = 955.68), moderate evidence for Euclidean triangles in the 4 m range (BF_{10} = 0.17) and anecdotal evidence for hyperbolic triangles in the 6 m range (BF_{10} = 2.12). In the far range from 8 m to 20 m, the results tended toward zero, indicating Euclidean geodesics with moderate support (BF_{10} = 0.18–0.33). Informed by the pilot results, the current experiment followed the same general procedure as the pilot except that it did not include a visible horizon behind the objects, had more trials per distance, had seven distances rather than ten, and excluded pointer starting positions within 30° of the target direction.

Based on these pilot results, the present study tested two preregistered directional hypotheses about the curvature of visual space at gaze height. For H1, we predicted that the mean angular error at 2 m would be significantly greater than zero indicating perceived elliptic curvature. For H2, we predicted that at 6 m, the mean angular error would be significantly less than zero, indicating hyperbolic curvature effects. As an exploratory measure, we investigated whether the mean angular error at 4 m, at intermediate distances (8–10 m), and at far distances (15–20 m) would remain close to zero, indicating Euclidean geometry.

2 Methods

The purpose of this study was to investigate the perceived visual geometry of the horizontal plane at gaze height in a simple VE using an exocentric pointing task through replication and extension of the experiment by Koenderink et al. [Koenderink et al. 2000] adapted to VR.

The methods, procedures, and protocols used in this experiment were approved by the University of Oulu's Ethics Committee of Human Sciences. The sample size, directional hypotheses, methods, and analyses were preregistered on the Open Science Framework before data was collected. The data and analysis script for the pilot experiment can be found on OSF.

2.1 Participants

We conducted an *a priori* power analysis using pilot study data (n = 20). The pilot showed a very large effect size of angular error at the closest distance (Cohen's d = 1.22 at 2 m), with smaller effects at intermediate distances and effects close to zero at far distances. We focused the power analysis on the 6 m effect ($d = -0.55$) as it was the smallest effect of interest in favor of an alternative hypothesis. To reach a power of at least 0.80 using a one-tailed, Bonferroni corrected alpha at .025 (consistent with two directional predictions at 2 and 6 m), we needed 29 participants, which we rounded to 30.

Preregistered exclusion criteria included technical interruptions, voluntary discontinuation, imperfect vision without corrective eyewear, and failure to follow task instructions. In total, we collected data from 32 participants, excluding two: one due to a tracking beacon failure, and the other due to having imperfect vision without using corrective lenses. Participants were recruited from the local university's participant pool, were at least 18 years old, and gave written informed consent for participation and data processing. They received a university merchandise item worth approximately 10 euros and were entered in a gift card raffle.

Before the task, participants completed a questionnaire on age, gender, vision, eye dominance, handedness, prior VR experience, and video game experience. Thirty participants were included in the analysis (mean age = 28.5 years, range = 18–40; 8 female, 22 male). Mean height was 170.7 cm (range = 122—192 cm), and average interpupillary distance (IPD), based on self-adjusted head-mounted display (HMD) settings, was 64.0 mm (range = 58—70 mm). All participants reported normal or corrected-to-normal vision and wore corrective lenses if needed. All participants except one were right-handed; 18 were right-eye dominant and 12 were left-eye dominant, as measured with the Miles Eye Dominance test [Miles 1930]. Most participants had limited prior VR experience: 23 reported either no prior use, only a few uses ever, or use once or twice per year.

2.2 Virtual environment design

The VE used in this experiment was built in Unity version 6000.1.17f1 using the XR Interaction Toolkit. It was intentionally minimal to prevent participants from using environmental objects, shadows, or horizon cues to align the pointer with the target. The scene consisted of a light gray-white plane and a gradient skydome that blended a blue sky into the light plane above the actual horizon, avoiding a visible horizon line behind the objects (see Figure 2). This design was informed by the pilot study, in which the horizon remained visible behind the objects and created a possible pointing aid.

The target and pointer were adjusted to the participant's HMD height while they looked straight ahead, matching the virtual objects' height to gaze height accounting for the known effects of height on distance perception in VR [Leyrer et al. 2011]. As shown in Figures 2 and 3, the target was an orange sphere with a diameter of one half-meter, and the pointer was a gray cube measuring one half-meter per side, with an orange arrow through its center. The objects did not cast shadows. The cube was constrained to rotate only within the horizontal plane, replicating the task from the original Koenderink experiment, but in a stripped-down, controlled setting.

Participants stood at the barycenter of an invisible equilateral triangle facing one side, as shown in Figure 3. The sphere and pointer were placed at the two visible vertices separated by a 120° angle in the participant's view (see Figure 3). Participants needed to turn their head back and forth to see both objects. The triangle's scale was adjusted across trials so the distance from the participant to the objects changed, while the triangle remained equilateral, the objects remained the same size, and the participant remained at the barycenter. From initial tests, it was determined that scaling the task objects across distances to maintain a constant visual angle, as done for the sphere in the original task by Koenderink et al. [Koenderink et al. 2000] and for both the sphere and pointer in [Koenderink et al. 2008], removed the impression of distance changes in an environment devoid of other spatial cues. Therefore, the objects were kept at a constant size across all distances.

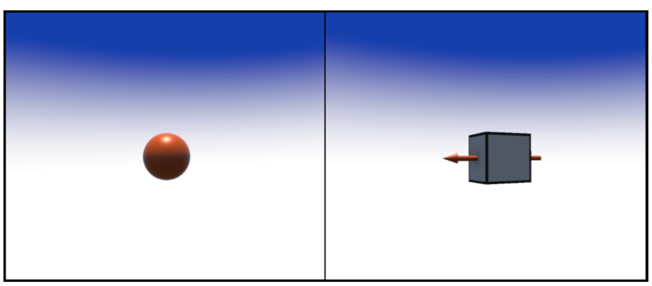

Figure 2: View of each object in the VE as the participant saw them, at eye height. The objects cannot be viewed simultaneously due to their 120° separation in the participant's view.

2.3 Experiment design

This experiment was a within-subjects design, and participants were naive to the purpose of the study. The equipment used was a Valve Index VR HMD with a tethered PC connection, Valve Index controllers, and two HTC Vive base stations for position tracking. Participants stood at a marked location on the laboratory floor with a free space of approximately two square meters.

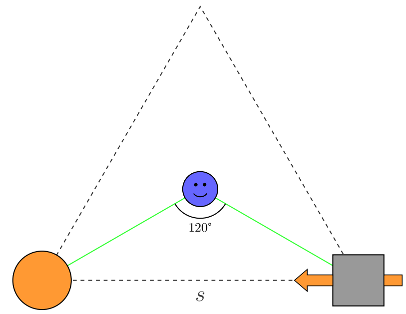

Figure 3: A diagram of the overhead view of the objects in equilateral their triangle configuration. The distances in each trial (s = 2–20 m) refer to the edges of the triangle and not the distance from the participant to either object.

The pilot study sampled 10 distances from 2–20 m in 2 m increments, with six trials per distance. Because distance effects remained near zero at the measured distances beyond 6 m, the present study sampled the farther distances more sparsely. We retained the distances in the 2–10 m range where finer sampling was likely to be more informative, and we changed the far range to include only 15 and 20 m and collected 10 trials per distance. This structure reduced time spent on uninformative trials in the far range and increased repetitions in the near range without increasing the total task length for participants. For the present study, the pointer and sphere objects appeared at the observer's eye height, at seven triangle side lengths: 2, 4, 6, 8, 10, 15, and 20 m. The sphere target appeared on either the right or left side of the observer's view, and each distance-by-side combination was repeated five times, resulting in 70 trials. Trials were presented in randomized order. The starting direction of the pointer was also randomized, excluding the 30° region on either side of the direct line between the pointer and target to prevent starting-position bias. The pointer cube was restricted to rotation within the horizontal plane at gaze height. Participants rotated the pointer cube horizontally until they believed it pointed directly at the sphere. They used the joystick on the controller to rotate the cube, a button to record the desired orientation, and the trigger button to advance to the next trial. They were allowed to turn their head and twist their torso, but were asked not to take any steps away from the starting position. Three unrecorded practice trials were included to ensure participants understood the task. The experiment was self-paced with no time limit, and took an average of 16.1 minutes to complete.

2.4 Geodesics and curvature

In our experiment, the angular error for each trial was calculated from the difference between the participant's final pointer direction and the veridical direction from the center of the pointer cube to the center of the sphere target, see Figure 1. The veridical direction was taken as 0°. Angular errors were defined as negative or positive when the participants pointed toward or away from themselves, respectively. These per-trial pointing errors are later used to estimate the angular excess of the implied triangle. The angular excess of a geodesic triangle is the degree to which the sum of a triangle's interior angles deviate from the Euclidean sum of 180°. 'Thin' triangles have negative angular excess, while 'inflated' triangles have positive angular excess, as the sum of internal angles is strictly less than 180° for thin triangles and strictly greater for inflated ones (see Figure 1). Koenderink and colleagues proposed using angular excess to characterize the geometry of visual space at gaze height [Koenderink et al. 2000]. As mentioned in the introduction, we assume that the visual space at gaze height constitutes an unknown two-dimensional Riemannian manifold that models how the physical Euclidean plane at gaze height is perceived. Angular pointing errors were assumed to be directed along the tangents to the geodesic segments joining the triangle's vertices at the corresponding vertices.

Let T_i be the geodesic triangle, defined as the union of the three geodesic segments joining three points, obtained for $i = 0, \ldots, 6$, in the settings of our experiment (by varying the length of the side, $s = 2, 4, 6, 8, 10, 15, 20$, see Figure 3). Let E_i denote the angular excess of the i-th triangle T_i in radians.

We must recall that the only data obtained from the experiment are angular excesses. Therefore, there is no unique choice for obtaining a triangle T_i starting from its angles. Koenderink and collaborators themselves stress that this is not a trivial issue, referring to it as 'a brute fact that is of considerable conceptual interest', when discussing the construction of geodesic triangles T_i from data obtained solely through the exocentric pointing task [Koenderink et al. 2008]. This raises questions about the assumed geometric structure of the space, as an infinite number of arbitrary choices are possible. One could choose to trace the triangles using circular arcs or Bézier curves, for example. This does not affect the values of E_i, but affects calculations concerning the curvature. Therefore, we choose to represent the geodesic triangles only for visualization purposes, using quadratic Bézier curves (see e.g., Figure 6).

In the following, we describe how the curvature was approximated for each triangle. As mentioned also in Koenderink et al. [2000], the angular excess of a geodesic triangle is directly related to a notion of curvature: the integral of the Gaussian curvature, the so-called *curvatura integra* or total curvature, via the Gauss-Bonnet theorem. Gauss showed that the total curvature of a geodesic triangle is equal to its angular excess. More precisely, let K be the Gaussian curvature of the visual space at gaze height. A consequence of the Gauss–Bonnet theorem for geodesic triangles is that

$$\iint_T K \mathrm{d}\sigma = \varphi_1 + \varphi_2 + \varphi_3 - \pi = E \qquad (1)$$

where φ_i, for $i = 1, 2, 3$, are the internal angles of a geodesic triangle T expressed in radians, E its angular excess, and $\iint_T K\mathrm{d}\sigma$ is the total curvature over the surface inside T, where $\mathrm{d}\sigma$ denotes the Riemannian area element, see e.g., [do Carmo 1976].

As a preliminary approach, one can assign a Gaussian curvature to each triangle by assuming that K is constant and equal to K_i over the surface whose boundary is T_i. Let us denote with A_i the area of the i-th geodesic triangle, $A_i = \iint_{T_i} \mathrm{d}\sigma$, for $i = 0, \ldots, 6$. Then, by applying Equation (1) one can obtain the values for the curvature as follows:

$$K_i = \frac{E_i}{A_i}, \qquad \text{for } i = 0, \ldots 6. \qquad (2)$$

It remains to compute the areas A_i of these triangles. The Gauss-Bonnet theorem requires the computation of an integral over a two-dimensional Riemannian manifold, which is unknown in our framework. Therefore, we decided to approximate the area by computing the surface of the corresponding Euclidean triangles sharing the same vertices of the T_is in $\mathbb{R}^2$. Under the assumption of constant Gaussian curvature over each triangle, this allowed us to determine an approximation of the curvature of the data for each triangle using (2), as detailed in Section 3.2.

2.5 Analyses

A linear mixed-effects model was used to predict angular error as a function of distance, with a random intercept term to capture subject-level differences. The confirmatory analyses were chosen to test the two directional hypotheses specified before data collection based on the pilot study. The pilot results indicated a large positive angular error at 2 m, a medium negative angular error at 6 m, and null effects at the remaining distances. Therefore, the preregistered confirmatory tests focused only on the 2 m and 6 m distances where there was a reason to expect directional effects, assessing whether the mean error at 2 m is greater than zero and whether the mean error at 6 m is less than zero. These tests were conducted as one-sided one-sample t-tests on participant-level model-estimated mean angular errors at each distance. The two confirmatory tests were Bonferroni corrected with an alpha level of $\alpha = .025$. No directional hypotheses were specified for the remaining distances, and these distances were not treated as confirmatory tests.

For completeness, as exploratory measures we also computed model-estimated means and 95% confidence intervals for each distance to describe the full pattern of results across all sampled distances. To estimate support for the null hypotheses relative to the alternative hypothesis at each distance, we calculated Bayes factors (for a primer, see [Wagenmakers et al. 2018]). We used a Bayesian one-sample t-test using a Cauchy prior (with a width of 1) for the participant-level mean angular errors.

Further exploratory analyses included modeling effects of distance and demographic and background information measures like vision type, eye dominance, VR and gaming frequency, gender, and age. Our goal here was to determine whether these factors had any effect on the participants' angular pointing errors across trials, including possible effects at distances that were not part of the confirmatory hypotheses. This exploratory model included distance as a categorical fixed effect and participant as a random intercept. We used the R package `buildmer` to perform model selection. The algorithm started with a maximal candidate model which included predictors for distance, target side, age, gender, height, vision measures, VR use, and gaming frequency. We also included interactions between distance and the demographic variables previously stated. Competing models were compared using backward elimination and likelihood ratio tests until the model with the best fit was found.

3 Results

Data from all 30 included subjects were analyzed. Visual inspection of the Q-Q plot of model residuals for normality revealed minor deviations at the tails, but no significant causes for concern.

Table 1: Cohen's d and Bayes Factors for Mean Error at each triangle side length.

Triangle side length (m)	Cohen's d	BF_{10}
2	1.29	177781.00
4	0.37	0.90
6	0.11	0.17
8	0.06	0.15
10	0.06	0.15
15	0.09	0.16
20	0.15	0.20

3.1 Angular error analysis by distance

The estimated mean angular error at 2 m was significantly greater than zero, $M = 8.52°$, ($SE = 1.45$), 95% CI [5.57, 11.47], $t(32.3) = 5.879$, $p < .001$ (one-sided, Bonferroni-corrected), consistent with H1. The estimated mean angular error at 6 m was not significantly less than zero, $M = 0.93°$, ($SE = 1.45$), 95% CI [−2.02, 3.88], $t(32.3) = 0.640$, $p = .737$ (one-sided, Bonferroni-corrected), contrary to H2.

Results of the average angular error by distance shown in Figure 4 indicate a very strong effect at close distances. The effect size of angular error by distance was calculated using Cohen's d. At 2 m, there was a very large effect ($d = 1.29$). At 4 m, the effect of error was small ($d = 0.37$). At 6 m, the effect size was very small ($d = 0.11$), and effect sizes remained practically negligible after 6 m (i.e., from 8 to 20 m; $d = 0.06$–0.15).

To determine the likelihood of the null hypothesis, that visual geometry is Euclidean, versus the alternative hypothesis, that visual geometry is curved, we calculate Bayes Factors. In Table 1, we see extremely strong evidence for the alternative hypothesis at the closest distance ($BF_{10} = 177781.00$). At 4 m, evidence was inconclusive ($BF_{10} = 0.90$). At 6 m, evidence was moderately in favor of the null ($BF_{10} = 0.17$). At all other measured distances, Bayes factors moderately favored the null ($BF_{10} = 0.15$–0.20), indicating that angular error did not reliably differ from zero between 6 and 20 m.

3.2 Arcplot and curvature

To depict geometric distortions deduced from the participants' pointing directions, we generated arcplots following the visualization approach introduced by Koenderink et al. [Koenderink et al. 2000]. However, while they used circular arcs to visualize geodesic triangles T_i, we used quadratic Bézier curves. With this method the average T_i among all trials is plotted for each side length s. The group mean arcplot, depicted in Figure 6, visualizes an elliptic triangle at 2 m, consistent with the significant positive angular error at this near distance. There is weaker evidence of an elliptic triangle at 4 m. At distances of 6 m up to 20 m, the angular errors remain positive, but align closely with the Euclidean configuration, indicating that for each triangle the visual space appears effectively Euclidean in the mid- and far-range distances tested here. In contrast to the original study, we do not report hyperbolic triangles at large distances. An array of the individual arcplots from our study can be found in the appendix in Figure 7.

As detailed in Section 2.4, under certain assumptions it is possible to compute, from the angular error data, the Gaussian curvature K_i associated with each triangle T_i, $i = 0, \ldots, 6$, using Equation (2). The plot in Figure 5 shows that, for each tested distance, the approximated curvature of the visual space at gaze height is positive for the 2 m and 4 m triangles, and is closer to the Euclidean triangle configuration at larger distances. This distance-by-distance behavior is consistent with the angular error shown in Figure 4. Note that the curvature approximation in Figure 5 should be interpreted differently from that in Figure 4. Angular error is measured in degrees and directly reflects the pointer error at each distance. Curvature normalizes this angular excess by the area of the corresponding triangle. Because triangle area grows quickly with side length, a similar angular error at larger distances corresponds to a much smaller curvature estimate. Thus, curvature values approach zero at larger side lengths even when small deviations remain visible in Figure 4.

When computing curvature based solely on angular excess, several modeling choices are involved, especially in how triangles and their areas are approximated. In this work, the curvature is estimated under modeling assumptions that include approximating triangles areas over the manifold with the area of the Euclidean corresponding ones in $\mathbb{R}^2$. Within this framework, the curvature is computed independently for each triangle, without coupling between neighboring elements. The reported values should therefore be interpreted as local empirical estimates obtained within this discretized procedure. Further investigations are required to provide a global description of Gaussian curvature as a function of distance.

3.3 Exploratory results

3.3.1 Mixed-effects modeling. The final model selected by `buildmer` included distance and sphere side as significant predictors ($R^2_{marginal} = 0.07$, $R^2_{conditional} = 0.63$). As in the confirmatory analysis, distance was a strong predictor of angular error, with all distances showing significantly lower error relative to the 2 m reference distance (all $p < .001$). A small but significant effect of sphere side was also retained in the final model ($\eta_p^2 = 0.004$), with targets appearing on the right side associated with slightly higher angular error compared to the left ($b = 0.77°$, $SE = 0.27$, $t(2063) = 2.85$, $p = .004$). To confirm the predictive power of the model apart from accounting for individual differences among participants, the final model was compared against a random intercepts null model. The final model provided a significantly better fit to the data ($\chi^2(7) = 371.31$, $p < .001$). Neither gaming frequency, VR frequency, gender, nor eye dominance were retained as significant predictors in the final model.

3.3.2 Questionnaire results. Participants rated the difficulty of the pointing task using a 5-point Likert scale for close, middle, and far distance ranges. Ratings increased with distance: close distances were rated as easy ($M = 1.50$, $SD = 0.94$), middle distances as moderate ($M = 3.00$, $SD = 0.89$), and far distances as difficult ($M = 4.10$, $SD = 0.92$).

In open-ended responses, 12 participants reported using imagined lines or geometric strategies, nine reported no strategy, and the remaining responses mentioned repeated cube positioning, head movements, or first pointing toward themselves. One participant reported having "found the trick," but ranked only 22nd out of 30 in mean pointing accuracy.

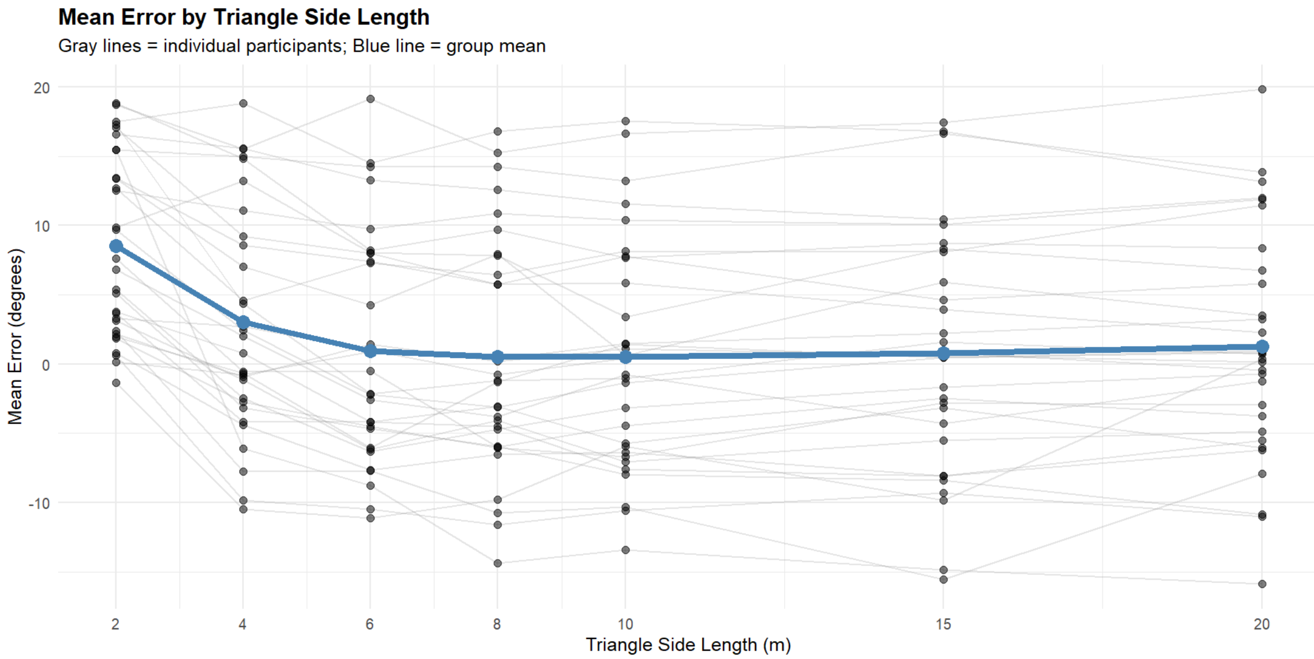

Figure 4: Mean angular error by triangle side length across all participants. The blue line represents the group mean angular error.

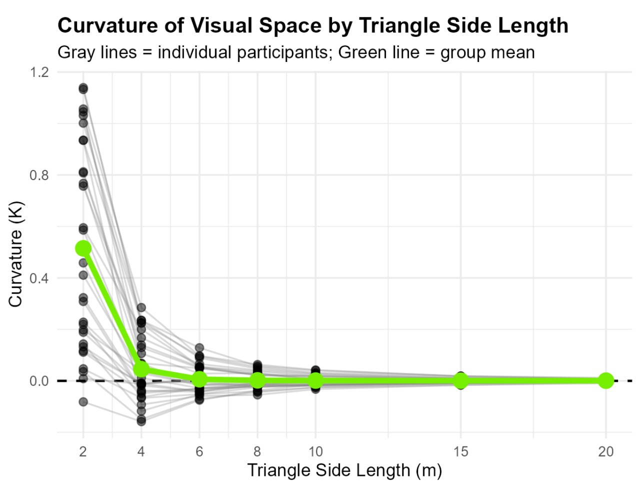

Figure 5: Estimated Gaussian curvature (K_i) associated to each triangle side length, plotted as the mean across trials for each participant and the overall group mean.

4 Discussion

Our results show a very robust effect of perceived elliptic triangles in the near visual space and near-Euclidean triangles at middle and far distances, consistent with the pattern observed in our pilot study. Most apparent in Figure 4, the strongest effect occurred at 2 m, evidence was weaker at 4 m, and angular error remained close to zero from 6 m to 20 m. However, they only partially replicate those of the original study [Koenderink et al. 2000]. The predicted near-distance effect was supported, whereas the negative angular excess at farther distances was not replicated in our VR environment with low visual cues.

Although the results from each study differ somewhat at the 4 and 6 m distances, there is not strong evidence for effects at these distances in either study. Across both studies, we can conclude that there is strong support for the effect of elliptic geodesics in the immediate near space and moderate support for Euclidean geodesics from 6 m up to 20 m. Our results seem to be in accordance with the visual fronto-parallel lines found in [Koenderink et al. 2002] using a bisection task (concave towards the observer, hence elliptic at near distances, 2 m range, and approximately straight, thus Euclidean, at larger distances, 10 m range).

One plausible explanation for this narrowly spread result could be the simple environment used. We intentionally excluded many of the depth cues available in the original meadow, apart from the size of the objects shrinking as they moved farther afield. With only a flat plane and two floating objects, participants may have relied primarily on simple angles between the objects rather than details in a natural setting. However, adding contextual cues would not necessarily have a uniform effect. Structured cues, such as walls, paths, or city blocks, could provide references for pointer alignment, whereas irregular natural cues could add depth and scale information without direct orientation information. This relates to earlier work demonstrating that perceived geometry and exocentric pointing can depend on the available cues and details of the scene [Doumen et al. 2008; Schoumans et al. 2000; Wagner 1985]. Thus, the near-space elliptic effect seen here may reflect either a robust feature of the task or a response to the reduced-cue environment. The use of VR may also have contributed to the discrepancy, since judgments can differ between real and virtual settings. Specifically, Peillard et al. [Peillard et al. 2019a] highlighted that egocentric distance and depth perception tend to be systematically underestimated in virtual environments compared to the real world. Although VR allowed us to standardize conditions across participants, this known perceptual compression may also have changed the perceptual information available for judging the relation between the pointer and target.

It is also important to stress that the results reported in [Koenderink et al. 2000] were obtained from a sample of three author-participants, which is prone to sampling bias; one could easily imagine drawing three participants at random, like those that make up the lower half of Figure 4. Thus, an explanation for the discrepancy is that the original far-distance hyperbolic effect may reflect sampling variability.

Individual differences also played a role in predicting spatial perception. In Figures 4 and 7, we can observe that there is substantial variability between individuals (though relatively little variability *within* individuals) when it comes to overall positive or overall negative angular errors. Some participants showed predominantly positive errors, some were mostly negative, and others that were nearly Euclidean across all distances. Despite the individual variability, the exploratory mixed-effects model did not retain gaming frequency, VR frequency, gender, or eye dominance as significant predictors, suggesting that the measured individual factors did not explain this variability. Future work should examine other possible sources of variability, such as spatial ability, simulator experience, and head or torso movements during pointing.

The exploratory mixed-effects model retained a small effect of sphere side, with targets on the right associated with slightly higher angular error. However, this effect was very small ($\eta_p^2 = 0.004$), not predicted by our hypotheses, and secondary to the much larger effect of distance. Because only one participant was left-handed, we interpret the side effect cautiously and do not treat it as evidence for a systematic asymmetry in the curvature of visual space.

5 Limitations and future work

A main limitation of the present study is that it cannot determine whether the discrepancy with the original study is due to the VE, the limited depth cues in the environment, the technology used, differences in statistical power, or perhaps a mixture of factors. Future work should directly compare real-world and VEs within the same set of participants in a well-powered study to determine whether any differences are attributable to VR and cues.

Recent work in natural outdoor environments provides important context for these findings, as Norman et al. [Norman et al. 2025] replicated Blank's paradigm [Blank 1958] using an exocentric distance estimation task with a sample of 30 participants. As in the present study, they observed high inter-observer variability. They further found that judgments of short distances (around 2 m) were often consistent with elliptic geometry, whereas judgments of larger distances (around 15.5 m) were consistent with hyperbolic geometry.

Building on the premise that visual perception is more veridically studied in ecological settings [Gibson 1979] and given that immersive technologies facilitate such approaches [Bülthoff and van Veen 2001], future work will include two detailed VEs in addition to the simple one. A natural scene, similar to the overgrown field where the Koenderink et al. study took place [Koenderink et al. 2000], and a cityscape with more explicit geometric cues, such as city blocks and linear architectural elements. The question is whether and to what extent these added cues influence the angular excess. Specifically, do environmental cues facilitate performance in the exocentric pointing task, and does structured urban geometry or irregular natural structure differently aid pointer alignment with the target?

Nonetheless, there is evidence that the exocentric pointing task is influenced by contextual elements. Schoumans et al. [Schoumans et al. 2000] showed that, in simulated environments, the presence of oriented planes systematically affected pointing performance, without reducing intra-observer variability. Doumen et al. [Doumen et al. 2008], found that poster boards as contextual objects in real settings influenced responses only when they served as direct references for orienting the pointer. This suggests that observers might achieve more veridical, near-Euclidean estimates when explicit reference cues are available, in accordance with [Wagner 1985].

Finally, although our design did not focus on the role of target location in the visual field, our results could be correlated with other studies showing that spatial judgments in VR vary depending on where an object appears relative to the observer [Peillard et al. 2019b]. In our experiment, the two objects were separated by a wide angle that required head and body movements to compare them, which could affect the final pointer setting. Future studies could examine how directional factors, head rotations, or asymmetries in viewing posture interact with exocentric pointing.

In the present work, we deliberately adopt a simple strategy that allows for a tractable estimation of curvature, though it should be regarded as a preliminary step. A more rigorous mathematical framework, likely requiring more advanced geometric tools and beyond the scope of this paper, would be needed to provide a more accurate characterization of the geometry of visual space. In particular, such an approach could enable the estimation of Gaussian curvature as a continuous function of distance. Future work should also investigate the relationship between these geometric properties and established results in egocentric distance perception [Gilinsky 1951; Peillard et al. 2019b], particularly whether a metric derived from exocentric pointing tasks is compatible with observed distance-compression or dilation effects.

In conclusion, a central difficulty lies not only in determining an appropriate mathematical structure for visual space, which appears context- and task-dependent [Schoumans et al. 2000], but also in developing methods to infer its curvature across diverse experimental data.

Acknowledgments

This work was supported by the European Research Council (project ILLUSIVE 101020977), the Research Council of Finland (projects BANG! 363637, and PERCEPT 322637), and the University of Oulu and Research Council of Finland (PROFI7 352788). The authors thank the anonymous reviewers for their insightful comments, which helped improve this manuscript.

References

Haley Adams, Jeanine Stefanucci, Sarah Creem-Regehr, and Bobby Bodenheimer. 2022. Depth perception in augmented reality: The effects of display, shadow, and position. In *2022 IEEE conference on virtual reality and 3D user interfaces (VR)*. IEEE, 792–801.

Richard B. Angell. 1974. The geometry of visibles. *Noûs* (1974), 87–117.

Albert A. Blank. 1953. The Luneburg Theory of Binocular Visual Space. *Journal of the Optical Society of America* 43, 9 (1953), 717–727.

Albert A Blank. 1958. Analysis of experiments in binocular space perception. *Journal of the Optical Society of America* 48, 12 (1958), 911–925.

Albert A. Blank. 1959. The Luneburg Theory of Binocular Space Perception. In *Psychology: A Study of a Science*, Sigmund Koch (Ed.). Vol. 1. McGraw Hill, 395–426.

Gerd Bruder, Fernando Argelaguet Sanz, Anne-Hélène Olivier, and Anatole Lécuyer. 2015. Distance Estimation in Large Immersive Projection Systems, Revisited. In *2015 IEEE Virtual Reality (VR)*. IEEE, 27–32.

Heinrich H. Bülthoff and Hendrik A. H. C. van Veen. 2001. Vision and action in virtual environments: Modern psychophysics in spatial cognition research. In *Vision and attention*. Springer, 233–252.

Sarah H. Creem-Regehr, Jeanine K. Stefanucci, and Bobby Bodenheimer. 2023. Perceiving Distance in Virtual Reality: Theoretical Insights from Contemporary Technologies. *Philosophical Transactions of the Royal Society B* 378, 1869 (2023), 20210456.

Raymond H. Cuijpers, Astrid M. L. Kappers, and Jan J. Koenderink. 2000a. Investigation of visual space using an exocentric pointing task. *Perception & Psychophysics* 62, 8 (2000), 1556–1571.

Raymond H. Cuijpers, Astrid M. L. Kappers, and Jan J. Koenderink. 2000b. Large Systematic Deviations in Visual Parallelism. *Perception* 29, 12 (2000), 1467–1482.

Raymond H. Cuijpers, Astrid M. L. Kappers, and Jan J. Koenderink. 2001. On the Role of External Reference Frames on Visual Judgements of Parallelity. *Acta Psychologica* 108, 3 (2001), 283–302.

Raymond H. Cuijpers, Astrid M. L. Kappers, and Jan J. Koenderink. 2003. The Metrics of Visual and Haptic Space Based on Parallelity Judgements. *Journal of Mathematical Psychology* 47, 3 (2003), 278–291.

Manfredo P. do Carmo. 1976. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, NJ.

Michelle J. A. Doumen, Astrid M. L. Kappers, and Jan J. Koenderink. 2008. Do reference surfaces influence exocentric pointing? *Acta psychologica* 128, 2 (2008), 310–317.

James J. Gibson. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin.

Alberta S. Gilinsky. 1951. Perceived size and distance in visual space. *Psychological review* 58, 6 (1951), 460.

Yahya Hmaiti, Mykola Maslych, Eugene M. Taranta, and Joseph J. LaViola. 2023. An Exploration of the Effects of Head-Centric Rest Frames on Egocentric Distance Judgments in VR. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 263–272.

Tarow Indow. 1974. On geometry of frameless binocular perceptual space. *Psychologia* 17, 1 (1974), 50–63.

Victoria Interrante, Brian Ries, and Lee Anderson. 2006. Distance Perception in Immersive Virtual Environments, Revisited. In *IEEE Virtual Reality Conference (VR 2006)*. IEEE, 3–10.

Jan J. Koenderink and Andrea J. van Doorn. 1998. Exocentric pointing. *Vision and action* (1998), 295–313.

Jan J Koenderink, Andrea J van Doorn, Astrid ML Kappers, Michelle JA Doumen, and James T Todd. 2008. Exocentric pointing in depth. *Vision Research* 48, 5 (2008), 716–723.

Jan J. Koenderink, Andrea J. van Doorn, Astrid M. L. Kappers, and Joseph S. Lappin. 2002. Large-Scale Visual Frontoparallels Under Full-Cue Conditions. *Perception* 31, 12 (2002), 1467–1475.

Jan J. Koenderink, Andrea J. Van Doorn, and Joseph S. Lappin. 2000. Direct Measurement of the Curvature of Visual Space. *Perception* 29, 1 (2000), 69–79.

Jan J. Koenderink, Andrea J. Van Doorn, and Joseph S. Lappin. 2003. Exocentric pointing to opposite targets. *Acta Psychologica* 112, 1 (2003), 71–87.

Markus Leyrer, Sally A. Linkenauger, Heinrich H. Bülthoff, Uwe Kloos, and Betty Mohler. 2011. 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*. 67–74.

Rudolf K. Luneburg. 1947. *Mathematical Analysis of Binocular Vision*. Princeton University Press.

Rudolf K. Luneburg. 1950. The metric of binocular visual space. *Journal of the Optical Society of America* 40, 10 (1950), 627–642.

Walter R. Miles. 1930. Ocular dominance in human adults. *The journal of general psychology* 3, 3 (1930), 412–430.

J. Farley Norman, Maria Carmichael, Evan Hagan, Yulia Mishchuk, Emily N. Bryant, Alejandro B. Ramirez, and Avery Sutton. 2025. Visual outdoor space perception. *Scientific Reports* 15, 1 (2025), 11549.

Etienne Peillard, Ferran Argelaguet, Jean-Marie Normand, Anatole Lécuyer, and Guillaume Moreau. 2019a. Studying exocentric distance perception in optical see-through augmented reality. In *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 115–122.

Etienne Peillard, Thomas Thebaud, Jean-Marie Normand, Ferran Argelaguet, Guillaume Moreau, and Anatole Lécuyer. 2019b. 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)*. IEEE, 227–236.

Rebekka S. Renner, Boris M. Velichkovsky, and Jens R. Helmert. 2013. The perception of egocentric distances in virtual environments-a review. *ACM Computing Surveys (CSUR)* 46, 2 (2013), 1–40.

Nicole Schoumans, Jan J Koenderink, and Astrid ML Kappers. 2000. Change in perceived spatial directions due to context. *Perception & Psychophysics* 62, 3 (2000), 532–539.

Nicole Schoumans and Jan J. Denier van der Gon. 1999. Exocentric pointing in three-dimensional space. *Perception* 28, 10 (1999), 1265–1280.

Jacek Turski. 2023. Riemannian geometries of visual space: Variable curvature and horizon. *Mathematical Methods in the Applied Sciences* 46, 8 (2023), 9298–9324.

Andrea van Doorn, Jan J. Koenderink, and Johan Wagemans. 2013. Exocentric Pointing in the Visual Field. *i-Perception* 4, 8 (2013), 532–542.

Koorosh Vaziri, Maria Bondy, Amanda Bui, and Victoria Interrante. 2021. Egocentric distance judgments in full-cue video-see-through VR conditions are no better than distance judgments to targets in a void. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 1–9.

Koorosh Vaziri, Peng Liu, Sahar Aseeri, and Victoria Interrante. 2017. Impact of visual and experiential realism on distance perception in VR using a custom video see-through system. In *Proceedings of the ACM symposium on applied perception*. 1–8.

Eric-Jan Wagenmakers, Maarten Marsman, Tahira Jamil, Alexander Ly, Josine Verhagen, Jonathon Love, Ravi Selker, Quentin F. Gronau, Martin Šmíra, Sacha Epskamp, et al. 2018. Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic bulletin & review* 25, 1 (2018), 35–57.

Mark Wagner. 1985. The Metric of Visual Space. *Perception & Psychophysics* 38, 6 (1985), 483–495.

A Appendix

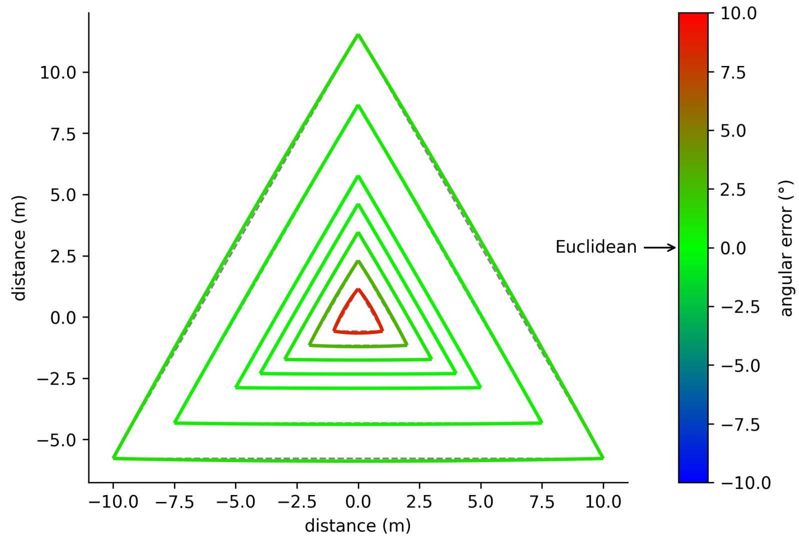

Figure 6: Group mean pointing error by distance arcplot across all participants, including distances 2, 4, 6, 8, 10, 15, and 20 m. Red denotes positive angular error (elliptic triangles), green indicates zero error (Euclidean), and blue indicates negative error (hyperbolic). For visualization purposes, triangles are represented using Bézier curves.

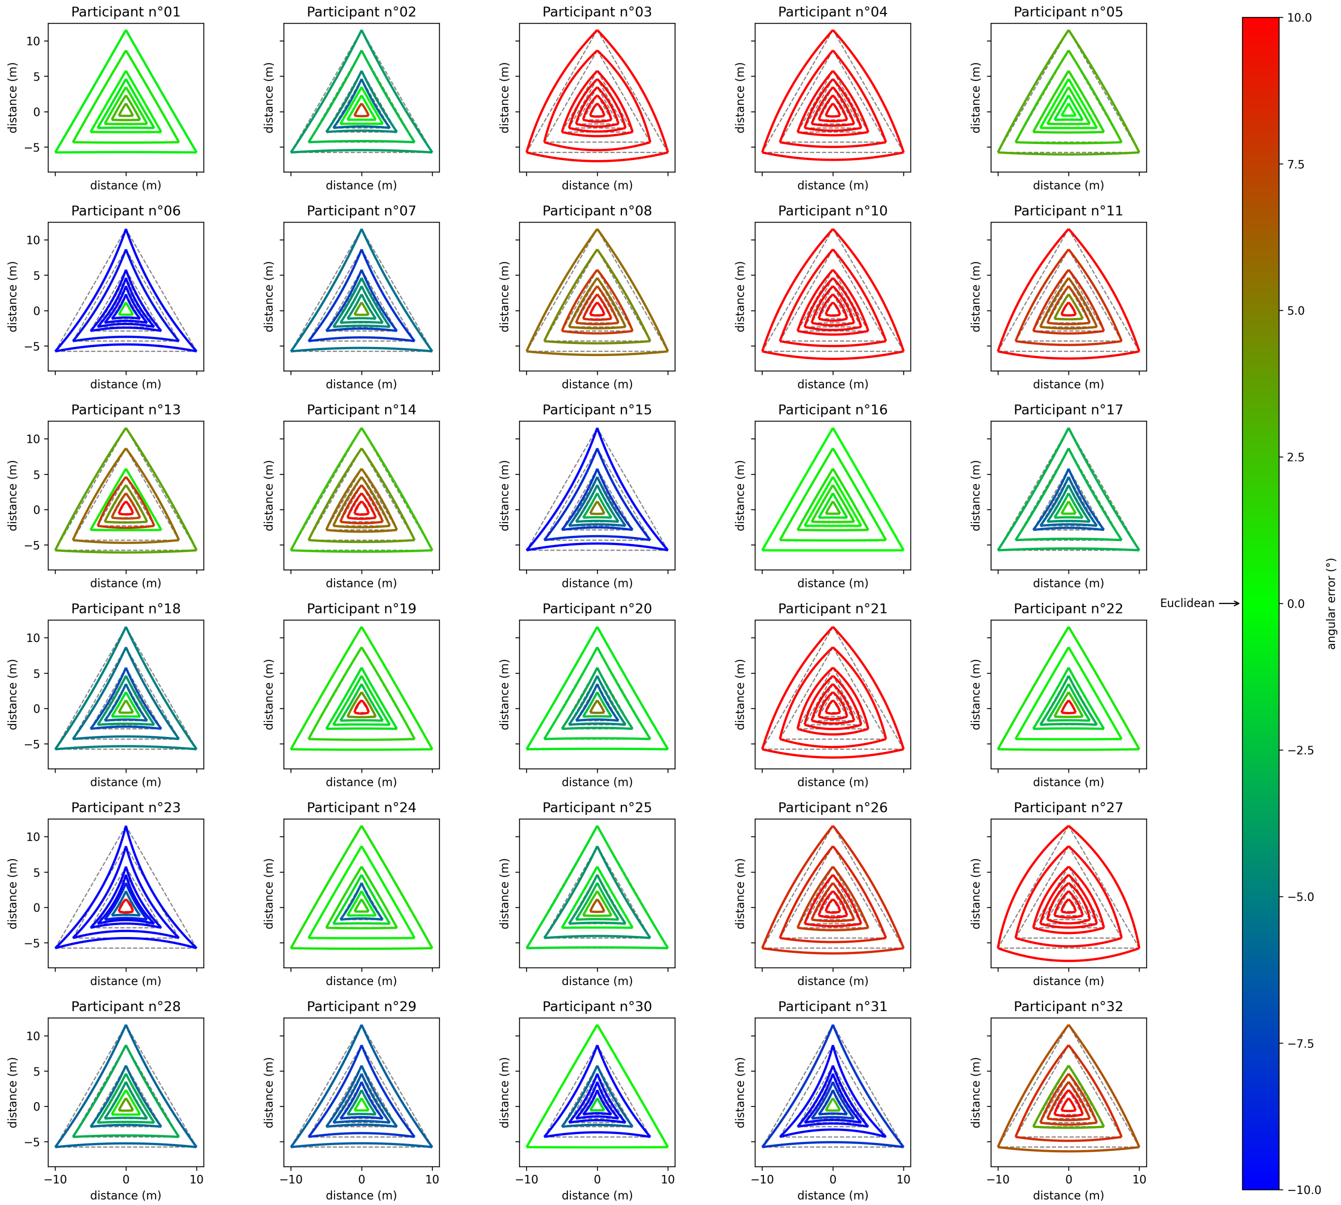

Figure 7: Mean error arcplot of all participants at chosen distances, 2, 4, 6, 10, 15, and 20 m. Red denotes a positive angular error corresponding to an elliptic triangle, whereas green indicates zero angular error (Euclidean) and blue denotes a negative angular error (hyperbolic). For visualization purposes, triangles are represented using Bézier curves.