

# Evaluating Foveated Frame Rate Reduction in Virtual Reality for Head-Mounted Displays

Christopher Flöter  
christopher.floeter@uni-ulm.de  
Visual Computing Group  
Ulm University  
Ulm, Germany

Sergej Geringer  
sergej.geringer@visus.uni-stuttgart.de  
Visualization Research Center  
University of Stuttgart  
Stuttgart, Germany

Guido Reina  
guido.reina@visus.uni-stuttgart.de  
Visualization Research Center  
University of Stuttgart  
Stuttgart, Germany

Daniel Weiskopf  
daniel.weiskopf@visus.uni-stuttgart.de  
Visualization Research Center  
University of Stuttgart  
Stuttgart, Germany

Timo Ropinski  
timo.ropinski@uni-ulm.de  
Visual Computing Group  
Ulm University  
Ulm, Germany



Figure 1: In our user study, participants tracked objects in virtual reality under different frame rates and scene environments. Frame  $n$  shows the fully rendered environment, in frame  $n + 1$ , the periphery is unchanged and only the foveal region is rendered freshly. Forest Asset courtesy of Oleh Lila, used with permission.

## ABSTRACT

Foveated rendering methods usually reduce spatial resolution in the periphery of the users' view. However, using foveated rendering to reduce temporal resolution, i.e., rendering frame rate, seems less explored. In this work, we present the results of a user study investigating the perceptual effects of foveated temporal resolution reduction, where only the temporal resolution (frame rate) is reduced in the periphery without affecting spatial quality (pixel density). In particular, we investigated the perception of temporal resolution artifacts caused by reducing the frame rate dependent on the eccentricity of the user's gaze. Our user study with 15 participants was conducted in a virtual reality setting using a head-mounted display. Our results indicate that it was possible to reduce average rendering costs, i.e., the number of rendered pixels, to a large degree before participants consistently reported perceiving temporal artifacts.

## CCS CONCEPTS

• **Computing methodologies** → *Perception*; **Virtual reality**; • **Human-centered computing** → *User studies*.

## KEYWORDS

Temporal resolution reduction, frame rate reduction, foveated rendering, virtual reality

## ACM Reference Format:

Christopher Flöter, Sergej Geringer, Guido Reina, Daniel Weiskopf, and Timo Ropinski. 2025. Evaluating Foveated Frame Rate Reduction in Virtual Reality for Head-Mounted Displays. In *2025 Symposium on Eye Tracking Research and Applications (ETRA '25)*, May 26–29, 2025, Tokyo, Japan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3715669.3725870>

## 1 INTRODUCTION

We address the problem of rendering performance for virtual reality (VR) delivered on head-mounted displays (HMDs), with applications ranging from computer games all the way to immersive analytics and visualization. Current HMD technology still falls short of matching human vision, which would require a pixel density of 60 pixels/degree (ppd), a refresh rate of 1800 Hz, and a field of view (FOV) of  $157^\circ \times 135^\circ$  [Cuervo et al. 2018]. Today's HMDs, despite their lower resolution, rely on gaze-directed [Levoy and Whitaker

1990] or foveated rendering [Guenther et al. 2012] to reduce computational costs and achieve sufficient rendering speed. However, such adaptive rendering is typically restricted to reducing the spatial resolution away from the fovea. In contrast, adapting the temporal resolution is comparatively unexplored [Mohanto et al. 2022], and existing methods for reducing temporal resolution [Dorr et al. 2005; Weier et al. 2016] are coupled to simultaneously reducing the spatial resolution.

Our goal is to further advance gaze-adaptive rendering by including foveated temporal resolution reduction, i.e., selectively reducing the image refresh rate (frame rate) in the periphery. Therefore, we want to establish a baseline for the perception of temporal artifacts in a virtual reality setting and determine whether it is possible to reduce the temporal resolution in a gaze-dependent way that is not perceivable to the user. Our method reduces the temporal resolution in an image region depending on the pixel’s distance to the user’s gaze. This defines concentric segments in screen-space around the gaze point that are rendered at different frame rates. Based on this setup, a user study ( $N = 15$ ) was conducted with HMDs to determine the perception of the resulting temporal artifacts and the extent to which they are considered distracting by the user when perceived. Figure 1 illustrates the setup of the user study. The contributions of this paper are:

- We discuss temporal rendering artifacts that occur due to our method of foveated temporal resolution reduction.
- We present the results of a user study in VR, reporting subjective user tolerance to the rendering artifacts.

Our results show we can reduce pixel rendering costs by up to 63.6% without users feeling uncomfortable.

## 2 RELATED WORK

*Foveated rendering* is often used to improve rendering performance by reducing rendering quality and workload in the periphery unnoticeable to the user [Mohanto et al. 2022]. In modern HMDs, gaze information is readily available from built-in eye tracking. In turn, beyond foveated rendering, virtual- and augmented reality applications use eye tracking to also provide gaze-based user interactions [Adhanom et al. 2023; Plopski et al. 2022].

Mohanto et al. [2022] classify foveated rendering methods into four categories: (1) *Adaptive resolution* methods lower the spatial render resolution outside the foveated region. An implementation was presented by Guenther et al. [2012]. This approach is most commonly used in consumer hardware. (2) *Geometric simplification* reduces the level of detail of objects dependent on gaze. It is often implemented in combination with adaptive resolution rendering. An implementation is described by Murphy and Duchowski [2001]. (3) *Shading simplification and chromatic degradation* reduce the computational effort for calculating the pixel color in a gaze-dependent fashion. An implementation is described by Stengel et al. [2016]. (4) *Temporal deterioration* techniques reduce the temporal resolution by reusing already rendered pixels from previous frames, therefore lowering the number of pixels to be rendered.

*Temporal resolution reduction* has the goal of re-drawing an image, or certain parts of it, at a lower update rate (frame rate) than strictly necessary. As such, this type of rendering technique need not be strictly gaze-dependent. Nehab et al. [2006] presented an approach

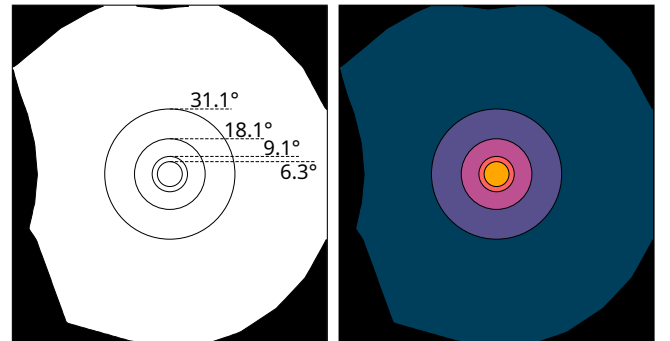
to re-project (re-use) a fragment when its value has not changed, saving computation costs by delaying pixel updates. When perfectly implemented, in theory, there would not be any reduction in render quality. Another method by Denes et al. [2019] decreases the render resolution every second frame. The temporal flickering caused by this reduction in quality, given a high enough refresh rate, causes no noticeable loss in quality for users. Weier et al. [2016] adapted the sampling density of rays in a ray tracing rendering pipeline based on the eccentricity of the user’s gaze. Missing pixels are generated using information from the previous frame in combination with the sample rays of the current frame. Dorr et al. [2005] defined gaze-dependent image regions and then decreased the temporal resolution by interpolating the currently rendered pixel with the previously displayed value. The latter two approaches [Dorr et al. 2005; Weier et al. 2016] facilitate post-processing effects, applying Gaussian blur on foveated regions to reduce the spatial resolution. Unfortunately, this makes it difficult to determine the extent to which the temporal artifacts affect the perception of the image quality, since they are always coupled with spatial resolution reduction. In this paper, we investigate a more basic rendering strategy compared to previous literature: we re-use pixels from previous frames without applying spatio-temporal re-projection.

## 3 METHOD: TECHNIQUE AND RENDERING ARTIFACTS


We investigate how reducing the temporal resolution in a gaze-dependent manner is perceived by the user in a VR setting, i.e., we combine the concept of temporal resolution (frame rate) reduction with foveated rendering.

### 3.1 Technique: Foveated Temporal Resolution Reduction

Given the gaze position on the screen, we define concentric circular image segments (foveal/peripheral regions) around the gaze point, depending on the angle spanned by the image pixels and the visual axis, i.e., gaze direction. We define the image segments using eccentricity angles as provided by Mohanto et al. [2022] (see Figure 2)—resulting in *five* image segments spanning from the



**Figure 2: (Left) Eccentricity mask defining fovea/periphery image regions, with diameters chosen according to Mohanto et al. [2022]. (Right) Image regions are color-coded according to frame rate reduction for a 12345 frame rate configuration.**



**Figure 3: Illustration of the rendering pipeline. Green boxes implement foveated rendering, blue boxes implement temporal resolution reduction. The images depict an object moving from the left to the right. HMD illustration courtesy of user Juhele “Jan Helebrant” via Open Clip Art Library, used under CC0 license. Suzanne Asset courtesy of TurboSquid, used with permission.**

fovea to the periphery. During rendering, we update different image segments at different frame rates: regions in the periphery are updated at slower frame rates than in the fovea. Figure 2 depicts the image segments based on eccentricity, the rendering pipeline for the method is illustrated in Figure 3, and Figure 4 shows temporal artifacts resulting from this rendering approach.

### 3.2 Temporal Rendering Artifacts

Updating different image regions at different intervals may introduce temporal rendering artifacts in the form of image discontinuities at region transitions. This effect is reinforced by our decision to forgo spatio-temporal re-projection, and instead presenting the image regions rendered at different frame rates verbatim to the user. In the following, we inspect possible causes of these artifacts, and how they might negatively impact the user’s experience. We assume that the VR scene consists of actively moving or stationary objects and a stationary environment (background).

*Object and Environment Tearing.* Assuming the user to be stationary, moving objects crossing region boundaries may appear to be cut off. Other factors influencing this effect include object *color*, *geometry*, *size*, *movement speed*, and *movement extent*. Assuming the user moves their head, image tearing of the scene environment (background) will appear across image regions, caused by the non-uniform temporal resolution. This effect is determined by the *visual variety* of the environment and the motion induced by *user behavior*.



**Figure 4: Temporal artifacts resulting from object movement (Left) and camera movement (Right) with our rendering method. HDRI Room/Attic (left) and Vestibule (right) Assets courtesy of user “Sergej Majboroda” via Poly Haven, used under CC0 license.**

The described rendering artifacts are exemplified in Figure 4. The perceived severity of these temporal artifacts depends on the frame rate reduction applied to the individual image regions, i.e., the local rendering delay, and the movement speed of the user and/or objects in the scene. Thus, the amount of tearing that occurs globally depends on the number of image regions with distinct frame rates.

*Image Homogeneity and Scene Motion.* The extent to which the scene affects the perception of the frame rate reduction depends on the image complexity, i.e., the cohesiveness of pixel areas with the same color, which in turn depends on the variety of scene objects and how they are shaded. For example, flat or cell shading will lead to larger areas of identically colored pixels than physically-based rendering. The latter would lead to more perceived tearing and temporal artifacts. Figure 5 (a)–(e) shows the five scenes we used for our user study, ranging from photorealistic HDRI scans and cell shading all the way to a constant-color skybox. The perceived movement of objects, and thus induced tearing, is also influenced by the movement of the camera. This is especially the case in a VR environment, where camera movement is immediately driven by the physical head movement of the user.

### 3.3 Implementation


We implemented our rendering method as a VR scene in the Unity engine [Unity Technologies 2025] version 2022.3, using a post-processing step to combine the latest frame with previously rendered frames according to the fovea/periphery image regions and pre-defined frame rate reduction configurations. The implementation we describe in the following also covers the setup and range of possible conditions in our conducted user study.

*Scene Setup.* We implement a scene with a static environment, rendering one of the textures shown in Figure 5 as a skybox. Further, the scene contains objects with distinct shape and color moving independently and at random in front of the user. Figure 1 depicts the VR scene with several moving objects in front of the user in a cell-shaded environment.

*Temporal Resolution Reduction in Image Regions.* Our implementation assigns different individual frame rates to different image regions (see Figure 2). Starting from the full frame update rate  $1/r$ , denoting an image update every frame (in VR usually at 90 Hz), we define the *frame update rate*  $1/r$  as (delayed) image updates every  $r$  frames, with  $r \in \{1, 2, 3, 4, 5\}$ . A *frame rate configuration* is defined as a tuple  $f = (f_1, f_2, f_3, f_4, f_5) \in \{1, 2, 3, 4, 5\}^5$  of frame update rates

**Table 1: Frame rate configurations (FRCs) and corresponding percentage of pixels that need to be drawn on average per frame.**

FRC	11111	11112	11122	12222	11223	11233	12233	12234	12334	12345
Percentage of pixels	100%	57.6%	52.2%	50.2%	36.4%	34.6%	34.4%	27.3%	26.7%	21.5%



**Figure 5: Screenshots of the environments: (a, b) HDRI scans, (c, d) pre-rendered scenes converted into cube maps, and (e) cube map with a constant color. (f) Object shapes moving in front of the user. HDRI Street Asset courtesy of users “Dimitrios Savva” and “Jarod Guest” via Poly Haven, used under CC0 license. HDRI Room/Attic Asset courtesy of user “Sergej Majboroda” via Poly Haven, used under CC0 license. Apartment Asset courtesy of Alexis Jose Guevara, used with permission. Forest Asset courtesy of Oleh Lila, used with permission.**

( $1/f_1, 1/f_2, 1/f_3, 1/f_4, 1/f_5$ ) for the five image regions. Using a shorter notation, the frame rate configuration  $11111 \equiv (1, 1, 1, 1, 1)$  denotes frame update rates  $1/f_1$  for all image regions, i.e., normal rendering without delays, whereas  $12345 \equiv (1, 2, 3, 4, 5)$  denotes an update rate of  $1/f_1$  in the central fovea, and decreasing update rates  $1/f_2, 1/f_3, 1/f_4, 1/f_5$  toward the periphery. For our user study, we consider frame rate configurations with only monotonically decreasing update rates ( $f_i \leq f_{i+1}$ ), where adjacent image regions are updated with at most one frame difference ( $f_{i+1} = f_i + 1$  or  $f_{i+1} = f_i + 0$ ). To limit the session time per participant in the user study we only consider 10 frame rate configurations (out of 16 possible configurations according to the above criteria): 11111, 11112, 11122, 11223, 11233, 12222, 12233, 12234, 12334, 12345.

*Metric for Amortized Pixel Rendering Savings.* To determine the average reduction of pixels rendered per frame for configuration  $f$ , we define the pixel rendering rate  $P = \frac{1}{a} \sum_i p_i = \frac{1}{a} \sum_i \frac{a_i}{f_i}$ . The amount of pixels rendered on average is  $p_i = \frac{a_i}{f_i}$ , with  $a_i$  being the number of pixels in image region  $i$ ,  $f_i$  the frame update rate in that region, and  $a$  the total number of pixels on the screen. The amortized pixel rendering rates for our frame rate configurations (FRCs) are reported in Table 1.

## 4 USER STUDY

The user study investigates the following research questions regarding foveated temporal resolution reduction:

- RQ1 To what extent are temporal artifacts perceived by the user in a VR setting wearing an HMD?
- RQ2 To what extent does the environment influence the perception of temporal artifacts?
- RQ3 To what extent does the behavior of the user influence the perception of temporal artifacts?

### 4.1 Study Design, Participants, and Apparatus

*Study Design.* The study was conducted in a controlled lab environment with a within-subject design. Participants completed tasks under different stimuli in a VR setting while wearing an HMD. Study tasks and stimuli were chosen to achieve a total study time of under 60 minutes. As *independent variables*, we controlled frame rate configurations (10 options) and scene environments (5 options) across two tasks (gaze follow/center task). As *dependent variables*, we collected subjective user responses on a 5-point Likert scale to questions regarding the perception of rendering artifacts, the grade of perceived distraction, and user discomfort. Investigating VR sickness [Kim et al. 2018; Tcha-Tokey et al. 2016; Wang et al. 2023] in our setting would also be desirable but beyond the targeted study duration.


*Participants.* We recruited 15 participants from the computer science and medical faculty. All participants had normal vision, they reported no visual impairment while wearing the HMD. Most participants had very little to no prior experience with VR headsets. Participants signed a consent form after being informed about the purpose of the study and their right to withdraw at any time.

*Apparatus.* To ensure interactive rendering rates of 90 Hz required for a VR setting, a computer with i7-4790 CPU and RTX 2080 GPU was used. The HMD used was a HTC VIVE Pro Eye with integrated eye tracking. With the frame rate configurations discussed previously, the lowest refresh rate for selected image regions was 18 Hz. The VR software for the study was the implementation described in Section 3.3.

### 4.2 Stimuli and Tasks

*Stimuli.* We employed stimuli consisting of a combination of the ten frame rate configurations and five scene environment settings, along with objects moving in front of the user (see Section 3.3).





**Figure 6: Boxplots showing results of user questions regarding Perception, Distraction, and Discomfort during the study. The x-axis denotes the tested frame rate configurations, the y-axis denotes user answers on the five-point Likert scale numerically coded from 0 to 4 (merging answers over all scene environments and participants per FRC). Left: Gaze follow task. Right: Gaze center task.**

Each stimulus was presented to a participant once in each of the two tasks, for a total of  $10 \times 5 \times 2 = 100$  stimuli (study rounds). For the definition of the fovea-centered image regions,  $1.1^\circ$  was added to the size of the original regions to accommodate for the tolerance in gaze determination of the used HMD. The number and size of the image regions were constant for the entire trial. The three moving objects were changed in color and geometry every round. During Task 1, all three objects had the same random color and geometry. For Task 2, the color and geometry of all objects were independent and random. During the study, the object speed and movement area were fixed for each round because increasing the speed of the objects, or the extent of their movement, would have changed the temporal artifacts in different image regions.

**Tasks.** We employed two different tasks based on the factors influencing temporal rendering artifacts discussed in Section 3.2. Prior research suggested that the perception of temporal artifacts depends on whether the user is tracking an object or focusing on a stationary point [Weier et al. 2016]. The *gaze follow task* (Task 1) consisted of three moving objects of the same shape and color. These objects moved independently, at random, in a specified region for 15 seconds. At the beginning of the round, one object was highlighted for two seconds. At the end of the round, the participant was asked to select the object highlighted in the beginning. This was done to encourage the user to follow a moving object. The *gaze center task* (Task 2) showed a small stationary circle for the entirety of the round. At the beginning of the round, the circle was green and turned red after 5 to 10 seconds. Once the circle turned red, the participant was asked to react as fast as possible by pressing a button on a controller. To ensure the presence of temporal artifacts, the moving objects from the first task were also present, whereby each of the three objects had a random shape and color.

### 4.3 Procedure

Participants were introduced to the VR controls and calibrated the HMD eye tracker. Participants then performed four test trials without temporal resolution reduction to get familiar with the VR scene and to get a frame of reference on the quality of the VR environment. For the duration of the trial, the participants were stationary and seated.

During the study, participants carried out the tasks under the stimuli described above, performing one task after the other with a randomized order of its stimuli. Learning effects were compensated by counter-balancing the order of the tasks. After every round, the participants provided feedback in the virtual environment, answering questions whether they could perceive any visual artifacts (FQ1 Perception), if the perceived artifacts were distracting (FQ2 Distraction), and if they were feeling any discomfort (FQ3 Discomfort), rating their subjective experience about the round on a 5-point Likert scale (“none,” “very little,” “moderate,” “strong,” and “very strong”). Apart from answering the feedback questionnaire, participants had no breaks between stimuli or tasks.

## 5 RESULTS

Figure 6 summarizes answers for the three measured dependent variables across both tasks. For the gaze follow task, discomfort stays low, for the gaze center task it is even negligible (green boxes). Over all tasks and trials, discomfort values of two and above (“moderate” or worse) were rare except for one specific participant. This is reflected in the outliers for the discomfort answers. Perception (blue boxes) and distraction (orange boxes) variables increase more quickly for the gaze follow task than for the gaze center task. Looking at the measured frame rate configurations more closely, the baseline FRC 11111, apart from outliers, predominantly shows the lowest values (“none”) for the perception, distraction, and discomfort variables across both tasks. For the gaze center task, this also applies to FRC 11112.

For the gaze follow task, the upper quartile of *perception* rises to “very little” as soon as the periphery gets less than full update rates, and becomes “moderate” as soon as at least three regions have reduced update rates. The upper quartile of perception only becomes “very strong” when the worst configuration is used (FRC 12345). For the gaze center task, the perception upper quartile is only “strong” for FRC 12345, the four preceding configurations still are “moderate” (only the best of these is “moderate” for the gaze follow task). The upper quartiles of *distraction* rise more slowly, capping at two “strong” configurations for the gaze follow task, and two “moderate” configurations for the gaze center task. The gaze center task then has “very little” distraction down to all four outer

regions with halved update rates. The gaze follow task needs the two outermost regions to have halved update rates for distraction to rise to “very little,” followed by three “moderate” configurations. *Discomfort* begins to be noticeable in the gaze follow task, starting from FRC 12233. Even with the worst configuration, FRC 12345, the upper quartile for discomfort still only reaches “very little.”

User feedback grouped by the five *scene environments* is very similar for identical frame rate configurations. The blank background seems to cause slightly lower (better) scores across all frame rates, but only for the gaze center task. Somewhat unexpectedly, for the gaze follow task, results for the blank background are not better than for other environments. We provide additional plots detailing these aspects, as well as others, in the supplemental material.

## 6 DISCUSSION

We investigated our first research question, (RQ1) perception of temporal artifacts by the user (see Section 4), by collecting user feedback across the dimensions perception of artifacts (FQ1), their distraction (FQ2), and user discomfort (FQ3). With our chosen frame rate configurations, we wanted to cover a large range of configurations to allow detailed analysis of user sensitivity to frame rate reduction, without having to invest too many resources into strictly finding just-noticeable differences (JNDs).

The results presented in Section 5 show that for both tasks, (1) users increasingly perceive frame rate reduction with smaller frame update rates and (2) frame rate reduction is perceived as more distracting or uncomfortable with smaller update rates toward the foveal region. For example, shifting from configuration 11223 to 11233 (one additional update rate 3 in the periphery) still shows similar results across all three feedback dimensions, with the latter configuration being rated slightly worse. At the same time, configurations 12222 and 11223 are rated similarly, although the former uses an update rate 2 in all regions except for the most central, whereas the latter shows the ground truth (update rate 1) in a larger foveal region and an update rate of 3 in the outermost periphery. However, for all of these three FRCs, less than 50% of users report no or barely perceivable artifacts and, at the same time, not being distracted or feeling uncomfortable (see median markers in the boxplots).

Regarding the influence of user behavior on the perception of temporal artifacts (RQ3), results of the *gaze follow* and *gaze center* user tasks show a clear tendency for users to perceive artifacts more, being distracted by them or even feeling uncomfortable, when they have to actively follow objects. A little bit surprising for us is the result from the gaze follow task that even for small frame update rates, e.g., configuration 12345, although users perceive artifacts and are moderately distracted by them, they still report low discomfort scores mostly in the range of “none” or “very little” discomfort. None of the participants had to interrupt or abort the trial due to VR sickness. Two of the participants experienced moderate symptoms of VR sickness, one of them reported they knew they were prone to it, the other had no prior VR experience. Other forms of discomfort reported by participants were due to the HMD feeling uncomfortable and eye strain due to focusing over an extended period of time. For the gaze center task, users mostly reported low discomfort across all FRCs. This could be attributed to our rendering method

producing temporal artifacts only after large head movements by the user, which were discouraged in the second task by design.

Regarding influence of the scene environment (RQ2), we found no significant differences in user feedback for frame rate configurations when grouped by scene environment. Surprisingly, the monotone-colored scene background was rated similar to the other environments across both tasks.

Literature suggests that VR sickness levels increase when exposed to a VR environment over a long period of time [Moss and Muth 2011; Stanney et al. 2003]. Thus, we assumed that the VR setting and our stimuli would cause increasing discomfort over the study length. We found that about half of the participants did not experience significant discomfort at all. Only two participants showed a clear increase in discomfort over time. We need further studies to ascertain whether there is a dependency or not. For example, the length of our stimuli (about 10–15 seconds) may be too short, and longer exposure to frame rate reduction, e.g., 1–2 minutes or more, may show greater discomfort or sickness effects on participants. There was no measurable learning effect observed between the different tasks (see the supplemental material).

## 7 ETHICAL CONSIDERATIONS AND PRIVACY

General ethical and privacy concerns related to measuring eye tracking data apply to this work as well. Because foveated rendering does not need to store such data privacy issues can be mitigated.

## 8 CONCLUSION AND FUTURE WORK

In this paper, we evaluated foveated frame rate reduction, a foveated rendering method where the frame update rate is varied in different regions of the rendered image. The results of our user study show that it is possible to substantially reduce the amount of rendered pixels in head-mounted VR before users consistently notice the altered frame rate updates and report discomfort. In our specific setting, a reduction by 63.6% (FRC 11223) is feasible. This observation is encouraging, indicating the potential of temporal resolution reduction.

However, further research is required as the intuitive judgment of our proposed technique would expect more and more severe discomfort and VR sickness feedback by users—while most of our participants reported barely feeling discomfort or being distracted by the rendering technique. In particular, extended user studies could include more participants and a wider range of stimuli and tasks. Another direction of future research could integrate spatio-temporal re-projection into our rendering method.

## ACKNOWLEDGMENTS

The authors wish to thank the anonymous reviewers for their constructive and critical feedback.

We also thank the credited asset authors and organizations for granting permission to use their works for publication in this paper.

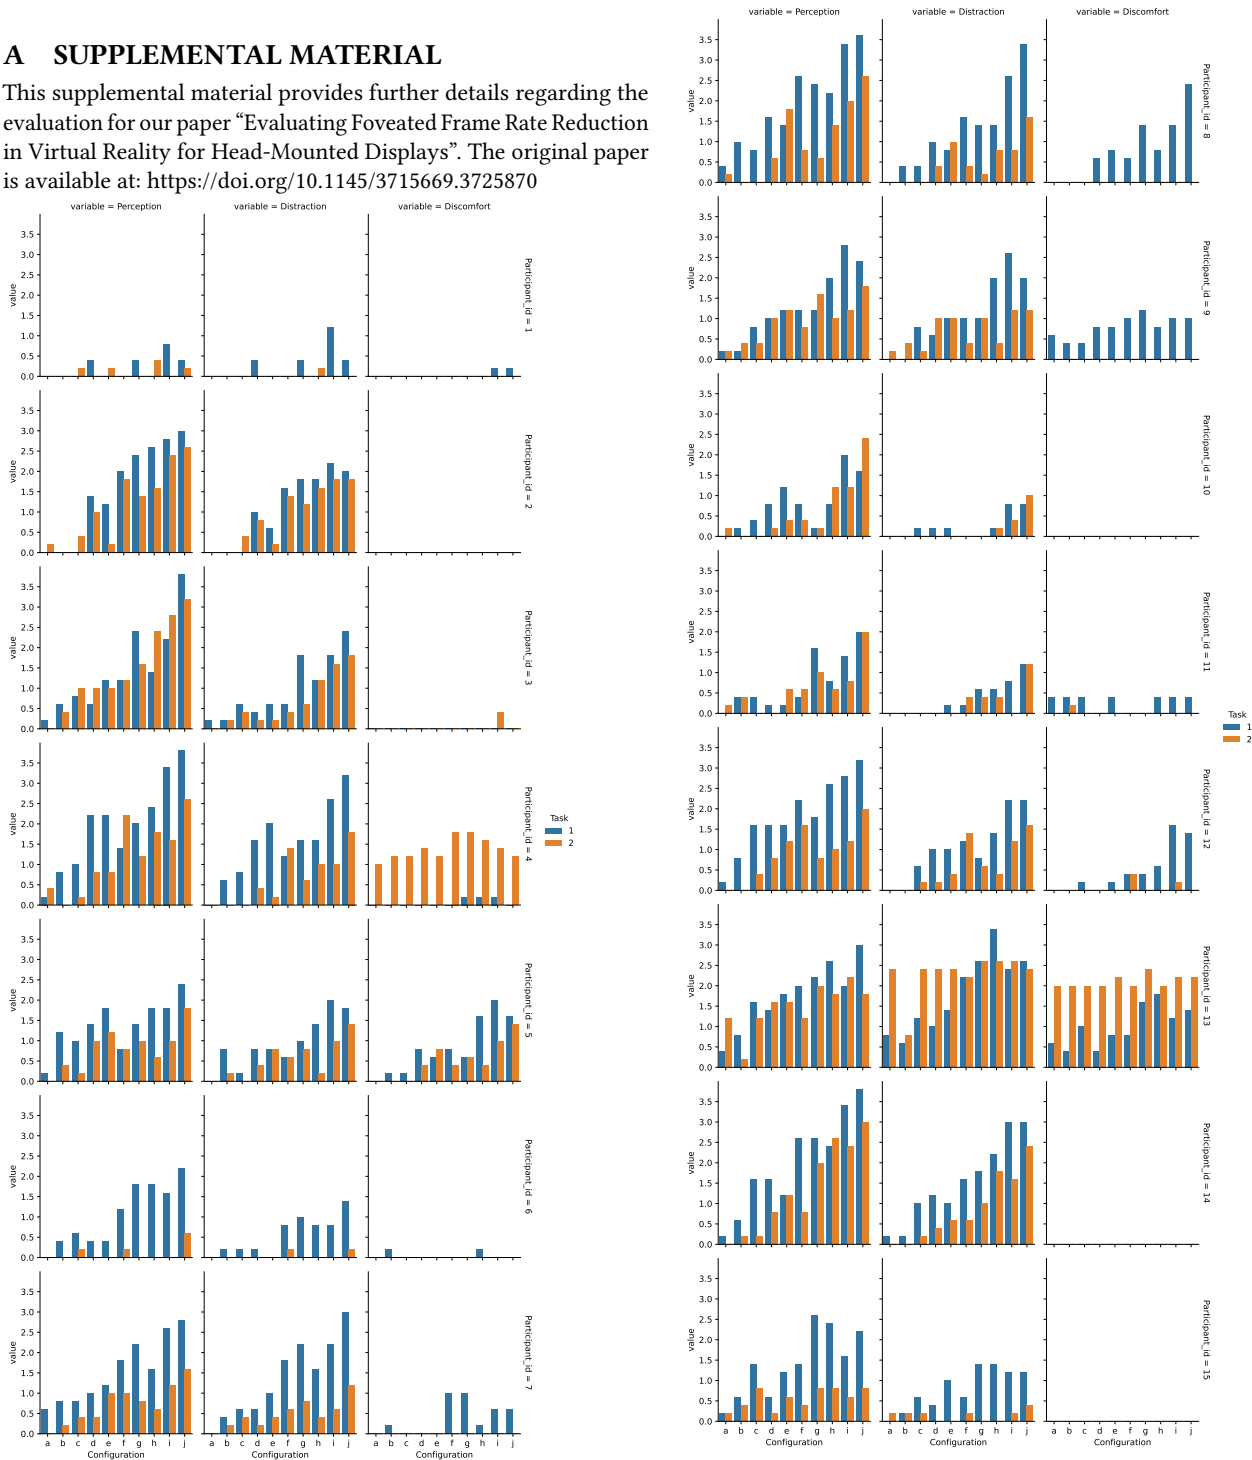
This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project-ID 251654672—TRR 161, and the Ministerium für Wissenschaft, Forschung und Kunst (MWK, Ministry of Science, Research and Arts) Baden-Württemberg, Germany—Project Virtuelle Kollaborationslabore Baden-Württemberg, Phase 1–3—KoLabBW.

## REFERENCES

- Isayas Berhe Adhanom, Paul MacNeilage, and Eelke Folmer. 2023. Eye tracking in virtual reality: A broad review of applications and challenges. *Virtual Reality* 27, 2 (2023), 1481–1505. <https://doi.org/10.1007/s10055-022-00738-z>
- Eduardo Cuervo, Krishna Chintalapudi, and Manikanta Kotaru. 2018. Creating the perfect illusion: what will it take to create life-like virtual reality headsets?. In *Proceedings of the International Workshop on Mobile Computing Systems and Applications (HotMobile)*. Association for Computing Machinery, New York, NY, USA, 7–12. <https://doi.org/10.1145/3177102.3177115>
- Gyorgy Denes, Kuba Maruszczyk, George Ash, and Rafal K. Mantiuk. 2019. Temporal resolution multiplexing: Exploiting the limitations of spatio-temporal vision for more efficient VR rendering. *IEEE Transactions on Visualization and Computer Graphics* 25, 5 (2019), 2072–2082. <https://doi.org/10.1109/tvcg.2019.2898741>
- Michael Dorr, Thomas Martinetz, Martin Böhme, and Erhardt Barth. 2005. Visibility of temporal blur on a gaze-contingent display. In *Proceedings of the Symposium on Applied Perception in Graphics and Visualization (APGV)*. Association for Computing Machinery, New York, NY, USA, 33–36. <https://doi.org/10.1145/1080402.1080407>
- Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D graphics. *ACM Transactions on Graphics* 31, 6, Article 164 (2012), 10 pages. <https://doi.org/10.1145/2366145.2366183>
- Hyun K. Kim, Jaehyun Park, Yeongcheol Choi, and Mungyeong Choe. 2018. Virtual reality sickness questionnaire (VRSQ): Motion sickness measurement index in a virtual reality environment. *Applied Ergonomics* 69 (2018), 66–73. <https://doi.org/10.1016/j.apergo.2017.12.016>
- Marc Levoy and Ross Whitaker. 1990. Gaze-directed volume rendering. *Computers and Graphics* 24, 2 (1990), 217–223. <https://doi.org/10.1145/91394.91449>
- Bipul Mohanto, ABM Tariqul Islam, Enrico Gobbetti, and Oliver Staadt. 2022. An integrative view of foveated rendering. *Computers and Graphics* 102 (2022), 474–501. <https://doi.org/10.1016/j.cag.2021.10.010>
- Jason D. Moss and Eric R. Muth. 2011. Characteristics of head-mounted displays and their effects on simulator sickness. *Human Factors* 53, 3 (2011), 308–319. <https://doi.org/10.1177/0018720811405196>
- Hunter Murphy and Andrew T. Duchowski. 2001. Gaze-Contingent Level Of Detail Rendering. In *Eurographics Short Presentations*. Eurographics Association, Eindhoven, Netherlands, 10 pages. <https://doi.org/10.2312/egs.20011004>
- Diego Nehab, Pedro V. Sander, and John R. Isidoro. 2006. The real-time reprojection cache. In *SIGGRAPH Sketches (SIGGRAPH)*. Association for Computing Machinery, New York, NY, USA, 185–es. <https://doi.org/10.1145/1179849.1180080>
- Alexander Plopski, Teresa Hirzle, Nahal Norouzi, Long Qian, Gerd Bruder, and Tobias Langlotz. 2022. The eye in extended reality: A survey on gaze interaction and eye tracking in head-worn extended reality. *ACM Computing Surveys* 55, 3, Article 53 (2022), 39 pages. <https://doi.org/10.1145/3491207>
- Kay M. Stanney, Kelly S. Hale, Isabelina Nahmens, and Robert S. Kennedy. 2003. What to expect from immersive virtual environment exposure: Influences of gender, body mass index, and past experience. *Human Factors* 45, 3 (2003), 504–520. <https://doi.org/10.1518/hfes.45.3.504.27254>
- Michael Stengel, Steve Grogork, Martin Eisemann, and Marcus Magnor. 2016. Adaptive image-space sampling for gaze-contingent real-time rendering. *Computer Graphics Forum* 35, 4 (2016), 129–139. <https://doi.org/10.1111/cgf.12956>
- Katy Tcha-Tokey, Emilie Loup-Escande, Olivier Christmann, and Simon Richir. 2016. A questionnaire to measure the user experience in immersive virtual environments. In *Proceedings of the Virtual Reality International Conference (VRIC)*. Association for Computing Machinery, New York, NY, USA, Article 19, 5 pages. <https://doi.org/10.1145/2927929.2927955>
- Unity Technologies. 2025. *Unity Engine*. Unity Technologies. <https://unity.com/>
- Jialin Wang, Rongkai Shi, Wenxuan Zheng, Weijie Xie, Dominic Kao, and Hai-Ning Liang. 2023. Effect of frame rate on user experience, performance, and simulator sickness in virtual reality. *IEEE Transactions on Visualization and Computer Graphics* 29, 5 (2023), 2478–2488. <https://doi.org/10.1109/tvcg.2023.3247057>
- Martin Weier, Thorsten Roth, Ernst Kruijff, André Hinkenjann, Arsène Péard-Gayot, Philipp Slusallek, and Yongmin Li. 2016. Foveated real-time ray tracing for head-mounted displays. *Computer Graphics Forum* 35, 7 (2016), 289–298. <https://doi.org/10.1111/cgf.13026>

## A SUPPLEMENTAL MATERIAL

This supplemental material provides further details regarding the evaluation for our paper “Evaluating Foveated Frame Rate Reduction in Virtual Reality for Head-Mounted Displays”. The original paper is available at: <https://doi.org/10.1145/3715669.3725870>



**Figure 7: Likert scale user feedback per participant (rows, wrapped at participant 7) per configuration (x-axes) per task (bar color), averaged over backgrounds and shapes configurations.**



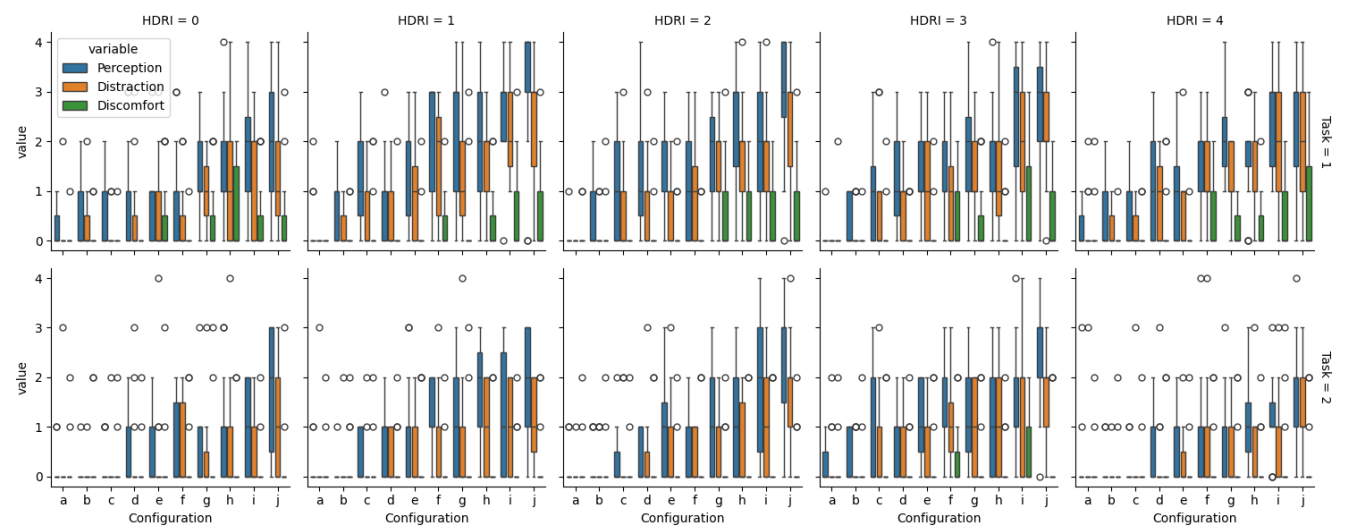


Figure 8: Boxplots showing results of user feedback grouped by scene environment. Top: Gaze follow task. Bottom: Gaze center task.




Figure 9: Discomfort rating occurrences per participant over all rounds.


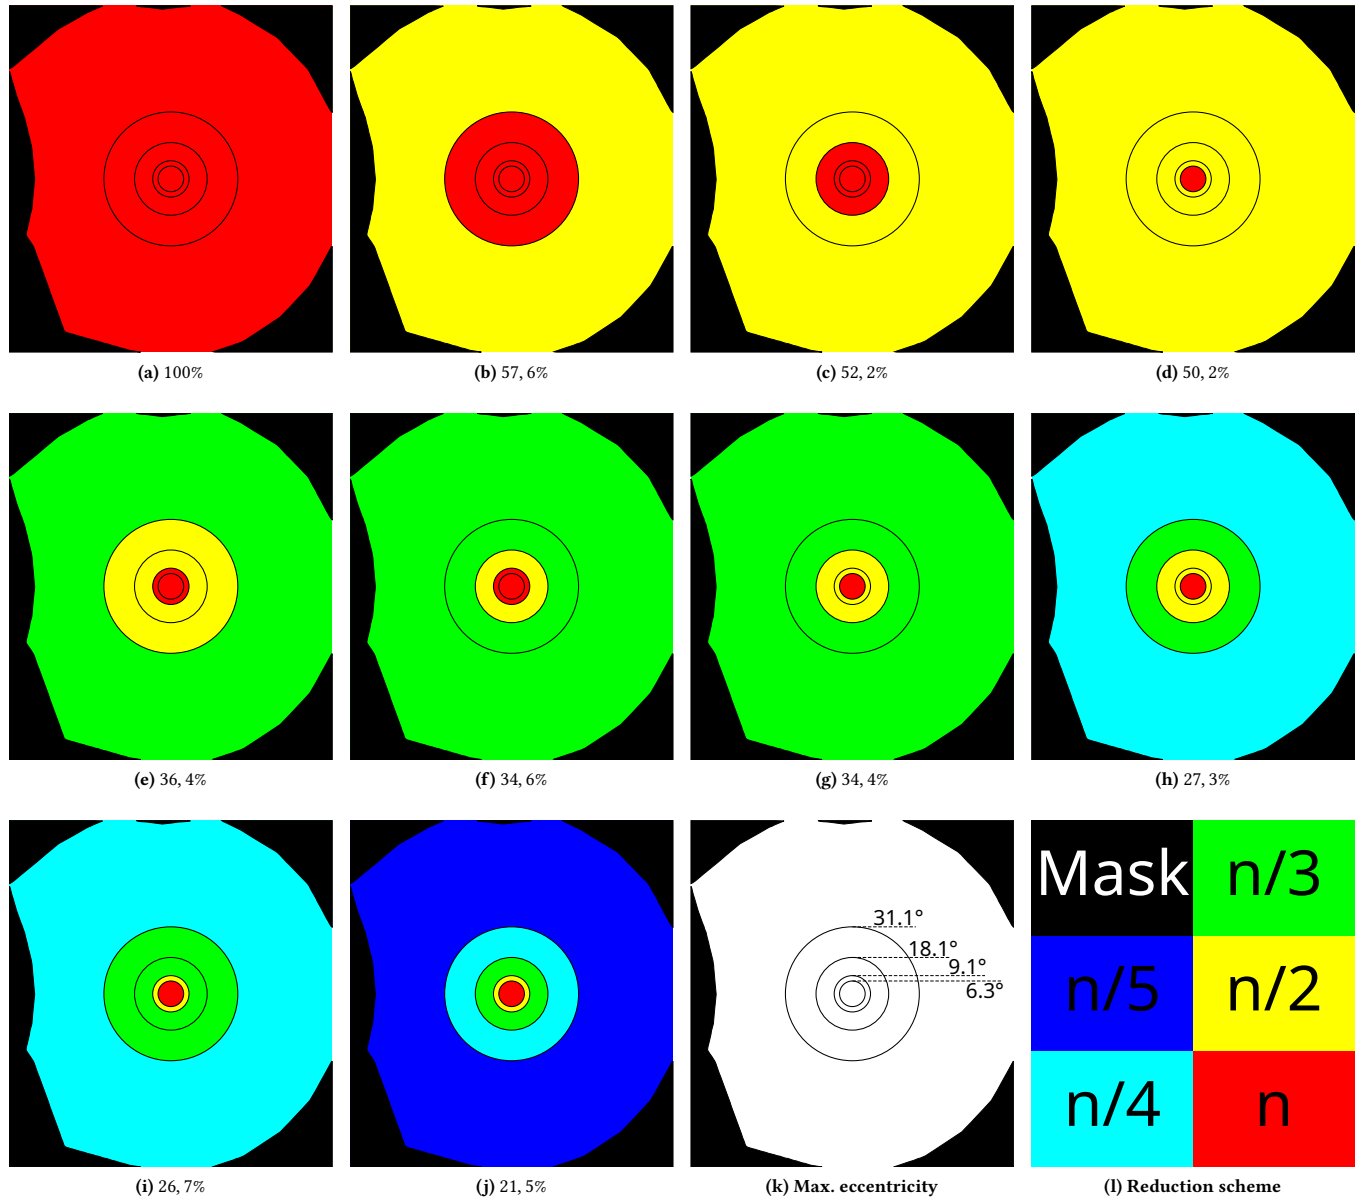
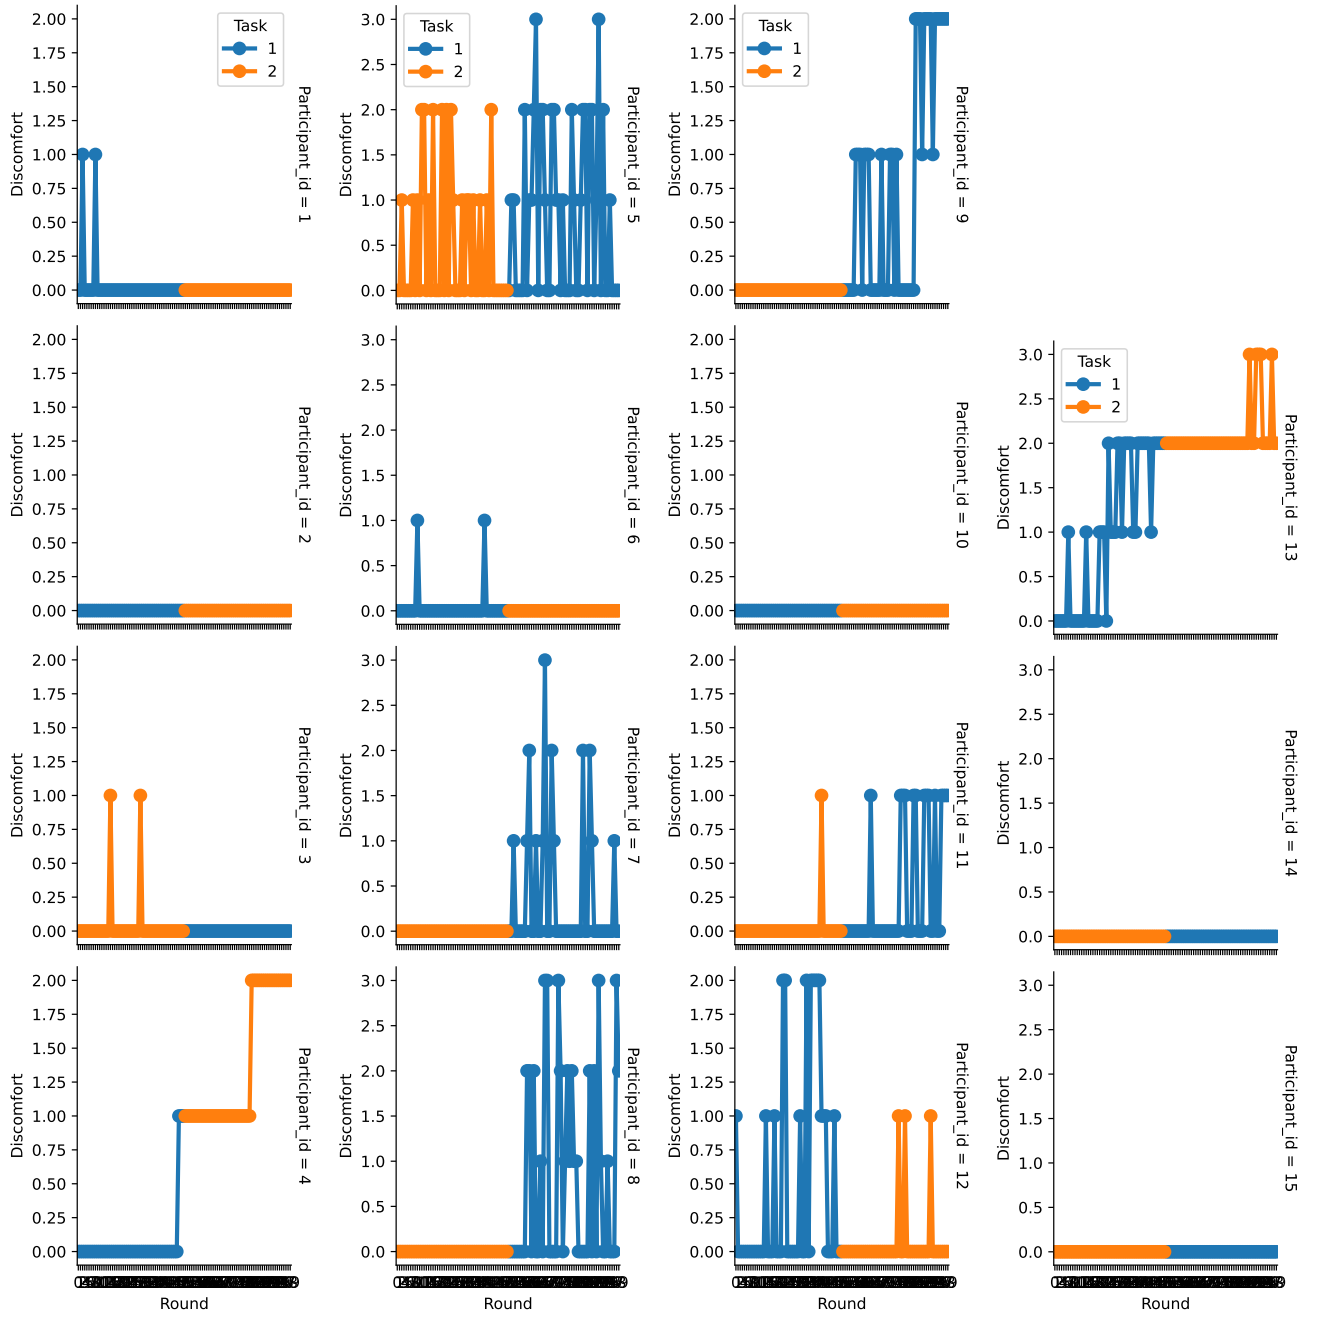


Figure 10: The geometry of the moving objects (left) and the color they can be (right). A Black border is added to ensure distinguishability of the objects even when overlapping.



**Figure 11: Visualization of the mask for each framework configurations used in the user study, with the user's gaze focused at the center of the display. The corresponding pixel reduction rate in %. The colors of the regions represent the effective frame rate for the region. The maximum degree of eccentricity for each region is given in (k).**



**Figure 12: Discomfort over rounds (x-axes) for each participant (rows), colored by task. Rows are wrapped at participants 4, 8, and 12. Note how only participants 4 and 13 have a clearly increasing discomfort over time.**