

The Relative Importance of Depth Cues and Semantic Edges for Indoor Mobility Using Simulated Prosthetic Vision in Immersive Virtual Reality

Alex Rasla

University of California, Santa Barbara
Santa Barbara, CA, USA
alexrasla@ucsb.edu

Michael Beyeler

University of California, Santa Barbara
Santa Barbara, CA, USA
mbeyeler@ucsb.edu

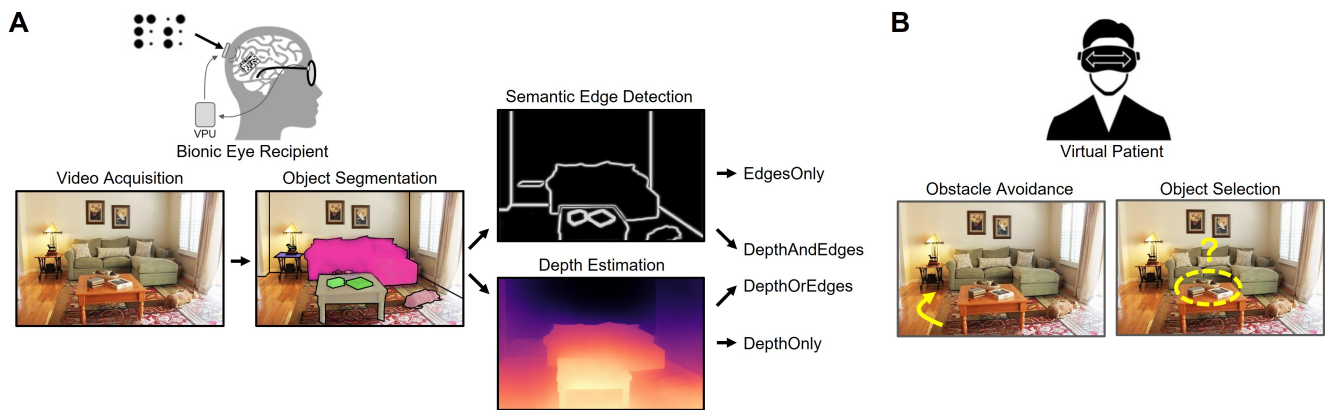


Figure 1: Scene simplification for bionic vision. A) Visual neuroprostheses (*bionic eyes*) electrically stimulate neurons in the visual system to restore a rudimentary form of vision to people living with incurable blindness (*inset*). To create meaningful artificial vision, the visual scene is simplified by extracting semantic edges and estimating relative depth before it is displayed, here illustrated on an indoor scene from the MS-COCO database. Semantic edges and depth cues may be visualized either independently (*EdgesOnly* and *DepthOnly* mode) or together (*EdgesAndDepth*). Alternatively, users may prefer the ability to flexibly switch between edges and depth cues (*EdgesOrDepth*). B) As a proof of concept, we used a neurobiologically inspired computational model to generate realistic predictions of simulated prosthetic vision, and asked sighted subjects (i.e., *virtual patients*) to avoid obstacles and select objects in an immersive virtual reality environment.

ABSTRACT

Visual neuroprostheses (*bionic eyes*) have the potential to treat degenerative eye diseases that often result in low vision or complete blindness. These devices rely on an external camera to capture the visual scene, which is then translated frame-by-frame into an electrical stimulation pattern that is sent to the implant in the eye. To highlight more meaningful information in the scene, recent studies have tested the effectiveness of deep-learning based computer vision techniques, such as depth estimation to highlight nearby obstacles (*DepthOnly* mode) and semantic edge detection to outline important objects in the scene (*EdgesOnly* mode). However, nobody has attempted to combine the two, either by presenting

them together (*EdgesAndDepth*) or by giving the user the ability to flexibly switch between them (*EdgesOrDepth*). Here, we used a neurobiologically inspired model of simulated prosthetic vision (SPV) in an immersive virtual reality (VR) environment to test the relative importance of semantic edges and relative depth cues to support the ability to avoid obstacles and identify objects. We found that participants were significantly better at avoiding obstacles using depth-based cues as opposed to relying on edge information alone, and that roughly half the participants preferred the flexibility to switch between modes (*EdgesOrDepth*). This study highlights the relative importance of depth cues for SPV mobility and is an important first step towards a visual neuroprosthesis that uses computer vision to improve a user’s scene understanding.

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

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXX.XXXX>

CCS CONCEPTS

• **Human-centered computing** → **Accessibility technologies**;
Virtual reality; *Empirical studies in visualization*.

KEYWORDS

bionic vision, virtual reality, simulated prosthetic vision, indoor mobility, scene simplification

ACM Reference Format:

Alex Rasla and Michael Beyeler. 2022. The Relative Importance of Depth Cues and Semantic Edges for Indoor Mobility Using Simulated Prosthetic Vision in Immersive Virtual Reality. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXX.XXXX>

1 INTRODUCTION

By the year 2050, roughly 114.6 million people will be living with incurable blindness [8]. Although some individuals can be treated with surgery or medication, there are no effective treatments for many people blinded by severe degeneration or damage to the retina, the optic nerve, or cortex. In such cases, an electronic visual prosthesis (*bionic eye*) may be the only option [16] (Fig. 1A). Analogous to cochlear implants, these devices electrically stimulate neurons in the early visual system to elicit neuronal responses that the brain interprets as visual percepts (*phosphenes*).

Current devices generally provide users with an improved ability to localize high-contrast objects and perform basic orientation & mobility tasks [2], but are not yet able to match the acuity of natural vision. Most current prostheses provide a very limited field of view (FOV); for example, the artificial vision generated by Argus II [26], the most widely adopted retinal implant thus far, is restricted to roughly 10×20 degrees of visual angle. This forces users to scan the environment with strategic head movements while attempting to piece together the information [14]. In addition, the limited number of electrodes (60 in Argus II) severely limits the number of independent phosphenes that the device can generate [7].

Consequently, researchers have suggested ways to simplify the visual scene before it is displayed using image processing and computer vision. One popular approach is to estimate relative depth in the scene [17, 27] and then make phosphenes appear brighter the closer they are to the observer, in order to highlight nearby obstacles. Here we refer to this method of substituting depth for intensity as *DepthOnly* mode (Fig. 1A). Another line of research suggests to extract semantic and structural edges instead [32, 33] (*EdgesOnly* mode), in order to give the user a sense of where important objects are in the scene.

However, nobody has studied how to best combine these two seemingly complementary sources of visual information. A straightforward approach would be to visualize both edge and depth cues at the same time (*EdgesAndDepth*). However, we hypothesized that users might instead prefer the ability to flexibly switch between the two (*EdgesOrDepth*). To assess the relative importance of depth and edge information, a side-by-side comparison is needed.

Due to the unique requirements of working with bionic eye recipients (e.g., constant assistance, increased setup time, travel cost), experimentation with different encoding methods remains challenging and expensive. Instead, embedding models of simulated prosthetic vision (SPV) in immersive virtual reality (VR) allows sighted subjects to act as *virtual patients* by “seeing” through the eyes of the patient, taking into account their head and eye movements as they explore an immersive virtual environment [21, 31, 35, 38]. This can speed up the development process by allowing researchers to test theoretical predictions in high-throughput experiments, the best of which can be validated and improved upon in an iterative process with the bionic eye recipient in the loop [22].

To this end, we make the following contributions:

- i. We embed a psychophysically validated SPV model [6] in VR to allow sighted participants to act as virtual bionic eye patients in an immersive virtual environment.
- ii. We explore the relative importance of different scene simplification strategies based on depth estimation [17] and semantic edge detection [33] as a preprocessing strategy for bionic vision (Fig. 1A).
- iii. We systematically evaluate the ability of these strategies to support obstacle avoidance and object identification (Fig. 1B) with a user study in immersive VR.

In sum, this is the first study to compare the relative importance of depth cues and semantic edges for bionic vision and an essential first towards a bionic eye that uses computer vision to improve a user’s scene understanding.

2 BACKGROUND

Retinal implants are currently the only FDA-approved technology to treat blinding degenerative diseases such as retinitis pigmentosa (RP) and age-related macular degeneration (ARMD). Most current devices acquire visual input via an external camera and perform edge extraction or contrast enhancement via an external video processing unit (VPU), before sending the signal through wireless coils to a microstimulator implanted in the eye or the brain (see Fig. 1A). This device receives the information, decodes it, and stimulates the visual system with electrical current. Two devices are already approved for commercial use: Argus II (60 electrodes arranged in a 6×10 grid, Second Sight Medical Products, Inc., [26]) and Alpha-IMS (1500 electrodes, Retina Implant AG, [34]). In addition, PRIMA (378 electrodes, Pixium Vision, has started clinical trials, with others to follow shortly [3, 15].

A common misconception is that each electrode in the grid can be thought of as a “pixel” in an image [11, 12, 25, 29, 32], and most retinal implants linearly translate the grayscale value of a pixel in each video frame to a current amplitude of the corresponding electrode in the array [26]. This is known as the scoreboard model, which implies that creating a complex visual scene can be accomplished simply by using the right combination of pixels, analogous to the images projected on the light bulb arrays of some sports stadium scoreboards [13]. On the contrary, recent work suggests that phosphenes vary in shape and size, and differ substantially across subjects and electrodes [6, 14].

Despite their potential to restore vision to people living with incurable blindness, the number of bionic eye users in the world is still relatively small (~ 500 retinal prostheses implanted to date). To investigate functional recovery and experiment with different implant designs, researchers have therefore been developing VR prototypes that rely on SPV. The classical method relies on sighted subjects wearing a VR headset, who are then deprived of natural viewing and only perceive phosphenes displayed in a head-mounted display (HMD). This viewing mode has been termed *transformative reality* [25] (as opposed to *altered reality* typically used to describe low vision simulations [4]), which allows sighted users to “see” through the eyes of the bionic eye recipient, taking into account their head and/or eye movements as they explore a virtual environment [22].

However, because most SPV studies rely on the scoreboard model [11, 12, 25, 29, 32], it is unclear how their findings would translate to real bionic eye recipients. Only a handful of studies have incorporated a great amount of neurophysiological detail into their setup [20, 21, 35, 37, 38], only three of which relied on an established and psychophysically validated model of SPV [21, 35, 38]. In addition, being able to move around as one would in real life has shown to significantly increase the level of immersion a user experiences [28]. However, the level of immersion offered by most SPV studies is relatively low, as stimuli are often presented on a screen [38, 39]. In contrast, most current prostheses provide a very limited FOV (e.g., Argus II: 10×20 degrees of visual angle), which requires users to scan the environment with strategic head movements while trying to piece together the information [14]. Furthermore, Argus II does not take into account the eye movements of the user when updating the visual scene, which can be disorienting for the user. Ignoring these human-computer interaction (HCI) aspects of bionic vision can result in unrealistic predictions of prosthetic performance, sometimes even exceeding theoretical limits for visual acuity (as pointed out by [10]).

3 RELATED WORK

Most retinal implants are equipped with an external VPU that is capable of applying simple image processing techniques to the video feed in real time. In the near future, these techniques may include deep learning-based algorithms aimed at improving a patient’s scene understanding [5].

Based on this premise, researchers have developed various image optimization strategies, and assessed their performance by having sighted observers (i.e., *virtual patients*) conduct daily visual tasks under SPV [1, 9, 12, 23, 27, 36]. These simulations allow a wide range of computer vision systems to be developed and tested without requiring implanted devices.

Current retinal prostheses are implanted in only one eye, and thus are unable to convey binocular depth cues. Previous work has therefore explored the possibility of obtaining depth information through additional peripherals, such as an RGB-D sensor, and studied behavioral performance of virtual patients typically navigating an obstacle course under SPV. For example, Ref. [29] used depth cues to generate a simplified representation of the ground to indicate the free space within which virtual patients could safely walk around, whereas Ref. [17] used deep neural networks to estimate per-pixel relative depth and then substituted depth for intensity. Depth cues were also shown to help avoid nearby obstacles that are notoriously hard to detect with other computer vision algorithms, such as branches hanging from a tree [24]. Ref. [27] used depth to increase the contrast of object boundaries and showed that this method reduced the number of collisions with ground obstacles. In addition, retinal prosthesis patients were shown to benefit from distance information provided by a thermal sensor when trying to avoid nearby obstacles and people [30].

Recently, with the development of deep learning in computer vision, semantic segmentation algorithms have become unprecedently effective. This can be used as another method to reduce visual clutter in prosthetic vision, be it in outdoor scenes [19] or

indoor scenes [32]. The latter study combined semantic and structural image segmentation to build a schematic representation of indoor environments, which was then shown to improve object and room identification in a SPV task [32]. Semantic segmentation has been applied to simplify both the outdoor scenes [19] and the indoor scenes [32] for retinal prostheses.

However, since the above algorithms were developed in isolation and tested on different behavioral tasks, a side-by-side comparison of their ability to aid scene understanding is still lacking. A notable exception is Ref. [17], which compared a number of scene simplification strategies that rely on depth extraction, semantic segmentation, and visual saliency on a single task. However, their study was limited to sighted participants viewing SPV videos on a screen, and thus did not account for the FOV restrictions that are common with current implants.

To address these challenges, we used a neurobiologically inspired computational model of bionic vision [6] to generate realistic predictions of SPV, and combined it with scene simplification strategies based on depth estimation [17] and semantic edge detection [33]. To allow for a fair comparison between algorithms, we asked virtual patients to avoid obstacles and identify objects in a number of immersive virtual environments.

4 METHODS

4.1 Virtual Patients

To simulate a bionic eye patient, we developed SPV simulations in Unity that were streamed in real time to a wireless head-mounted VR headset (HTC VIVE Pro Eye with wireless adapter, HTC Corporation). We followed the procedure outlined in Fig. 1A to simulate different scene simplification strategies and tested them on an obstacle avoidance and object selection task (Fig. 1B). All our simulations were run on an Intel i9-9900k processor (C# code) and an NVIDIA RTX 2070 Super GPU with 16GB of DDR4 memory (shader code). The entire SPV workflow was thus as follows:

- i. Image acquisition: Utilize Unity’s virtual camera to acquire the scene at roughly 90 frames per second and downscale to a target texture of 86×86 pixels.
- ii. Scene simplification: Extract semantic edges and/or depth cues to simulate different scene simplification strategies (see Section 4.5).
- iii. Electrode activation: Determine electrode activation based on the visual input as well as the placement of the simulated retinal implant. In the current study, a 3×3 Gaussian blur was applied to the preprocessed image to average the grayscale values around each electrode’s location in the visual field. This gray level was then interpreted as a current amplitude delivered to a particular electrode in the array.
- iv. Phosphene model: Use Unity shaders to convert electrode activation to a visual scene in real time (see Section 4.4).
- v. Phosphene rendering: Render the elicited phosphenes in the HMD of the VR system.

4.2 Participants

We recruited 18 participants with normal or corrected-to-normal vision from the Anonymous Participant Pool at Anonymous Institution to act as virtual patients in our experiment. Participants ranged




Figure 2: Room layouts. Participants started in the center along the bottom wall. Participants were instructed to walk towards the other end of the room while avoiding obstacles. At the end of the room, there was either one table with three objects on it, or three tables with one object on each. Participants had to identify the medium-sized cube located one on of the tables.

from 18 – 20 in age, with 6 participants identifying as male and 12 identifying as female. Of these participants, 4 had never used VR before, 12 had used VR 1 – 5 times before, 1 had used it 10 – 20 times, and 1 had used it 20+ times. Potential participants were excluded if they reported that they were prone to cybersickness. The study was approved by Anonymous Institution’s Institutional Review Board.

4.3 Rooms

Participants were asked to navigate six different virtual rooms that were filled with 3-7 obstacles to avoid and a target object to select among distractor objects (Fig. 2). Participants always started along the bottom wall and were instructed to walk towards the other end of the room while avoiding obstacles. Upon collision with an object, a “thud” sound was played through the VIVE headphones.

Once they had passed all obstacles, a “chime” sound was played through the VIVE headphones to indicate the end of the obstacle avoidance portion of the task. Participants then had to navigate to one of three tables and identify a medium-sized cube located on it (Fig. 3). Each room had either one table with three objects on it, or three tables with one object each. While the arrangements of objects was pseudo-randomized on each trial (consisting either of a sphere, a cylinder, and a medium cube; or a small cube, a medium cube, and a large cube), participants always had to select the medium cube.

4.4 Simulated Prosthetic Vision

To generate realistic simulations of phosphene appearance, we adapted code from the VR-SPV open-source toolbox [21]. This toolbox provides a C#/shader implementation of the psychophysically validated axon map model [6] that describes phosphenes using

two shape parameters, ρ (specifying phosphene size) and λ (describing phosphene elongation). We chose $\rho = 300$ and $\lambda = 550$ for our simulations, which is roughly in the middle of the range of values reported by Argus II users [6]. To support the real-time simulation of a 20×15 electrode array implanted in the retina, we pre-calculated an initial mapping of each electrode’s effects on the scene before starting the experiment.

4.5 Modes

We considered four different scene simplification modes (Fig. 4):

- i. *EdgesOnly*: To simulate semantic and structural edge segmentation [33], we drew white outlines around the ground-truth object boundaries of the obstacles, tables, and objects in the room (Fig. 4, *top*). In addition, we highlighted the edges of the room to give participants an indication of the room layout.
- ii. *DepthOnly*: To simulate depth to intensity substitution [17], we estimated the ground-truth depth for each pixel on the Unity screen using raycasting. Inverse depth was then linearly mapped to phosphene brightness, so that close objects appeared the brightest, and the far away wall was black (Fig. 4, *middle*).
- iii. *EdgesAndDepth*: This mode visualized both edge and depth information at the same time (Fig. 4, *bottom*).
- iv. *EdgesOrDepth*: In this mode, participants randomly started in either *EdgesOnly* or *DepthOnly* mode and were able to toggle between modes via press of a button on the VIVE controller.

4.6 Procedure

To prevent learning effects, the order in which the modes were presented and the order in which the rooms were presented was




Figure 3: Example views of the table with different objects on it. The correct object to identify was always the medium-sized cube (indicated with a white dashed circle) among two distractor objects. The arrangement of the objects was pseudo-randomized on each trial. Users had to confirm their selection by pressing a button on the VIVE controllers.

pseudo-randomized across participants. To reduce cognitive load and context-switching, all rooms for a particular mode were presented in a block. This allowed participants to get comfortable with one mode at a time. Participants saw each combination of room and mode only once.

Before the start of each mode, participants were allowed to explore a tutorial room (not part of Rooms A–F) that contained a single obstacle to avoid and a medium-sized cube on a table without distractor objects. Participants were free to navigate the room for as long as they wanted, until they felt comfortable with the SPV mode. When they indicated that they were ready to start the experiment, they had to grab the medium-sized cube off the table and confirm their selection with a button press. They were then given 30 seconds to walk back to the starting position, indicated by a red circle on the floor. At this point the experiment started with a pseudo-randomly selected room.

During the experiment, participants were given a maximum of four minutes per trial. A trial was concluded by selecting an object from one of the tables in the room. If participants exceeded the time limit, the trial ended automatically, SPV was turned off, and the participant was given 30 seconds to walk back to the starting position, before the next trial started.

In *EdgesOrDepth* mode, participants were allowed to toggle between *EdgesOnly* and *DepthOnly* modes as often as they wanted. Each trial started in the mode that was last active during the last trial (or for Trial 1: whichever was last active in the tutorial room).




Figure 4: SPV modes used for scene simplification. *Top:* In *EdgesOnly* mode, only semantic and structural edges are visualized. *Middle:* In *DepthOnly* mode, per-pixel ground-truth depth is inverted and linearly translated to grayscale level. *Bottom:* In *EdgesAndDepth* both edges and depth are visualized. A fourth mode, *EdgesOrDepth*, gave users the ability to toggle between *EdgesOnly* and *DepthOnly* modes by pressing a button on the VIVE controller.

This way, if participants preferred one mode over the other, they could just stay on that mode without having to constantly switch back and forth between trials.

4.7 Data Collection and Analysis

During the whole experiment, we recorded participant position (sampled every 0.5 seconds), time and place of obstacle collisions, and object selection. Upon collision with an obstacle, we applied a 3 second timeout window so that prolonged collisions with the same object was only counted once.

We considered the number of collisions (the fewer the better) and time taken (the shorter the better) as the main metrics to quantify obstacle avoidance performance. The obstacle avoidance part of the task ended as soon as participants entered the top third of the room. From that time forward, participants were considered to be in the object selection part of the task, even if they accidentally crossed the line back into the obstacle avoidance portion of the room.

We considered the fraction of correctly selected objects (accuracy; the higher the better) and time taken (the shorter the better) as the main metrics to quantify object selection performance. A trial ended as soon as participants selected an object. Trials where time ran out before a selection could be made were counted as incorrect object selections.

To determine whether performance was significantly different across modes and rooms, we ran a regression analysis for each metric mentioned above using the Ordinary Least Squares (OLS)




Figure 5: Obstacle avoidance (OA) performance, measured by success rate (i.e., the number of trials with zero collisions, Panel A) and time taken (Panel C), and object selection (OS) performance, measured by accuracy (i.e., the fraction of trials where the correct object was selected, Panel B), and time taken (Panel D). The dashed line in Panel B indicates chance performance (33 %). Vertical bars are the standard error of the mean (SEM). Statistical significance was determined using paired t -tests, corrected for multiple testing using the Holm-Sidak method (*: $p < .05$).

model in the statsmodel Python module as $Metric \sim C(Mode) + C(Room)$. We also considered gender and VR experience as factors affecting performance, but found no statistically significant differences and thus subsequently removed these variables from the model. p -values were calculated using paired t -tests of the OLS model, which corrects for multiple testing using the Holm-Sidak method.

5 RESULTS

5.1 Obstacle Avoidance

Fig. 5 summarizes the obstacle avoidance (OA) and object selection (OS) performance of all participants for the different scene simplification modes.

Participants performed the OA portion of the task most successfully using the *EdgesAndDepth* mode, completing on average 89.8 % of trials without colliding (“success rate”, Fig. 5A). However, this was not significantly better than with the *DepthOnly* (87.9 %) and *EdgesOrDepth* modes (87.0 %; t -test $p > 0.05$, corrected for multiple comparisons with the Holm-Sidak method). On the other hand, participants performed worst with the *EdgesOnly* mode, as evidenced by a significantly lower success rate (64.8 %) and among the longest times taken (25 s, Fig. 5C). Participants were significantly faster using *DepthOnly* (17 s) than they were with *EdgesOrDepth* (22 s) and *EdgesAndDepth* (25 s; Fig. 5C).

5.2 Object Selection

Participants performed the OS portion of the task most successfully using the *EdgesOrDepth* mode, on average locating and selecting the correct object in 54.6 % of trials (Fig. 5B). However, this was not significantly better than with the *EdgesOnly* (46.3 %) and *DepthOnly* modes (50.9 %). Performance with the *EdgesAndDepth* mode (40.7 %) was only slightly better than chance (33 %).

All participants took approximately 40 s to walk towards the correct table and select the correct object, no matter which scene simplification mode was used (Fig. 5D). Since the location of the correct object was pseudo-randomized on each trial, it is possible

that most of the time for the OS portion of the task was spent simply walking back and forth between the three tables.

5.3 Participant Paths

The paths that participants took are summarized in Fig. 6, showing a birds-eye view of the six different rooms. Here, each line represents the path of a participant, with colors getting more saturated as time progresses, and collisions indicated by black crosses.

The poor OA performance of the *EdgesOnly* mode first reported in Fig. 5 can also be appreciated here, with participants causing more collisions compared to the other modes across all rooms. This is especially evident in Room F: Even though participants attempted to avoid the obstacle maze by moving in a slalom fashion, there are only two obstacles that never led to collisions.

It is interesting to note that participants seemed to explore Room A the most, with some participants accidentally walking away from the table in the top-left corner of the room and returning to the OA portion of the task. On the other hand, most participants walked mostly straight ahead in Room B (except in *EdgesOrDepth* mode) and Room E, thus avoiding most obstacles, although most participants still collided with the obstacle that was closest to the starting position. Room F was by far the most challenging, leading to a large number of collisions in all four modes.

5.4 Relative Importance of Depth vs. Edges

One potential indicator of the relative importance of depth and edge cues was *EdgesOrDepth* mode: here participants had the ability to flexibly switch between *EdgesOnly* and *DepthOnly* information at the push of a button. Participants were free to toggle modes as often as they saw fit.

We recorded both the number of times that participants toggled modes as well as the amount of time spent in each mode. The results are shown in Fig. 7. To our surprise, participants spent most of their time in *DepthOnly* mode (Fig. 7, top), both during the OA as well as the OS portion of the task.

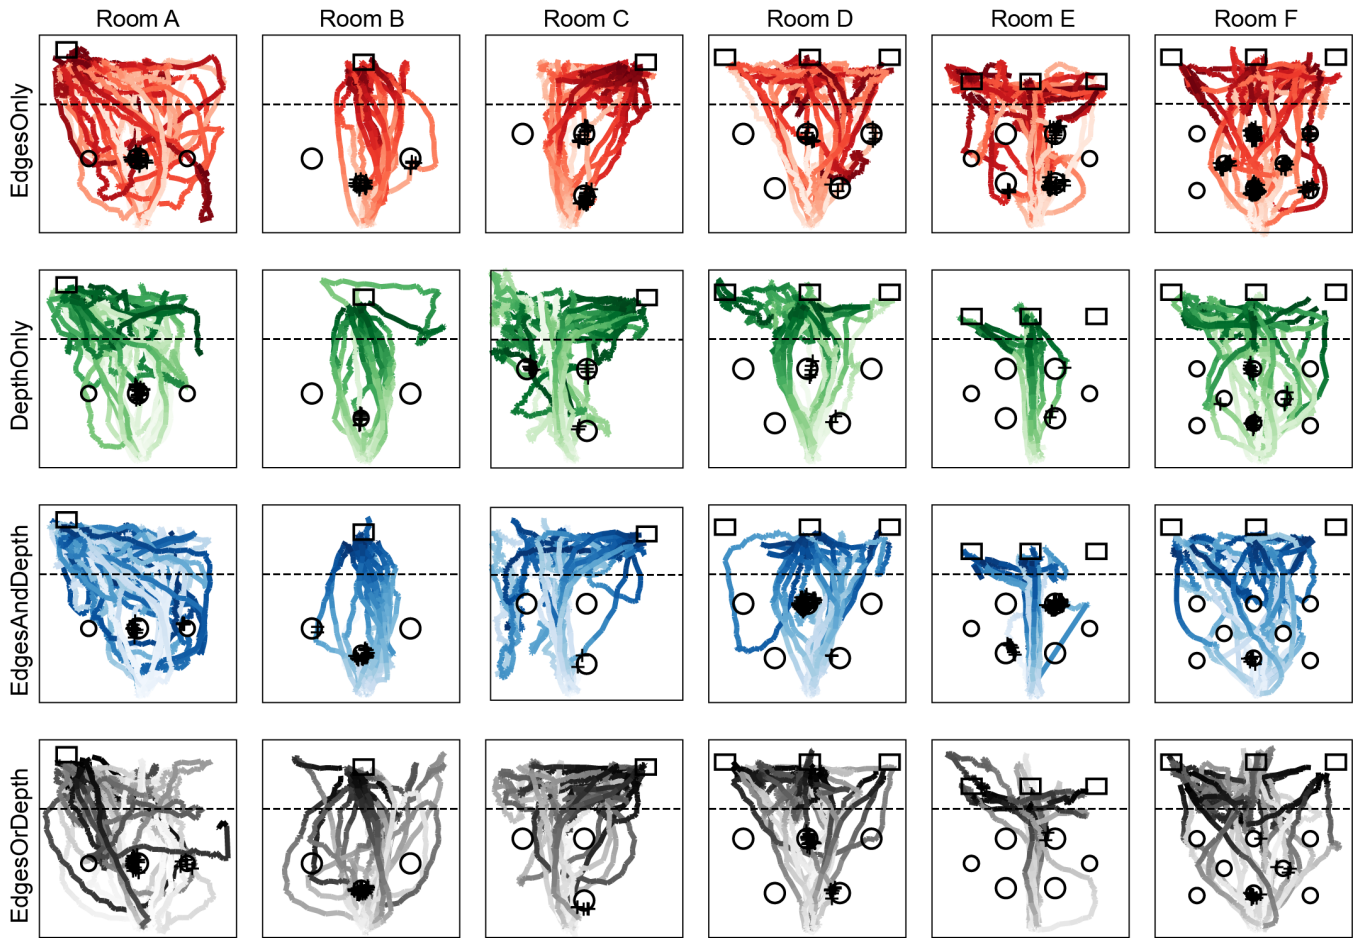


Figure 6: Birds-eye view of paths taken by all subjects in the different rooms (columns) using the different scene simplification strategies (rows). Color of paths gets more saturated as time moves on. Collisions are indicated with a black + sign. Circles indicate the location of the obstacles, and rectangles are the table. The task switched from obstacle avoidance to object selection as soon as the participant crossed the (to them invisible) horizontal dashed line.

In addition, most participants did not often switch between modes ((Fig. 7, bottom). The experiment was set up so that each trial started on the mode that was last active (either during the preceding trial or in the tutorial room), so that participants who preferred one mode over the other would not have to constantly switch back and forth. During the OA portion of the task, participants tended to toggle modes either never (perhaps because the trial started in their preferred mode), once (perhaps because they trial started in their nonpreferred mode), or twice (perhaps to briefly view the other mode, but then promptly switch back to the preferred one).

The only exception was Room A, where participants tended to switch modes much more often. Since the orders of the rooms was pseudo-randomized, it is not entirely clear why Room A would prompt participants to switch modes more often than other rooms.

5.5 User Preferences

Upon completion of the experiment, we asked participants which mode they preferred for the OA portion of the task, the OS portion of the task, and overall (Fig. 8).

Consistent with Fig. 5A, participants preferred *DepthOnly* and *EdgesOrDepth* to avoid obstacles; $\chi^2(1, N = 18) = 8.22, p < 0.05$. However, *EdgesAndDepth* did also result in strong OA performance, yet was only preferred by two out of 18 participants. Most participants preferred the ability to switch between depth and edge cues by means of the *EdgesOrDepth* mode.

When it came to selecting objects, user preferences were mixed (Fig. 8, center) and not significantly different from a uniform random distribution: $\chi^2(1, N = 18) = 0.667, p = 0.88$. This is consistent with the somewhat lackluster performance of participants in the object selection portion of the task (Fig. 5B), suggesting that none of the tested modes may be ideal for object recognition.

Lastly, overall preference mirrored the OA result, with most participants preferring *DepthOnly* and *EdgesOrDepth* mode (Fig. 8,




Figure 7: EdgesOrDepth: Fraction of time spent on each mode (top) and number of mode toggles (bottom) for each room, for the obstacle avoidance portion of the task (left) and the object selection portion of the task (right).

right). However, this result was not significantly different from a random uniform distribution, although it approached significance ($\chi^2(1, N = 18) = 7.33, p = 0.06$). Notably, none of the subjects preferred *EdgesOnly* mode overall, and only four subjects preferred *EdgesAndDepth* mode. The majority of participants was split between being comfortable using *DepthOnly* mode and preferring the flexibility to switch between modes using *EdgesOrDepth* mode (even though most time in *EdgesOrDepth* was spent looking at depth cues).

Overall these results further corroborate the relative importance of depth cues for both obstacle avoidance and object identification.

6 DISCUSSION

6.1 Relative Importance of Depth Cues and Semantic Edges for Bionic Vision

Here we asked sighted participants to act as virtual bionic eye recipients by navigating virtual indoor environments using simulated prosthetic vision. To the best of our knowledge, this is the first study that directly compares the relative importance of depth cues and semantic edges for bionic vision.

As expected, we found that participants were significantly better at avoiding obstacles using depth-based methods as opposed to relying on edge information alone (Fig. 5A). These results are consistent with previous studies highlighting the importance of highlighting nearby obstacles for orientation & mobility tasks [24, 27]. However, the exact way in which depth cues were presented seemed to carry little importance, as evidenced by similar performance of the

DepthOnly, *EdgesAndDepth*, and *EdgesOrDepth* modes. Nevertheless, it is interesting to note that participants completed the obstacle avoidance portion of the task the fastest using *DepthOnly* mode (Fig. 5C), even though *EdgesAndDepth* provided the same amount of depth cues at all times.

Theoretically, participants should have performed the same (at least asymptotically) with the *EdgesOrDepth* mode, as it provides at least the same amount of information as *DepthOnly* mode. In its original conception, *EdgesOrDepth* mode was supposed to provide all the functionality of the individual modes (i.e., *DepthOnly* and *EdgesOnly*) while giving the user the power and flexibility to choose whichever mode they deem best for a particular situation. Instead, it is possible that the ability to switch added unnecessary overhead in terms of time taken to complete the task and opened up the possibility that users may perform the task in a suboptimal mode. From an HCI perspective, this is an important insight to consider for future designs of bionic eyes that rely on computer vision for scene simplification [5].

6.2 Depth and Semantic Edges May Not Be Sufficient for Object Recognition

Depth cues were also important for object selection (Fig. 5B), with participants performing best using the *DepthOnly* and *EdgesOrDepth* modes. However, performance was more lackluster as compared to the obstacle avoidance portion of the task. Participants performed near chance levels using *EdgesAndDepth*, and were only slightly better with the other three modes. Only *DepthOnly* and *EdgesOrDepth* were significantly better than *EdgesAndDepth*, and time taken to complete this subtask was similar across modes.

Overall these results suggest that presenting both edge and depth cues at once may hinder participants' ability to segregate important objects from the background. Yet none of the tested modes seem to provide sufficient information for participants to excel at this portion of the task.

6.3 Users Prefer Depth Information Over Edges

The relative importance of depth information was further corroborated by Fig. 7, which indicates that when participants were given a choice between edges and depth (i.e., in *EdgesOrDepth* mode), they largely preferred to navigate the room by relying on depth information alone.

In obstacle avoidance and overall, most participants preferred *DepthOnly* and *EdgesOrDepth* over the other two modes (Fig. 8), but opinions were largely split: Approximately half of participants preferred the flexibility of switching between edges and depth, whereas the other half was content with *DepthOnly* mode. The lackluster object selection performance is also evidenced by the user preference chart, which was not significantly different from random.

Overall these results suggest that while depth information seems to be the most beneficial for obstacle avoidance and object selection, there may not be one scene simplification made that is preferred by the majority of bionic eye recipients.

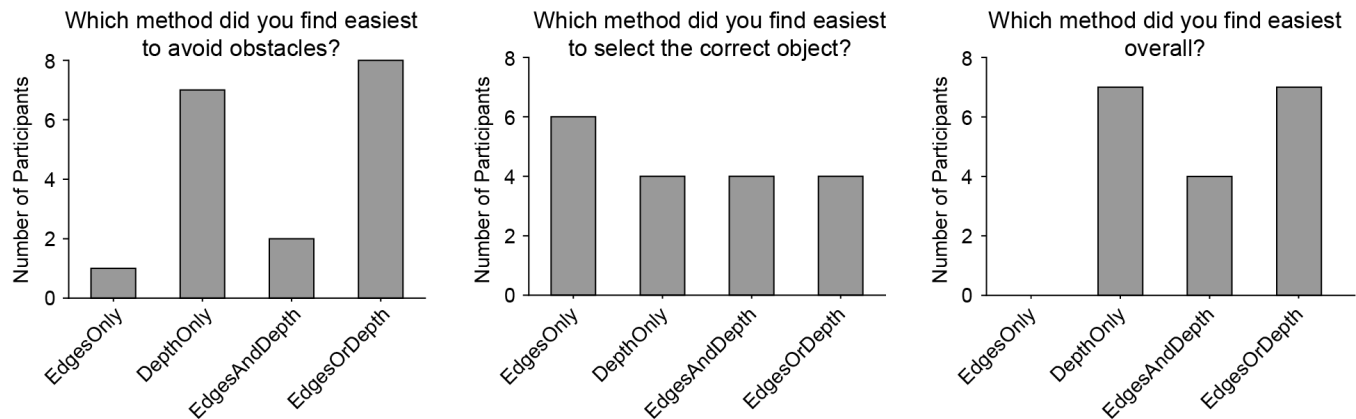


Figure 8: Reported user preferences for the different SPV modes.

6.4 Limitations and Future Work

Although the present study effectively uses immersive VR to address previously unanswered questions about SPV, there are a number of limitations that should be addressed in future work as outlined below.

First, it is important to point out that the study was performed on sighted undergraduate students readily available at Anonymous Institution. Their age, navigational affordances, and experience with low vision may therefore be drastically different from real bionic eye recipients, who not only tend to be older (and prolific cane users), but also receive extensive vision rehabilitation training.

Second, we acknowledge that the navigating with simulated prosthetic vision was a challenging task even for experienced users of VR. Sometimes, as the experiment progressed into the later stages, fatigued and frustrated participants would become increasingly reckless and charge straight ahead through the obstacle course (without any real-life consequences of a collision) or simply select the first object they encountered. Therefore, a larger participant pool might have been able to average out such behavior and reveal more pronounced differences between the tested SPV modes.

Interestingly, we found vast individual differences in task performance, which were not unlike those reported in the literature [18]. Subjects who did well with one mode tended to do well across all modes (data not shown), suggesting that some people were inherently better at adapting to prosthetic vision than others. Future work should therefore zero in on the possible causes of these individual differences and compare them to real bionic eye users. Studying these differences could identify training protocols to enhance the ability of all device users.

7 CONCLUSION

In sum, this is the first study to compare the relative importance of depth cues and semantic edges for bionic vision and constitutes an essential first towards a bionic eye that uses computer vision to improve a user's scene understanding.

REFERENCES

[1] W. I. Al-Atabany, T. Tong, and P. A. Degenaar. 2010. Improved content aware scene retargeting for retinitis pigmentosa patients. *Biomed Eng Online* 9 (Sept.

2010), 52. <https://doi.org/10.1186/1475-925X-9-52>

[2] Lauren N. Ayton, Nick Barnes, Gislin Dagnelie, Takashi Fujikado, Georges Goetz, Ralf Hornig, Bryan W. Jones, Mahiul M. K. Muqit, Daniel L. Rathbun, Katarina Stingl, James D. Weiland, and Matthew A. Petoe. 2020. An update on retinal prostheses. *Clinical Neurophysiology* 131, 6 (June 2020), 1383–1398. <https://doi.org/10.1016/j.clinph.2019.11.029>

[3] Lauren N. Ayton, Peter J. Blamey, Robyn H. Guymer, Chi D. Luu, David A. X. Nayagam, Nicholas C. Sinclair, Mohit N. Shivdasani, Jonathan Yeoh, Mark F. McCombe, Robert J. Briggs, Nicholas L. Opie, Joel Villalobos, Peter N. Dimitrov, Mary Varsamidis, Matthew A. Petoe, Chris D. McCarthy, Janine G. Walker, Nick Barnes, Anthony N. Burkitt, Chris E. Williams, Robert K. Shepherd, Penelope J. Allen, and for the Bionic Vision Australia Research Consortium. 2014. First-in-Human Trial of a Novel Suprachoroidal Retinal Prosthesis. *PLOS ONE* 9, 12 (Dec. 2014), e115239. <https://doi.org/10.1371/journal.pone.0115239> Publisher: Public Library of Science.

[4] Min Bao and Stephen A. Engel. 2019. Augmented Reality as a Tool for Studying Visual Plasticity: 2009 to 2018. *Current Directions in Psychological Science* 28, 6 (Dec. 2019), 574–580. <https://doi.org/10.1177/0963721419862290> Publisher: SAGE Publications Inc.

[5] Michael Beyeler and Melani Sanchez Garcia. 2022. Towards a Smart Bionic Eye: AI-Powered Artificial Vision for the Treatment of Incurable Blindness. <https://doi.org/10.31219/osf.io/fpdyy>

[6] Michael Beyeler, Devyani Nanduri, James D. Weiland, Ariel Rokem, Geoffrey M. Boynton, and Ione Fine. 2019. A model of ganglion axon pathways accounts for percepts elicited by retinal implants. *Scientific Reports* 9, 1 (June 2019), 1–16. <https://doi.org/10.1038/s41598-019-45416-4>

[7] Michael Beyeler, Ariel Rokem, Geoffrey M Boynton, and Ione Fine. 2017. Learning to see again: biological constraints on cortical plasticity and the implications for sight restoration technologies. *Journal of Neural Engineering* 14, 5 (Oct. 2017), 051003. <https://doi.org/10.1088/1741-2552/aa795e>

[8] Rupert R. A. Bourne, Seth R. Flaxman, Tasanee Braithwaite, Maria V. Cicinelli, Aditi Das, Jost B. Jonas, Jill Keefe, John H. Kempen, Janet Leasher, Hans Limburg, Kovin Naidoo, Konrad Pesudovs, Serge Resnikoff, Alex Silvester, Gretchen A. Stevens, Nina Tahhan, Tien Y. Wong, Hugh R. Taylor, Rupert Bourne, Peter Ackland, Aries Arditi, Yaniv Barkana, Banu Bozkurt, Tasanee Braithwaite, Alain Bron, Donald Budenz, Feng Cai, Robert Casson, Usha Chakravarthy, Jaewan Choi, Maria Vittoria Cicinelli, Nathan Congdon, Reza Dana, Rakhi Dandona, Lalit Dandona, Aditi Das, Iva Dekaris, Monte Del Monte, Jenny Deva, Laura Dreer, Leon Ellwein, Marcela Frazier, Kevin Frick, David Friedman, Joao Furtado, Hua Gao, Gus Gazzard, Ronnie George, Stephen Gichuhi, Victor Gonzalez, Billy Hammond, Mary Elizabeth Hartnett, Minguang He, James Hejtmancik, Flavio Hirai, John Huang, April Ingram, Jonathan Javitt, Jost Jonas, Charlotte Joslin, Jill Keefe, John Kempen, Moncef Khairallah, Rohit Khanna, Judy Kim, George Lambrou, Van Charles Lansingh, Paolo Lanzetta, Janet Leasher, Jennifer Lim, Hans Limburg, Kaweh Mansouri, Anu Mathew, Alan Morse, Beatriz Munoz, David Musch, Kovin Naidoo, Vinay Nangia, Maria Palaiou, Maurizio Battaglia Parodi, Fernando Yaacov Pena, Konrad Pesudovs, Tunde Peto, Harry Quigley, Murugesan Raju, Pradeep Ramulu, Serge Resnikoff, Alan Robin, Luca Rossetti, Jinan Saaddine, Mya Sandar, Janet Serle, Tueng Shen, Rajesh Shetty, Pamela Sieving, Juan Carlos Silva, Alex Silvester, Rita S. Sitorus, Dwight Stambolian, Gretchen Stevens, Hugh Taylor, Jaime Tejedor, James Tielsch, Miltiadis Tsilimbaris, Jan van Meurs, Rohit Varma, Gianni Virgili, Jimmy Volmink, Ya Xing Wang, Ning-Li Wang, Sheila West, Peter Wiedemann, Tien Wong, Richard Wormald, and Yingfeng Zheng.

2017. Magnitude, temporal trends, and projections of the global prevalence of blindness and distance and near vision impairment: a systematic review and meta-analysis. *The Lancet Global Health* 5, 9 (Sept. 2017), e888–e897. [https://doi.org/10.1016/S2214-109X\(17\)30293-0](https://doi.org/10.1016/S2214-109X(17)30293-0) Publisher: Elsevier.
- [9] Justin R. Boyle, Anthony J. Maeder, and Wageeh W. Boles. 2008. Region-of-interest processing for electronic visual prostheses. *Journal of Electronic Imaging* 17, 1 (Jan. 2008), 013002. <https://doi.org/10.1117/1.2841708> Publisher: International Society for Optics and Photonics.
- [10] Avi Caspi and Ari Z. Zivotofsky. 2015. Assessing the utility of visual acuity measures in visual prostheses. *Vision Research* 108 (March 2015), 77–84. <https://doi.org/10.1016/j.visres.2015.01.006>
- [11] S. C. Chen, G. J. Suaning, J. W. Morley, and N. H. Lovell. 2009. Simulating prosthetic vision: I. Visual models of phosphenes. *Vision Research* 49, 12 (June 2009), 1493–506.
- [12] Gislin Dagnelie, Pearse Keane, Venkata Narla, Liancheng Yang, James Weiland, and Mark Humayun. 2007. Real and virtual mobility performance in simulated prosthetic vision. *Journal of Neural Engineering* 4, 1 (March 2007), S92–S101. <https://doi.org/10.1088/1741-2560/4/1/S11>
- [13] Wm H. Dobbelle. 2000. Artificial Vision for the Blind by Connecting a Television Camera to the Visual Cortex. *ASAIJ Journal* 46, 1 (Feb. 2000), 3–9. https://journals.lww.com/asaiojournal/Fulltext/2000/01000/Artificial_Vision_for_the_Blind_by_Connecting_a_2.aspx
- [14] Cordelia Erickson-Davis and Helma Korzybska. 2021. What do blind people “see” with retinal prostheses? Observations and qualitative reports of epiretinal implant users. *PLOS ONE* 16, 2 (Feb. 2021), e0229189. <https://doi.org/10.1371/journal.pone.0229189> Publisher: Public Library of Science.
- [15] Laura Ferlauto, Marta Jole Ildelfonsa Airaghi Leccardi, Naig Aurelia Ludmilla Chenais, Samuel Charles Antoine Gilliéron, Paola Vagni, Michele Bevilacqua, Thomas J. Wolfensberger, Kevin Sivula, and Diego Ghezzi. 2018. Design and validation of a foldable and photovoltaic wide-field epiretinal prosthesis. *Nature Communications* 9, 1 (March 2018), 1–15. <https://doi.org/10.1038/s41467-018-03386-7>
- [16] Eduardo Fernandez. 2018. Development of visual Neuroprostheses: trends and challenges. *Bioelectronic Medicine* 4, 1 (Aug. 2018), 12. <https://doi.org/10.1186/s42234-018-0013-8>
- [17] Nicole Han, Sudhanshu Srivastava, Aiwen Xu, Devi Klein, and Michael Beyeler. 2021. Deep Learning–Based Scene Simplification for Bionic Vision. In *Augmented Humans Conference 2021 (AHs'21)*. Association for Computing Machinery, New York, NY, USA, 45–54. <https://doi.org/10.1145/3458709.3458982>
- [18] Yingchen He, Susan Y. Sun, Arup Roy, Avi Caspi, and Sandra R. Montezuma. 2020. Improved mobility performance with an artificial vision therapy system using a thermal sensor. *Journal of Neural Engineering* 17, 4 (Aug. 2020), 045011. <https://doi.org/10.1088/1741-2552/aba4fb> Publisher: IOP Publishing.
- [19] Lachlan Horne, Jose Alvarez, Chris McCarthy, Mathieu Salzmann, and Nick Barnes. 2016. Semantic labeling for prosthetic vision. *Computer Vision and Image Understanding* 149 (Aug. 2016), 113–125. <https://doi.org/10.1016/j.cviu.2016.02.015>
- [20] Horace Josh, Collette Mann, Lindsay Kleeman, and Wen Lik Dennis Lui. 2013. Psychophysics testing of bionic vision image processing algorithms using an FPGA Hatpack. In *2013 IEEE International Conference on Image Processing*. 1550–1554. <https://doi.org/10.1109/ICIP.2013.6738319> ISSN: 2381-8549.
- [21] Justin Kasowski and Michael Beyeler. 2022. Immersive Virtual Reality Simulations of Bionic Vision. In *Augmented Humans 2022 (AHs 2022)*. Association for Computing Machinery, New York, NY, USA, 82–93. <https://doi.org/10.1145/3519391.3522752>
- [22] Justin Kasowski, Nathan Wu, and Michael Beyeler. 2021. Towards Immersive Virtual Reality Simulations of Bionic Vision. In *Augmented Humans Conference 2021*. ACM, Rovaniemi Finland, 313–315. <https://doi.org/10.1145/3458709.3459003>
- [23] Heng Li, Xiaofan Su, Jing Wang, Han Kan, Tingting Han, Yajie Zeng, and Xinyu Chai. 2018. Image processing strategies based on saliency segmentation for object recognition under simulated prosthetic vision. *Artificial Intelligence in Medicine* 84 (Jan. 2018), 64–78. <https://doi.org/10.1016/j.artmed.2017.11.001>
- [24] P. Lieby, N. Barnes, C. McCarthy, Nianjun Liu, H. Dennett, J. G. Walker, V. Botea, and A. F. Scott. 2011. Substituting depth for intensity and real-time phosphene rendering: Visual navigation under low vision conditions. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, Boston, MA, 8017–8020. <https://doi.org/10.1109/IEMBS.2011.6091977>
- [25] Wen Lik Dennis Lui, Damien Browne, Lindsay Kleeman, Tom Drummond, and Wai Ho Li. 2011. Transformative reality: Augmented reality for visual prostheses. In *2011 10th IEEE International Symposium on Mixed and Augmented Reality*. 253–254. <https://doi.org/10.1109/ISMAR.2011.6092402> ISSN: null.
- [26] Y. H. Luo and L. da Cruz. 2016. The Argus(R) II Retinal Prosthesis System. *Prog Retin Eye Res* 50 (Jan. 2016), 89–107. <https://doi.org/10.1016/j.preteyeres.2015.09.003>
- [27] Chris McCarthy, Janine G. Walker, Paulette Lieby, Adele Scott, and Nick Barnes. 2014. Mobility and low contrast trip hazard avoidance using augmented depth. *Journal of Neural Engineering* 12, 1 (Nov. 2014), 016003. <https://doi.org/10.1088/1741-2560/12/1/016003> Publisher: IOP Publishing.
- [28] Marco Pasch, Nadia Bianchi-Berthouze, Betsy van Dijk, and Anton Nijholt. 2009. Immersion in Movement-Based Interaction. In *Intelligent Technologies for Immersive Entertainment (Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering)*, Anton Nijholt, Dennis Reidsma, and Hendri Hondorp (Eds.). Springer, Berlin, Heidelberg, 169–180. https://doi.org/10.1007/978-3-642-02315-6_16
- [29] Alejandro Perez-Yus, Jesus Bermudez-Cameo, Jose J. Guerrero, and Gonzalo Lopez-Nicolas. 2017. Depth and Motion Cues with Phosphene Patterns for Prosthetic Vision. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. 1516–1525. <https://doi.org/10.1109/ICCVW.2017.179> ISSN: 2473-9944.
- [30] Roksana Sadeghi, Arathy Kartha, Michael P. Barry, Paul Gibson, Avi Caspi, Arup Roy, and Gislin Dagnelie. 2019. Thermal and Distance image filtering improve independent mobility in Argus II retinal implant. *Journal of Vision* 19, 15 (Dec. 2019), 23–23. <https://doi.org/10.1167/19.15.23> Publisher: The Association for Research in Vision and Ophthalmology.
- [31] Melani Sanchez-Garcia, Ruben Martinez-Cantin, Jesus Bermudez-Cameo, and Jose J. Guerrero. 2020. Influence of field of view in visual prostheses design: Analysis with a VR system. *Journal of Neural Engineering* 17, 5 (Oct. 2020), 056002. <https://doi.org/10.1088/1741-2552/abb9be> Publisher: IOP Publishing.
- [32] Melani Sanchez-Garcia, Ruben Martinez-Cantin, and Josechu J. Guerrero. 2019. Indoor Scenes Understanding for Visual Prosthesis with Fully Convolutional Networks. In *VISGRAPP*. <https://doi.org/10.5220/0007257602180225>
- [33] Melani Sanchez-Garcia, Ruben Martinez-Cantin, and Jose J. Guerrero. 2020. Semantic and structural image segmentation for prosthetic vision. *PLOS ONE* 15, 1 (Jan. 2020), e0227677. <https://doi.org/10.1371/journal.pone.0227677>
- [34] K. Stingl, K. U. Bartz-Schmidt, D. Besch, A. Braun, A. Bruckmann, F. Gekeler, U. Greppmaier, S. Hipp, G. Hortdorfer, C. Kernstock, A. Koitschev, A. Kusnyerik, H. Sachs, A. Schatz, K. T. Stingl, T. Peters, B. Wilhelm, and E. Zrenner. 2013. Artificial vision with wirelessly powered subretinal electronic implant alpha-IMS. *Proc Biol Sci* 280, 1757 (April 2013), 20130077. <https://doi.org/10.1098/rspb.2013.0077>
- [35] Jacob Thomas Thorn, Enrico Migliorini, and Diego Ghezzi. 2020. Virtual reality simulation of epiretinal stimulation highlights the relevance of the visual angle in prosthetic vision. *Journal of Neural Engineering* 17, 5 (Nov. 2020), 056019. <https://doi.org/10.1088/1741-2552/abb5bc> Publisher: IOP Publishing.
- [36] Victor Vergnieux, Marc J.-M. Macé, and Christophe Jouffrais. 2017. Simplification of Visual Rendering in Simulated Prosthetic Vision Facilitates Navigation. *Artificial Organs* 41, 9 (Sept. 2017), 852–861. <https://doi.org/10.1111/aor.12868> Publisher: John Wiley & Sons, Ltd.
- [37] Milena Vurro, Anne Marie Crowell, and John S. Pizaris. 2014. Simulation of thalamic prosthetic vision: reading accuracy, speed, and acuity in sighted humans. *Frontiers in Human Neuroscience* 8 (2014), 816. <https://doi.org/10.3389/fnhum.2014.00816>
- [38] Lihui Wang, Fariba Sharifan, Jonathan Napp, Carola Nath, and Stefan Pollmann. 2018. Cross-task perceptual learning of object recognition in simulated retinal implant perception. *Journal of Vision* 18, 13 (Dec. 2018), 22. <https://doi.org/10.1167/18.13.22>
- [39] Zhao Ying, Geng Xiulin, Li Qi, and Jiang Guangqi. 2018. Recognition of virtual maze scene under simulated prosthetic vision. In *2018 Tenth International Conference on Advanced Computational Intelligence (ICACI)*. 1–5. <https://doi.org/10.1109/ICACI.2018.8377543>