SHIFTING GAZES WITH VISUAL PROSTHESES: LONG-TERM HAND-CAMERA COORDINATION

by

Michael Patrick Barry

A dissertation submitted to Johns Hopkins University in conformity with the requirements for the degree of Doctor of Philosophy.

Baltimore, Maryland

March 2018

Purpose: Prosthetic vision is young, and many aspects of its use remain unexplored. Hand-camera coordination, the prosthetic correlate of hand-eye coordination, relies heavily on how the camera is aligned with the eye. It is unknown whether users of prostheses can adapt to using misaligned cameras, or whether requirements for proper alignment remain constant over time.

Methods: Four blind subjects implanted with Argus II retinal prostheses participated in this study. Each subject attempted to touch a single 4°–7° white target that was randomly located on an otherwise black touchscreen in a target localization task. Touch response accuracy was used to determine the necessary adjustment to eye-camera alignment, the optimal camera alignment position (OCAP). Subjects attended over 100 sessions across up to 5.3 years.

S1–S3 were given misaligned cameras for over 1 year. Adaptation was measured through changes in localization errors. Outside that period of intentional misalignment, cameras were aligned to maximize localization accuracy. During the final year, localization tasks were performed in alternation with eye tracking. S2–S4 also participated in 1-day experiments with simultaneous eye tracking and target localization.

Results: Subjects were not able to significantly reduce localization error when cameras were misaligned. When trying to maximize localization accuracy, necessary OCAPs changed significantly over time. OCAP trend directions within days and trial runs matched changes between the beginnings of days and runs. Changes between the end of a day or run and the beginning of the next tended to point in the opposite direction of the previous trend, indicating a reset of OCAP changes.

Changes in eye orientations correlated significantly with changes in OCAPs. Eye-orientation trends displayed the same reset behavior between days and runs as OCAPs. Simultaneous eye tracking and localization showed agreement between eye-orientation and localization-error trend directions. Adjusting camera alignment with eye-tracking data slowed changes in localization errors.

Conclusions: Users of current visual prostheses cannot passively adapt to camera misalignments. OCAPs are not constant with time. Prosthesis users who desire maximum pointing accuracy will require regular camera realignments. Camera alignments based on eye tracking can reduce both transient and long-term changes in localization that are related to eye movements.

Readers: Gislin Dagnelie, Robert W. Massof, Xiaoqin Wang

Acknowledgements

While many people directly and indirectly contributed to the work described in this dissertation, none were more essential than the participants who donated their time and effort to the underlying experiments. These participants risked themselves, both physically and emotionally, in agreeing to help the field of visual prosthetics by having retinal prostheses surgically implanted. The prosthesis recipients did not regain anything like normal vision by undergoing the procedure, but generously hoped to provide knowledge and experience necessary to develop better visual prostheses. Beyond the already lengthy demands of participating in the feasibility study for the Argus II retinal prosthesis, these participants volunteered for additional studies that explored finer aspects of prosthetic vision. S1, S2, and S3 each devoted over 100 hours just to this dissertation's research, and S3 went so far as to dedicate over 400 hours. Their work was exhausting and repetitive, and I cannot thank them enough for their generosity and reliability.

Gislin Dagnelie, Ph.D. has been an outstanding mentor throughout this entire process. He introduced me to the realm of visual prostheses as an undergraduate student, and has guided me unwaveringly through the completion of my graduate education. Gislin has been a source of insight for not only the science behind vision restoration and low-vision rehabilitation, but also finer matters of culture and the simplicities of optimism and kindness.

Robert Massof, Ph.D. and Xiaoqin Wang, Ph.D. kindly agreed to serve on my thesis committee and provided valuable feedback that improved both the substance and presentation of this dissertation. I did not update them on research progress as often as may have been optimal for debating experiment minutiae, but they graciously still took the time to understand everything

and point out weak spots in my work. This dissertation's topic did not directly relate to either of their primary research focuses, but they still found time in their busy schedules to help me improve and complete my project.

While working under Dr. Dagnelie, many people affiliated with the lab provided guidance and support through numerous projects. Veronika Mueller was my first direct mentor, who introduced me to prosthetic vision simulations as part of her graduate studies. H. Christiaan Stronks, Ph.D. shared much of the implantee-testing experience, and offered additional perspectives on working through Ph.D. training. Olukemi Adeyemo held everything together by coordinating both subjects and experimenters, making sure everyone was where they needed to be and that they had all the necessary resources. Liancheng Yang kept all of the custom equipment operational and wrote the programs used for the tower-mounted eye-tracking tests.

Low-vision rehabilitation specialist Duane Geruschat, Ph.D. was heavily involved with other projects that examined the abilities of Argus II users, and was always very inclusive and supportive during my years of training. Dr. Geruschat helped me appreciate that the artificial vision experienced by prosthesis users is not as good as we tend to hope or desire. Pamela Jeter, Ph.D., while our research areas had minimal overlap, always brought a nice atmosphere to lab meetings and helped me to recognize the place of prosthetic vision in the wider realm of low vision. I also had the pleasure of working with many trainees and students who assisted the lab over shorter time periods, including Lauren Dalvin, M.D., Jordan Matelsky, Sophia Diaz-Aguilar, Ali Abdulkarim, Akudo Umeh, Reem El-Husseiny, Alfred Vinnett, and Neeraj Ochaney. These students each brought a unique perspective as they assisted me with experiments, and they certainly improved the data collection process. Roksana Sadeghi has provided similar assistance

in recent months, and I wish her all the best as she continues her Ph.D. training under Dr. Dagnelie.

Second Sight Medical Products, Inc., beyond providing the technological basis for this dissertation's research, supplied essential training, materials, and experimental support. I learned much about how to test and configure the Argus II from Second Sight's former and current employees, including Arup Roy, Ph.D., Jessy Dorn, Ph.D., Jianing Wei, Ph.D., Avi Caspi, Ph.D., Varalakshmi Wuyyuru, Uday Patel, Ph.D., Ashish Ahuja, Ph.D., and Punita Christopher, Ph.D. The original target localization program, used as the basis for most measurements in this dissertation, was developed by Second Sight for the Argus II Feasibility Study. Jordan Matelsky customized the program for my studies in 2014, and Varalakshmi Wuyyuru facilitated the development of a more polished research program that has been used for target localization since 2015.

Gayatri Kaskhedikar, Ph.D. and Liancheng Yang spent much time determining the best ways to use Dr. Dagnelie's tower-mounted eye tracker. I had the fortune using the eye tracker after they worked out most of the bugs, which made adaptations for my experiments significantly easier. Dr. Roy and Dr. Caspi collaborated with Paul Rosendall, Jason Harper, Kapil Katyal, and others at the Johns Hopkins University Applied Physics Laboratory to integrate the use of eye-tracking glasses into Argus II stimulation. Although few target localization tests with the eye-tracking glasses were performed, the data from that setup were critical for explaining other results in this dissertation.

My entry into Ph.D. training, and the success I found along the way, would not have been possible without the inspiration and guidance I received from numerous educators throughout my academic career. My neuroscience professors, particularly Stewart Hendry, Ph.D., Linda

Gorman, Ph.D., and Brenda Rapp, Ph.D., exposed me to the wonderful complexities of the nervous system. They fueled my desire to further investigate the nervous system and its workings, particularly how one might interact with it through neuroprostheses. Takashi Yoshioka, Ph.D., as the director of my B.A./M.S. program, expanded my insight on research and publications, as well as bolstered my comfort with research presentations. Long before going college, I also had the benefit of having Ronald Heitmann and Linda Barnfield as middle school teachers. Mr. Heitmann and Mrs. Barnfield were committed to teaching their students about the world and helping them grow as individuals. Through efforts like the "We Didn't Start the Fire" project and tours of Europe, they contributed greatly to my perseverance and general thirst for knowledge.

Beyond any other influence, I owe my parents, Patrick and Linda Barry, the most for my academic achievements. My mom shared with me her love of science and biology, and my dad bestowed his interest in engineering and capacity for evaluating situations. I could not have excelled without their depths of kindness, generosity, and wisdom. Alongside them, my granduncle Curtis McFarland, Ph.D. has long stood as a particular inspiration for achieving a doctoral degree. My aunts, uncles, and many cousins have served as a pool of unending support and grounding. Irma Zhang, offering constant confidence, patience, and encouragement, anchored my sanity throughout each of the challenges I faced. Thank you to all who helped make this work possible!

Contents

Figures

Tables

1 Introduction

Blindness isolates and hampers the activities of millions around the world. Some individuals, while being legally blind and suffering severe visual deficits, can still perceive light and discern general shapes. Their vision is simply not capable of capturing and processing information with the quality assumed for many societal functions, such as reading signs or navigating traffic. More profound visual deficits can leave one with only bare light perception. Those with such extensive vision loss may be able to detect light or distinguish patterns of motion, but lack form vision. Rather than having difficulty reading a sign, such people would be challenged to visually tell the difference between a city street and a beach. Complete vision loss constitutes no light perception, and those afflicted cannot detect bright flashes of light even under the best conditions. Many with bare and no light perception have few options for treatment or recovery, even in developed communities. Visual prostheses, electronics that attempt to bypass damaged parts of the visual system, offer at least some hope of improvement for the profoundly blind. Prosthetic stimulation of the visual system is still very young, however, and many challenges for the technology have not yet been explored.

Evolution has crafted an intricate system of molecular interactions to translate photonic energy into neural activity. Fine and complex connections among neurons in the visual system then build response profiles from simple light detection to highlighting points of contrast and eventually to indicating the presence of a face or an object approaching the viewer. Neurons in the retina, before any signals leave the eye, already display assortments of preferences for spatial contrasts in colors or brightness, or sensitivities biased toward moving or static stimuli. Chapter 2 introduces the anatomy of the eye and the subtleties of natural visual processing. Juxtaposed to the fine-tuned, delicate architecture of the normal visual system, Chapter 3

discusses the comparatively neolithic designs of visual prostheses. Despite our knowledge of circuitry in the retina and the rest of visual system, therapeutic stimulation technology does not allow for much more than sending electrical currents indiscriminately throughout large areas that contain many different neurons. Each electrical pulse reaches neurons of numerous types that occupy multiple layers of tissue. Each layer would normally have a distinct role in information processing, and activity in one neuron could indicate conditions exactly opposite those signaled by its neighbor's activity. Current visual prostheses have to ignore these details, and function simply on the principle that electricity can force neurons to become active. Fortunately, the elicited neural activity usually creates an impression of light. Unfortunately, that unnatural blast of neural activity does not carry precise information, and only allows recipients to point to sources of light, detect movement, or, in the best cases, recognize high-contrast shapes.

Even the fairly simple task of pointing to a light, however, is more complex than it appears. One might assume that, given a prosthesis recipient can perceive a light, he or she should be able to point to where it is. Indeed, a prosthesis user can perceive a light coming from a particular direction and point to it. Nevertheless, the alignment of that perceived direction with the actual direction to the light source is far from guaranteed. The apparent location of the light can be very different from its true location, particularly when a prosthesis uses a camera that is separated from the eye. The camera may be pointing in one direction, while the eye is pointed elsewhere. Although the prosthesis will provide stimulation only based on the camera's view, the user will perceive the light to be wherever the eye is aimed. To avoid misalignment between the camera and the eye, the part of the camera's view used to create stimulation can be chosen to coincide with where the eye is looking. The prosthesis setting that controls this alignment is called the camera alignment position (CAP), and is detailed in Chapter 4.

When a CAP is not properly configured, the camera is considered to be misaligned with the eye. Misalignment can be observed through localization errors, such as when a user points too far to the right or too high when trying to indicate the location of a light. Suppose that a user points to the right of a target light. One could make an adjustment by shifting the CAP to the right by the angular difference between the target and the response. Such a modification would require the user to move the camera farther to the left to see the target, and would thereby cause the user to point more to the left. The modified CAP that would have led the user to point directly at the target would be considered the optimal camera alignment position (OCAP). For any individual localization task trial, the OCAP for the specific conditions of that trial would cause the user to point with perfect accuracy. When multiple trials are considered, and when only one CAP can be chosen but no single CAP could achieve perfect accuracy, the OCAP is considered the CAP that would provide the maximum accuracy possible. OCAPs are conceptual constructs for understanding what is necessary to maximize prosthesis-user point accuracy, and are the basis for most of the following research. OCAPs can be considered to exist for any given set of conditions in a localization task, but because the determining factors for accuracy in prosthetic vision are unknown, OCAPs are only estimated and considered post hoc in this text.

Among the unknowns related to pointing with prosthetic vision, it is unclear whether a user with a misaligned camera can improve pointing accuracy over time. A CAP set 10° too far to the left might be considered analogous to prism glasses that impose visual shifts 10° to the right. In both cases, the viewer would initially be inclined localize lights and objects about 10° to the right of their true locations. In the case of prisms, viewers with otherwise normal vision are known to adapt to the perceptual offset rather quickly. Here and in later chapters, any consistent reduction in the effect of an applied condition is considered adaptation. In this context, adaptation to misalignment would be consistent improvement in localization accuracy without

altering the original source of misalignment. Given time to practice visuomotor coordination, viewers with prism glasses can learn to localize and interact with objects accurately within minutes or hours. Such adaptation, however, is dependent on viewers' ability to view their own actions with the shifted vision. Users of visual prostheses have very poor vision, and likely do not receive enough information to intuitively perceive their own bodies in action. Prosthesis users should therefore, without specific training or dedicated auditory feedback, not be able to adapt to misaligned cameras. Pointing with misaligned cameras was tested with and without automated verbal feedback, and the results of analyses on user accuracy and OCAPs are given in Chapter 5. All observed prosthesis users failed to show adaptation in absence of verbal feedback. Because they lack the ability to adapt, users of visual prostheses will require measures to ensure cameras are properly aligned, or will need special training in how to compensate for misalignments.

In light of their failure to adapt to misalignment, one might question how well proper camera alignments persist over time. If users could adapt to misalignments, that adaptation would serve as an anchor for any established alignment. Alignment dictates that CAPs and OCAPs have the same values, and a misalignment would develop if the OCAP ever changed to be different from the CAP. If users could adapt to misalignment, the process of adaptation would drive the OCAP back to the CAP. OCAPs would thus never be expected change from the CAP. Lacking adaptation, however, the question whether OCAPs change becomes far more important. Without adaptation on the user's side of camera alignment, the system would need to be consistently updated to match the CAP to the OCAP. Prosthesis users performed localization tasks for up to 5.3 years, and CAPs were set as close as possible to OCAPs at each visit. Results are detailed in Chapter 6. OCAPs indeed changed over time, such that only 16% or fewer of possible combinations of 30-day OCAP averages had differences that were not statistically significant.

OCAPs moved along trends within days and trial runs, and the directions of those trends tended to match changes between days or between runs.

It is known that the direction of the eye's gaze is important in localizing any visual input. For example, if someone perceives a flash while the eyes are pointing to right, the flash will be seen as though it came from the right. Eyes pointing to the left would similarly place the flash to the viewer's left. This makes perfect sense in normal vision, as eyes not pointed in the direction of the flash would fail to see the light. In prosthetic vision, however, a flash can be perceived any time an electrode provides current. The orientation of the eyes is thus very important to how light is perceived, but it does not change anything in how a prosthesis stimulates the visual system. Details on how the brain factors eye orientation into visual perception, and effects that have been observed in prosthetic vision, are discussed in Chapters 2 and 4.

When OCAPs change, one might consider changing eye orientation as a possible cause. While sighted individuals can calibrate their perception of eye orientation, such as when the eyes are fixed straight ahead, based on visual input, the blind do not have that kind of feedback. Average eye orientations for a blind person could therefore change subtly over time, without the person ever feeling or correcting the changes. If average eye orientations were to change, so would the average locations of elicited percepts. Percept locations would change in both egocentric space and in the camera's field of view, thus changing a user's OCAP. Chapter 7 provides details on how user eyes were tracked both separately from the normal localization task and with eye-tracking glasses during target localization. Despite considerable variability, distinct relationships were found between eye orientations and OCAPs. Both displayed the same behavior in day or trial run trends, and trends of the 2 measurements tended to move in the same directions.

When CAPs were automatically adjusted based on eye orientation, changes in localization errors were significantly slower than seen with normal, unadjusted CAPs.

These observations of camera alignment apply not only to most current retinal prostheses, but also to any visual prostheses that will interface with the visual system inside the skull. Prostheses that use stimulators in the thalamus or visual cortex will have similar problems in handling misalignments between the camera and the eyes. Until prosthesis technology leaps forward to offer the blind the ability to clearly see their own actions, prosthesis users will not have the information to intuitively adapt to camera misalignments. Even when cameras are initially aligned, OCAPs do not stay constant. Observed OCAP movements motivated the realization that average eye orientations in blind individuals change over time. Just as individual eye movements create transient misalignments and OCAP changes, shifts in average eye orientation drive much of long-term OCAP changes. In order to maintain maximal localization accuracy for users of most visual prostheses, OCAP changes will need to be taken into account through regular camera realignments, perhaps with the aid of eye tracking.

2 The visual system

Vision, as experienced by humans, is related to activity in the brain's visual cortex. Visual prostheses attempt to recreate proper visual cortex activity when the natural visual system fails to do so. In order to understand the magnitude of challenges in vision restoration, and why visual prostheses fail to recreate normal vision, one should appreciate the intricacies of how the visual system normally functions. Light enters the eyes, and the light captured by the eye's retina is the brain's only source of visual information. The retina processes some of the visual information and then sends it to brain via the optic nerve. Stimulation from the optic nerve elicits intricate patterns of brain activation that correspond to detected stimuli. Neurons in visual cortex interact with neurons in other brain areas to stitch visual percepts together with those of other modalities to form our general perception of the world.

Transforming signals from light to action potentials, and ultimately perception, involves many complex processes. Failure at any stage of signal transduction or transmission can result in reduced visual perception. Visual function, activity that requires visual information, is naturally hampered by visual deficits, but the extent and dynamics of disabilities are not always well-known. The challenge of restoring visual function is even more complex and less understood.

2.1 Anatomy of the eye

Each eye is roughly ellipsoid in shape, save the bulge of the cornea and the optic nerve. Most of the eye's volume is occupied by the vitreous body, a gel that helps maintain the structure of eye and positions of the lens and retina. Most of the eye is covered by the sclera, connective tissue that principally protects and determines the shape of the eye. Below the sclera, the choroid is

the vasculature for the outer retina. The anterior portion of the eye is mostly dedicated to optics, manipulating the paths and amount of light entering the eye. At the back of eye, juxtaposed to the vitreous, the retina is a layer of neural tissue that detects and transduces photon signals and extends as far forward as the ora serrata.

The anterior-most portion of the eye is the cornea. The cornea is composed of several layers of transparent tissue, and has generally high collagen composition. It provides most of the eye's refractive power, as well as protecting the eye's interior from infection. Light enters the eye by first passing through the cornea. The difference in air's and the cornea's refractive indices, combined with the cornea's curvature, causes light paths to converge. If the curvature of the cornea is altered, as in refractive surgery, the degree of convergence can be increased or decreased to aid in focusing light paths onto the retina.

After passing through the cornea, light enters the aqueous humor of anterior chamber. As the name implies, the aqueous humor is mostly water with some solutes. Intraocular pressure is modulated by the production and drainage of the aqueous humor. In the anterior chamber, aqueous humor has already passed through the posterior chamber, around the lens and iris, and flows out of the eye through the trabecular meshwork. Increases in resistance through the trabecular meshwork reduces aqueous humor outflow and thus increases intraocular pressure. Abnormal increases in intraocular pressure can cause damage to the optic nerve known as glaucoma.

The iris prevents a portion of the light passing through the aqueous humor from reaching the lens. This barrier is heavily pigmented and serves to limit the amount light that reaches the retina. The iris connects to the sclera at its peripheral boundaries, and only permits light to pass through the pupil, a circular hole defined by its central boundary. Sphincter and dilator muscles

can alter the shape of the iris to respectively reduce or increase the diameter of the pupil. Gross light levels falling upon the retina, as captured by intrinsically-photosensitive retinal ganglion cells (ipRGCs), trigger constriction of the pupil in the pupillary light reflex.

Light that passes through the aqueous humor of pupil next strikes the lens. The lens supplements the refractive power of the cornea, permitting selective focusing within a range of distances. Zonule fibers suspend and connect the lens to the ciliary body. The ciliary muscle in the ciliary body controls the curvature of the lens, and thus the lens's focusing power. As a ring-shaped muscle concentric with the lens, contraction of the ciliary muscle removes tension from the zonule fibers and allows the lens to take on a more spherical shape. This increase in lens curvature provides more focal power for viewing closer targets. Relaxation of the ciliary muscle increases tension on the zonule fibers, which in turn decreases the curvature of the lens. The ability of the lens to become more spherical decreases with age, eventually to the point of presbyopia, when corrective lenses may be required for short-distance focusing.

The ciliary body also includes the ciliary epithelium, which produces the aqueous humor. The aqueous humor flows out of the ciliary epithelium into the posterior chamber, which mostly takes up volume behind the periphery of the iris. On the opposite side of the ciliary body and zonule fibers lies the vitreous body. Like the aqueous humor, the vitreous body is mostly composed of water, but is much more viscous. The vitreous body is not replenished over time, and exhibits changes with aging.

After being focused by the cornea and lens, light passes mostly uninterrupted through the vitreous body to the posterior end of the eye. Light then meets and interacts with the retina. Less than a millimeter thick and lining more than half of the eye's interior, the retina is responsible for detecting light and generating visual signals for the brain. Given how the cornea

and lens invert paths of light, the retina corresponds to an inverted visual field. The optical center of the retina is colinear with the centers of the cornea and the lens, and this line is considered the optical axis of the eye. Spreading out from the optical axis, the superior retina captures the inferior visual field, and the temporal retina captures the nasal visual field. One degree of visual field corresponds to about 280 μm of retina near the optical axis, but the μm/° relationship drops with retinal eccentricity to less than 150 μm/° in the far periphery[1].

Although not lying directly on the eye's optical axis, the fovea is the part of the retina that corresponds to the center of vision. Eyes move to land photons from visual targets on the fovea, rather than along the optical axis. The retinal circuitry and corresponding cortical representations are best suited for high-acuity viewing with the fovea. The alternative axis defined by the fovea and the target of interest is considered the visual axis of the eye.

The retina consists of several layers of tissue. The innermost layer, that which light would strike first, is the nerve fiber layer. The nerve fiber layer is separated from the vitreous by the inner limiting membrane and is formed by axons of retinal ganglion cells (RGCs). These axons converge at the optic disc, nasal to the macula, and pass through fenestrated tissue continuous within the sclera called the lamina cribrosa. No photoreceptors are present at the optic disc, and this creates the physiologic blind spot in each eye. Beyond the lamina cribrosa, the RGC axons join to form the optic nerve. Although the optic nerve is myelinated beyond the retina-optic nerve junction, axons within the retina are typically not myelinated[2]. Myelination of retinal axons has been associated with visual deficits, such as visual field defects, and when more severe, anisometropic myopia, amblyopia, and strabismus[3].

Beneath the nerve fiber, the somata of the RGCs and displaced amacrine cells form the ganglion cell layer. Cells in this layer participate in the final stages of retinal processing of light. Amacrine

cells have short neurites and modulate the activity nearby RGCs, bipolar cells, or other amacrine cells. All RGCs project axons into the nerve fiber layer to conduct signals to the brain, mostly reacting to inputs from other retinal neurons. Amacrine cells are not well studied, but may include as many as 40 or more different cell types[4]. RGCs also exhibit considerable variation, with at least 17 different types in the primate retina[5]. The neurons of the ganglion cell layer extend neurites farther along the path of light into the inner plexiform layer, interacting with neurites from the inner nuclear layer.

The inner nuclear layer is the next farthest layer from vitreous body. This layer consists primarily of interneurons and Müller cells. Although some amacrine cells are found in the ganglion cell layer, most exist in the inner nuclear layer. Bipolar cells receive input from photoreceptors and accordingly stimulate RGCs. Horizontal cells, having lateral connections and orientation, modulate the output of photoreceptors based on the activity of neighboring photoreceptors. Most of these interneurons, like photoreceptors and unlike RGCs, do not produce action potentials; only some bipolar and amacrine cells show such spiking activity. Müller cells are retinal glia with somata in the inner nuclear layer, but have processes extending to and forming both the inner and outer limiting membranes. These membranes, respectively at the border of the vitreous body and between the photoreceptor somata and inner segments, are barriers that partially insulate the extracellular environment of the retina. Müller cells also envelope the neurons of the retina, providing support similar to that of astrocytes elsewhere in the nervous system.

The interneurons of the inner nuclear layer form connections with photoreceptors through the outer plexiform layer. All synapses in the outer plexiform layer involve at least one photoreceptor synaptic terminal, either a cone pedicle or rod spherule. These terminals attach

to the main bodies of the photoreceptors. Photoreceptor cell bodies principally contain the photoreceptor nuclei, and collectively form the outer nuclear layer of the retina. Cones, one of two general types of photoreceptors, have cell bodies that line the outer edge of the outer nuclear layer. Cell bodies of rods, the other type of photoreceptor, can be found scattered throughout the remaining inner span of the outer nuclear layer.

The outermost parts of the neural retina, the photoreceptor inner and outer segments are dedicated to photon capture and signal transduction. The outer limiting membrane is formed by Müller cells joining with the inner ends of the inner segments, creating an extracellular barrier between the levels of the photoreceptor cell bodies and inner segments. This barrier contains matter related to phototransduction to the interphotoreceptor space, preventing diffusion around the rest of the neural retina. The inner segment itself supports phototransduction with mitochondria, ribosomes, and opsin production.

The outer segment is connected to the inner segment by a cilium and carries out the actual process of photon capture and transduction. Photons have traveled mostly uninterrupted from the lens surface and through the rest of retina before reacting here with the photoreceptors. The outer segment consists of stacked membrane discs that contain light-reactive photopigments. These photopigments, opsins, indirectly trigger changes in photoreceptor transmembrane potential, and thus create a neural signal.

Rods and cones derive their names from the shapes of their outer segments. In both, membrane discs are generated at the base of the cilium, close to the inner segment, move toward the outer tip of the outer segment, and are ultimately shed and phagocytosed. In rods, the discs separate from the outer segment membrane and the cilium. The discs float stacked within the outer segment membrane, each about the same diameter, giving rods a cylindrical appearance. Unlike

the in the rods, discs of the cones remain connected to the outer segment membrane and the cilium. As the cone outer segment grows, in proportion with shedding, it grows more in length than in width. This pattern of growth gives cone outer segments their conical shape.

Retina providing central vision tends to be structurally different from peripheral retina. While rods may dominate most of the retina, cone density spikes in the area of central vision and rod density falls to 0 near the center. Photoreceptors in this specialized region are typically elongated and packed more tightly together. These features provide greater ability to capture incident photons and achieve more spatial specificity, particularly in bright conditions. Inner retinal vasculature is also less prominent in this region, reducing chances of blood vessels interfering with photon paths.

In humans and other primates, no inner retinal vasculature is apparent within a radius of approximately 80 µm from the retina's visual center[6, 7]. This lack of blood vessels permits greater elasticity in the avascular zone compared to the rest of the retina. Circuitry in the central retina also lends to greater flexibility, as most bipolar cells connect to only one cone, and most RGCs to only one bipolar cell[8]. This contrasts with broader, converging and lateral connections prevalent elsewhere in the retina. Intraocular pressure acts on the developing retina to deform this elastic patch and push aside layers of the inner retina[9]. The resulting pit is known as the fovea centralis. At the center of the fovea, the retina only consists of cone cell bodies, inner and outer segments, and Müller cell processes. All other structural elements, even the cone pedicles, are pushed aside and the cones have their maximal exposure to incoming photons. This region lacking cone pedicles is considered the foveola and spans approximately 200 µm in diameter[10].

While inner retinal vasculature is absent in foveola, photoreceptors there and throughout the retina are metabolically sustained by the retinal pigment epithelium (RPE). The RPE is a single

layer of epithelial cells that delivers nutrients to the photoreceptors and removes metabolic waste from the subretinal space. High pigment concentrations allow the RPE to absorb photons not captured by photoreceptors and prevent photons from scattering within the eye. Light absorption by the RPE also helps to prevent damage related to photooxidation in the RPE and deeper tissues[11].

Photooxidation takes a significant toll on photoreceptors as they detect and transduce light signals. Photoreceptor discs at the tip of the outer segment, near the RPE, are the photoreceptor's oldest discs and most likely to contain damaged proteins and lipids and toxic products of photooxidation. Photoreceptors shed their outer discs daily, and the RPE removes them through phagocytosis. Molecules critical to phototransduction that are not synthesized by the photoreceptors get recycled for continued use by the photoreceptor.

The RPE specifically supports phototransduction by isomerizing the chromophore retinal and maintaining extracellular ion concentrations. Retinaldehyde, often shorted to retinal, is the key molecule that allows photoreceptors to react to photons. When a proper retinal isomer, 11-*cis*-retinal, bound to an opsin absorbs a photon, conformational changes trigger signal cascades that change photoreceptor membrane permeability to sodium ions. The change in flow of ions alters the photoreceptor's transmembrane potential, and thereby changes the photoreceptor's neurotransmitter output and creates a neural signal. Before a retinal molecule can be reused to for photon capture, the RPE must restore it to its 11-*cis* configuration. In order for photoreceptors to continue signaling photon captures reliably, the RPE counteracts changes in subretinal ion concentrations by altering the transport of ions across its membrane.

The blood supply of the RPE, and thereby the photoreceptors, is provided by the choroid. The choroid lines the eye outside the RPE, containing several layers of tissue that span 100–200 µm

in thickness[12]. Bruch's membrane constitutes the choroid's inner surface, itself several layers of tissue that separate the choroid capillaries from the RPE. Progressively larger blood vessels permeate the choroid as distance from the RPE increases. Great oxygen tension gradients are maintained to effectively drive oxygen through Bruch's membrane and the RPE to the photoreceptors. Beyond the choroid vasculature, large melanocytes provide further protection from photons not initially captured by photoreceptors or the RPE. The human choroid contains more melanin than the RPE[13], but its function may lean more toward control of reactive oxygen species than capturing stray photons[14]. Closer to the sclera, nonvascular smooth muscles may reduce the thickness of the choroid as they contract, thereby adjusting the position of the retina in response to poor focusing.

The lamina fusca marks the divide between the choroid and the sclera. It contains melanocytes and fibroblasts, along with myelinated axons. There is some disagreement as to whether the lamina fusca is outermost layer of the choroid or the innermost layer of the sclera[12, 15]. Beyond the lamina fusca lies the stroma of the sclera, consisting of many collagen fibrils that give the sclera its structural strength. Extraocular muscle tendons ultimately attach to the sclera in the stroma. Metabolic requirements of the stroma are relatively low, and no capillaries directly supply this structural layer under normal conditions. Instead, nutrients for the stroma are supplied by the choroid and vasculature in the superficial layers of the sclera. The episclera, a well vascularized layer of connective tissue, surrounds the outside of the stroma, most prominently in the anterior half of the eye. Tenon's capsule sits above the episclera as a layer of collagen bundles that separates the sclera from orbital fat.

2.2 Phototransduction

The structure and function of the eye and its constituent tissues revolve around detecting light information that is useful to the organism. All such detection relies on one retinaldehyde isomer: 11-*cis*-retinal. All mammal photoreceptors package this isomer into transmembrane proteins called opsins. When 11-*cis*-retinal absorbs a photon, the configuration of retinaldehyde about its double bond between the 11th and 12th carbon atoms flips to create all-*trans*-retinal. This isomerization induces a conformational change in the bound opsin, which in turn activates a signaling pathway in the photoreceptor.

Figure 1 – Phototransduction cascade

A photon activates rhodopsin by photoisomerizing 11-*cis*-retinal. Active rhodopsin subsequently activates the α subunit of transducin. Transducin α binds to the γ subunit of phosphodiesterase (PDE), relieving the γ-subunit inhibition on the α and β subunits. PDE hydrolyzes cytosolic cyclic guanosine monophosphate (cGMP) to 5'-GMP. The reduced concentration of cGMP causes cyclic-nucleotide gated (CNG) cation channels to close. Guanylate cyclase activating protein (GCAP) allows guanylate cyclase to produce more cGMP. G-protein receptor kinase (GRK) rhodopsin kinase phosphorylates rhodopsin to permit deactivation by arrestin. Recoverin inhibits the activity of rhodopsin kinase, and thereby impedes the deactivation of rhodopsin. Figure created by Veleri et al.[16], used under the Creative Commons Attribution License 3.0.

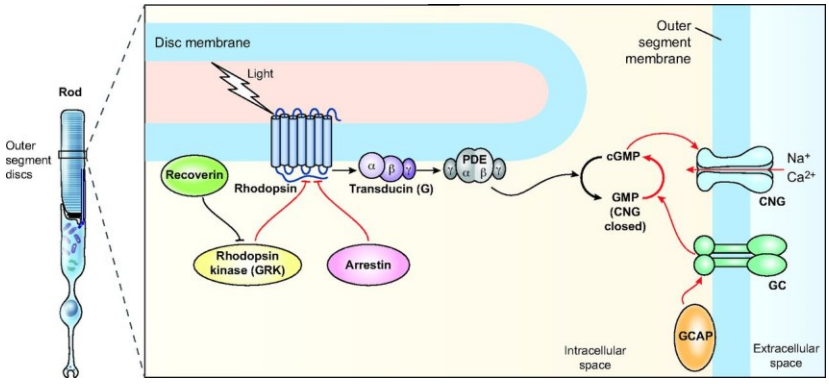

11-*cis*-retinal, when part of an unprotonated Schiff base in solution, maximally absorbs photons with ultraviolet wavelengths of about 360 nm[17]. Protonation of the Schiff base increases maximal absorption wavelength into the visible spectrum, to about 440 nm. Retinaldehyde

binds to opsins through protonated Schiff base linkage, and other interactions with the opsin can further adjust the maximal absorption wavelength, referred to as opsin shift[18]. In rods retinaldehyde is bound to rhodopsin, shifting maximal absorption in humans to 493 nm[19]. Each cone can have any one of three cone opsins, which cause retinaldehyde to absorb maximally at 426, 530, or 552–557 nm[20]. Melanopsin in ipRGCs shifts retinaldehyde maximal absorption to 479 nm[21]. These absorption peaks are consistent within each photoreceptor type, and no photoreceptor uses more than one type of opsin. In the case of cones, signals generated by different opsin types in separate photoreceptors are compared to gain wavelength information about stimulus photons.

All opsins are class-A G-protein coupled receptors (GPCRs)[22]. Being GPCRs, opsins consist of seven transmembrane helices and attach to a guanine nucleotide-binding protein[23]. Molecular structure details and binding patterns vary amongst opsin types. Rhodopsin is the best studied opsin, and it is known to form dimers in the rod disc membranes[24]. Each dimer can bind to one transducin, the heterotrimeric G-protein involved in the next step of phototransduction. Dimers may be further organized into parallel tracks with rod discs[25]. Simulation studies suggest that transducin molecules may bind to and rapidly dissociate from inactive rhodopsin dimers within tracks, and transducin may be largely trapped within a track until it encounters an active dimer.

Photon capture by a rhodopsin-bound retinaldehyde induces conformational changes in rhodopsin, with multiple labeled phases such as batho-, lumi-, and metarhodopsin[26]. Metarhodopsin is considered the active state of rhodopsin, and is further divided in two principal phases. Metarhodopsin I, specifically metarhodopsin Ib, releases guanosine diphosphate (GDP) from the alpha subunit (Tα) of a bound transducin[27]. This involves metarhodopsin I deprotonating its Schiff base and movement of transmembrane helices 5 and 6

to become metarhodopsin II. Metarhodopsin II facilitates binding of guanosine triphosphate (GTP) to Tα, thereby activating transducin.

Configurations of metarhodopsin I and II can interchange at up to kilohertz frequencies[28], but the rate of action in rods is typically limited by reactant concentrations[29]. In rods, rhodopsin can rapidly bind, activate, and release hundreds of transducin molecules per second. This rate, however, may be limited to the supply of transducin within a rhodopsin track[25].

Deactivation of rhodopsin occurs quickly, with a time constant of approximately 40 ms[30]. GPCR kinase 1 phosphorylates rhodopsin, which reduces its activity with transducin and increases affinity for the arrestin-1 protein. Arrestin-1 easily binds after 3 rhodopsin C-terminal sites are phosphorylated, and blocks further interaction with transducin. The Schiff base link with all-*trans*-retinal is hydrolyzed, separating rhodopsin and retinaldehyde. New 11-*cis*-retinal can bind to rhodopsin, and arrestin detaches[31]. Phosphatases slowly dephosphorylate rhodopsin to restore its full signaling capacity[32].

Rhodopsin is estimated to activate 8–16 transducin molecules on cytoplasmic side of the rod disc in response to capturing a photon[25, 33]. Once activated, the Tα-GTP subunit separates from the still-joined beta (Tβ) and gamma (Tγ) subunits, and both pieces of transducin separate from rhodopsin. Tα-GTP and Tβγ, through lipid modifications on the Tα and Tγ, briefly remain in close contact with the disc membrane[34]. The activated Tα acts on a membrane-bound enzyme: phosphodiesterase 6 (PDE6). Tα specifically binds to a gamma subunit of PDE6 (PDE6γ) and relieves PDE6γ's inhibitory effect on PDE6. Tα and PDE6γ remain linked to the rest of PDE6, and the entire complex stays anchored to the disc membrane[35].

As Tα removes the inhibitory effect of PDE6γ, PDE6γ drastically increases the affinity of Tα-GTP for the regulator of G protein signaling 9 (RGS9) protein[36]. RGS9 promotes the GTPase activity of Tα, causing Tα to hydrolyze its bound GTP to GDP and an inorganic phosphate ion[37]. The resulting Tα-GDP dissociates from PDE6 and restores the inhibition of PDE6γ. Tα-GDP then binds to a Tβγ complex, reforming inactive transducin. Approximately 3–4 times more Tβγ exists in the rod outer segment than Tα, which is thought to facilitate the restoration of inactive transducin[38].

PDE6 has 2 gamma subunits, and is maximally activated when both are bound to Tα-GTPs[39]. When active, the catalytic subunits of PDE6 hydrolyze cytosolic cyclic guanosine monophosphate (cGMP) to 5'-GMP[40]. In rods, the catalytic subunits form a heterodimer (PDE6αβ), whereas the catalytic subunits in cone PDE6 form a homodimer ($PDE6\alpha'_2$)[41]. Just as rhodopsin activated multiple transducin molecules, PDE6 further amplifies the signal of photon capture by rapidly hydrolyzing many cGMP molecules. When PDE6 is maximally activated, it is able to hydrolyze up 8000 cGMP molecules per second, close to the limit imposed by the diffusion of cGMP[42]. The rate of transducin deactivation is 2–3 times slower than that of rhodopsin[43], and PDE6 is not always fully activated, so one can estimate approximately 2000 cGMP molecules hydrolyzed per photon capture[33].

Cytosolic cGMP interacts with cyclic nucleotide-gated (CNG) cation channels in the rod and cone plasma membranes. In rods, these channels are formed as heterotetramers, with 3 A1 subunits (CNGA1) and 1 B1a subunit (CNGB1a)[44, 45]. Each subunit has its own binding region for cGMP, and the CNG channel requires 2–4 bound cGMP molecules to open. Binding more cGMP molecules increases the stability of the open state, and thus the probability that a CNG channel will be open[44]. The probability of a CNG channel opening with no bound cGMP is extremely low.

cGMP molecules rapidly associate and dissociate with CNG channels, so a reduction in cGMP concentration quickly reduces the probability of channels being open, and thereby the expected number of open channels. CNG channels, opened by cGMP, admit various cations into the rod outer segment, including Ca^{2+}. While the channels are open, Ca^{2+} binds to guanylate cyclase activating proteins (GCAPs) 1 and 2 and inhibit their activity. When the channels close, however, this Ca^{2+}-driven inhibition is relieved and the GCAPs activate guanylate cyclase (GC) on the disc membranes. GC counteracts the cGMP hydrolysis of PDE6 by synthesizing more cGMP from GTP[46].

Photon capture thus results in the closure of CNG channels in the rod plasma membrane. The maximal number of open CNG channels is achieved in the dark, but this is only about 10^4 open channels[47], or about 1% of the rod CNG channels[48]. One captured photon can reduce the number of open CNG channels by 5%[49, 50], or several hundred channels.

As the CNG channels exist on the plasma membrane surrounding the stack of discs, and the discs themselves limit the diffusion of cGMP along the long axis of the rod, the highest concentration of closed CNG channels appears at the same axial point as the photon-activated rhodopsin. The light-induced drop in cGMP concentration, and thus open channels, appears as a point sink when plotted against axial position. Drops in cGMP and open channels spread beyond the proximal discs, but are markedly reduced. The ultimate effect on rod intracellular composition from activation in a single compartment still spreads sufficiently to create a detectable signal in the rod's activity[46].

While open, CNG channels permit passage of various cations across the rod's plasma membrane. A consequence of this ion flux, as more positive ions flow inward than outward, is a net rise in transmembrane potential, with the extracellular potential as the reference. As the

resting potential for the photoreceptor lies around −65 mV, positive changes in membrane potential are considered depolarization. Negative changes in potential, as when more CNG channels close, are hyperpolarizations or repolarizations. Maximal depolarization in darkness can bring the transmembrane potential up to −30 mV, and full closure of CNG channels in light hyperpolarizes the cell to its resting potential of −65 mV[51].

Across this physiological range in transmembrane potential, CNG channel ion currents do not display much variance. Inward current primarily consists of Na^+, Ca^{2+}, and Mg^{2+} ions, while K^+ contributes a small outward current. The current is dominated by Na^+ influx, while Ca^{2+} and Mg^{2+} each contribute 14–20% of the total inward current[52]. Although Na^+ dominates the current, flux of divalent ions through CNG channels is disproportionately high relative to ion concentrations. Heightened CNG channel affinity for divalent cations facilitates Ca^{2+}-mediated feedback, and limits CNG channel conductivity by transiently blocking the channel pore[53].

A single-photon response can result in a modest hyperpolarization of about 1 mV in rods. This is, however, much greater than the 5 μV hyperpolarization in cones[54], reflecting the greater utility of rods over cones in low-light conditions. These hyperpolarizations spread from the outer segment to the photoreceptor terminal and alter the release of synaptic vesicles from photoreceptors. Both rods and cones contain L-type voltage-gated Ca^{2+} channels in their terminals[55]. These channels open with greater probability at more positive transmembrane potentials, such that 10–15% may be open at −35 mV. Ca^{2+} admitted by CNG is mostly confined to the outer segment, so Ca^{2+} flow into the terminals is almost completely dependent on these voltage-gated channels[56]. Hyperpolarization of 1 mV, as in the rod single-photon response, is sufficient to reduce terminal intracellular Ca^{2+} concentration by about 20%[57].

Ca^{2+} entry into photoreceptor terminals causes the release of synaptic vesicles. These vesicles contain L-glutamate, and release this neurotransmitter into the synaptic cleft by merging with the plasma membrane and exocytosing their contents. In cones, this relies entirely on synaptic ribbons, to which numerous vesicles are preloaded when release is not being triggered. At light-offset depolarization, all ribbon-loaded vesicles are quickly released, and further release is limited by Ca^{2+}-dependent loading of new vesicles[58]. Fewer ribbons are present in rod spherules, but rods also contain non-ribbon release sites for slow, continuous glutamate release in darkness. As Ca^{2+} influx falls during protracted depolarization, rods utilize Ca^{2+}-induced Ca^{2+} release from the terminal endoplasmic reticulum to maintain synaptic activity[59]. Improving the consistency of glutamate release this way can help improve the signal-to-noise ratio for a rod's single-photon response.

Released glutamate elicits responses from post-synaptic neurons, and forms the basis of the neural signal derived from photon capture. Specifically, photon capture is signaled by a reduction in the release of glutamate, as rods and cones continually release glutamate in the dark. Once the photon signal has been transduced into a neural signal, other neurons can work to combine that signal with others to form visual percepts.

2.3 Retinoid cycle

Rods lack the ability to produce 11-*cis*-retinal. Cones in some species may be able to indirectly recycle limited amounts of all-*trans*-retinal into 11-*cis*-retinal, but still rely on external retinoid metabolism[60]. Without support of the RPE and Müller cells, ciliary photoreceptors would not have the key photoactive ingredient to for capturing photons. Retinal support cells play the

critical role of initially producing 11-*cis*-retinal and recycling all-*trans*-retinol into 11-*cis*-retinal or 11-*cis*-retinol[61].

Figure 2 – Retinoid cycle

All-*trans*-retinol (atRol) is delivered from the choroidal blood supply by transthyretin (TTR) and retinol binding protein 4 (RBP4). A receptor protein (STRA6) moves atRol into the RPE cytoplasm. Lecithin retinol acyltransferase (LRAT) esterifies atRol to retinyl esters, which accumulate as retinosomes. The RPE65 enzyme hydrolyzes and reisomerizes retinyl esters to 11-*cis*-retinol (11cRol). 11-*cis*-retinol dehydrogenases (11cRDH) oxidize 11cRol to 11-*cis*-retinal (11cRal). Interphotoreceptor retinoid-binding proteins (IRBP) assist with the movement of 11cRal through the interphotoreceptor matrix (IPM) to the photoreceptors. In the rod outer segment (ROS) shown, 11cRal binds to chromophore-lacking rhodopsin and is photoisomerized to all-*trans*-retinal(atRal). All-*trans*-retinol dehydrogenases (atRDH) reduce atRal to atRol, which returns through the IPM to the RPE. Figure created by Malechka et al.[62], used under the Creative Commons Attribution-Noncommercial-No Derivatives License 4.0.

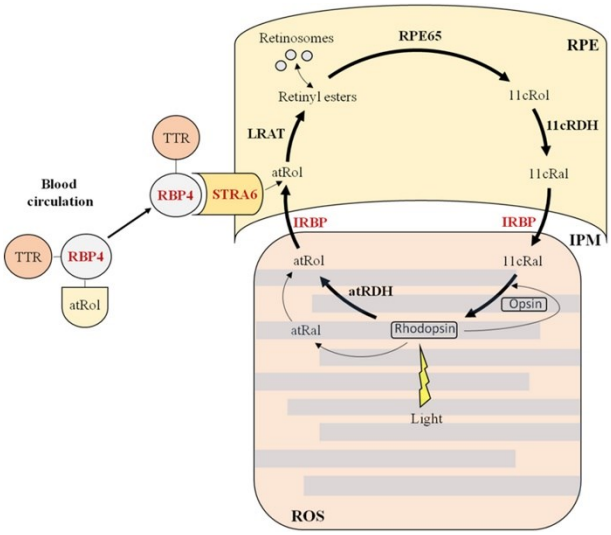

Retinaldehyde and its related compounds, called retinoids, are derived in animals from dietary intake. Animals lack the ability to synthesize retinoids, and these substances are highlighted for consumption as vitamin A. Dietary all-*trans*-retinyl esters and carotenoids such as β,β-carotene are mostly converted to all-*trans*-retinol within the intestines. Newly formed and dietary all-*trans*-retinol then gets esterified to retinyl esters inside the intestinal mucosa by the enzyme lecithin retinol acyltransferase (LRAT). Retinyl esters get packaged into chylomicrons and are circulated through lymph to the bloodstream[63].

Retinyl esters are primarily stored in the liver. Retinyl ester hydrolases in the liver convert retinyl esters back to all-*trans*-retinol based on retinoid demand. All-*trans*-retinol in the liver binds to retinol binding protein 4 (RBP4), and is subsequently released into blood plasma. The plasma contains the protein transthyretin, which fluctuates between monomer and homotetramer forms. The binding of one RBP4 to a transthyretin tetramer stabilizes the tetramer and forms the retinol transport complex[64]. Formation of this complex prevents the renal excretion of RBP4 and retinol, increasing the delivery of retinol to the eye[65].

Once the retinol binding complex reaches the choroidal vasculature of the eye, it binds to a receptor protein on plasma membrane of the RPE. This receptor corresponded to a gene identified as being stimulated by retinoic acid, and both the gene and receptor were therefore named stimulated by retinoic acid 6 (STRA6)[66, 67]. STRA6 removes all-*trans*-retinol from RBP4 and admits the retinol into the RPE cytoplasm[68].

In the cytoplasm, free all-*trans*-retinol gets transported by cellular retinol-binding protein 1 (CRBP1) to the RPE endoplasmic reticulum (ER). LRAT is attached to the membrane of the smooth ER and converts the all-*trans*-retinol back to retinyl esters. The retinyl esters subsequently combine to form retinosomes[69], a process which facilitates net uptake and local storage of retinoids by the RPE.

Retinosomes tend to collect near the plasma membrane of the RPE. For the next step in conversion toward 11-*cis*-retinal, however, the retinosome retinyl esters must return to the RPE smooth ER to interact with the RPE-specific 65 kDa protein (RPE65), a retinoid isomerase. Retinosome retinyl esters may serve as direct substrates for RPE65 activity, or may require hydrolysis, transport to ER, and re-esterification by LRAT before interacting with nearby RPE65. The precise path from retinosome storage to conversion by RPE65 is not yet known[61, 70]. Once

delivered to RPE65, the enzyme simultaneously hydrolyzes the retinyl ester and changes the retinoid conformation from all-*trans* to 11-*cis*, forming 11-*cis*-retinol[65]. Without RPE65 activity, the retinoid cycle stalls and retinosomes enlarge as the retinyl ester content in RPE cells builds[71].

Before use for phototransduction, 11-*cis*-retinol must be oxidized to 11-*cis*-retinal. Retinol dehydrogenases (RDHs) are enzymes that perform conversions between retinol and retinal, with varying substrate affinities. RDHs are all membrane-bound enzymes, as would be appropriate to interact with generally hydrophobic retinoids[72]. RDH5, RDH10, and RDH11 act on 11-*cis*-retinoids, and attach to the membrane of the RPE ER. Each of these RDHs uses nicotinamide adenine dinucleotide (NAD) and/or NAD phosphate (NADP) as a coenzyme, and coenzyme affinity can effectively determine each RDH's tendency for retinol oxidation or retinal reduction. NAD predominately exists in its oxidized form (NAD^+) in the RPE, whereas the reduced form of NADP (NADPH) dominates its oxidized form[61]. RDH11 preferentially interacts with NADP, but can still bind NAD for oxidizing reactions, and act on both all-*trans* and 11-*cis*-retinoids. RDH5 and RDH10 selectively use NAD, and are thus both dedicated to retinol oxidation. RDH10, however, shows more affinity for all-*trans*-retinol, and RDH5 principally drives oxidation of 11-*cis*-retinol to 11-*cis*-retinal[73]. RDH5 has further been hypothesized to form complexes with RPE65, which might streamline the conversion of retinyl esters to 11-*cis*-retinal, but such complexes have not yet been purified *in vitro*[61, 74].

Just as all-*trans*-retinol was delivered to the ER membrane by a binding protein, cellular retinaldehyde binding protein (CRALBP) carries 11-*cis*-retinal from the ER to the apical RPE plasma membrane near the photoreceptors. CRALBP's binding pocket has a curvature ideally suited for 11-*cis*-retinoids, but will also accommodate 9-*cis*-retinoids. Acidic substances, such as

certain phospholipids in the plasma membrane, may reduce the affinity of CRALBP for retinal and trigger the retinal's release[75].

Rods and cones secrete interphotoreceptor retinoid-binding protein (IRBP) into the interphotoreceptor matrix that separates outer segments from the RPE. IRBP's role in the retinoid cycle is not clear, but it has been believed to assist the flow of retinoids among rods, cones, RPE, and Müller cells[68]. IRBP may also serve to protect retinoids from photodegradation and promote the release of retinoids from plasma membranes[76]. With the help of IRBP and/or by passive diffusion, 11-*cis*-retinal crosses from the RPE to photoreceptor outer segments to be integrated into the process of phototransduction.

Inside the photoreceptor outer segment, 11-*cis*-retinal joins with an inactive, unoccupied opsin, granting the opsin sensitivity to visible light. The exact path from the photoreceptor plasma membrane to the opsin, however, is not yet clear[68]. After photoisomerization to all-*trans*-retinal, as discussed in Section 2.2, all-*trans*-retinal dissociates from the opsin. Most of the dissociated all-*trans*-retinal diffuses toward the photoreceptor cytoplasm, but some instead exits into the intradiscal membrane leaflet[77].

All-*trans*-retinal that enters the intradiscal space can require some assistance to cross the disc membrane to the cytoplasm. Retinal, by itself, is a hydrophobic compound and can easily diffuse across lipid membranes. About 40% of disc membrane is made of phosphatidylethanolamine (PE), however, and PE readily reacts with retinal to form *N*-retinylidene-PE (NRPE). Unlike retinal, NRPE cannot independently cross disc membranes[78]. Left unchecked, NRPE can irreversibly bind to a second retinal, which leads to the buildup of toxic diretinal compounds in the disc membrane[79]. The rim of each photoreceptor disc has transmembrane proteins, ATP-binding cassette transporter 4 (ABCA4), that can remove NRPE from the disc interior. ABCA4 has

binding sites for NRPE, as well as all-*trans*-retinal[80] and PE, and flips its orientation to release substrates from the disc interior toward the cytoplasm. Once exposed to the cytoplasm, NRPE can hydrolyze back to separate all-*trans*-retinal and PE.

Attached to the cytosolic leaflet of the disc membrane, RDH8 reduces all-*trans*-retinal to all-*trans*-retinol. Like other RDHs inclined toward reduction reactions, RDH8 preferentially uses NADP as its cofactor. RDH14 is also located in photoreceptor outer segments, displays affinity for NADP cofactors, and can act on all-*trans*-retinal, as well as 9-*cis*-, 11-*cis*-, and 13-*cis*-retinal[73]. Its specific localization and importance to retinoid recycling, however, is not yet known[68]. RDH12 and RDH13, along with RDH11 in rods, are RDHs located in the inner segments of photoreceptors that can reduce all-*trans*-retinal that diffuses out of the outer segments.

All-*trans*-retinol is moved without substantial structural changes from the photoreceptor disc membranes to the RPE ER for recycling to 11-*cis*-retinal. All-*trans*-retinol is known to mostly move by passive diffusion within the disc and plasma membranes of rods[81]. Minimal amounts of all-*trans*-retinol is dissolved in the photoreceptor cytoplasm. All-*trans*-retinol may transfer between disc and plasma membranes with the assistance of a lipophilic carrier, which could permit rapid equilibration of retinol among these membranes. From the plasma membrane, IRBP facilitates release of retinol into the interphotoreceptor matrix[82]. Once all-*trans*-retinol passes through the RPE plasma membrane, CRBP1 transports it to the ER to reenter earlier stages of the classical retinoid cycle, beginning with esterification by LRAT.

Cones have access to the classical retinoid cycle, but also seem to recycle retinoids by another path. After photoisomerization of 11-*cis*-retinal to all-*trans*-retinal in cones, the short-chain dehydrogenase/reductase retSRD1 reduces all-*trans*-retinal to all-*trans*-retinol in conjunction with RDH8[83]. Instead of moving to the RPE for reisomerization, all-*trans*-retinol can also be taken

up by Müller cells. Dihydroceramide desaturase-1 in the Müller cell ER isomerizes all-*trans*-retinol to 11-*cis*-retinol[84]. The mass-action drive for this isomerization may be maintained by the esterification of 11-*cis*-retinol to 11-*cis*-retinyl esters by multifunctional *O*-acyltransferase[85], creating Müller cell retinoid stores. CRALBP transfers 11-*cis*-retinol out of Müller cells[86], and cones absorb the retinol for further processing. It is known that the necessary oxidation to 11-*cis*-retinal occurs within cone outer segments, and RDH14 may provide this service[87]. A cone RDH also seems capable of oxidizing 11-*cis*-retinol without NAD or NADP, and can instead reduce all-*trans*-retinal in the process. This cone-specific retinoid cycle can produce 11-*cis*-retinal as much as 240 times faster than the classical retinoid cycle, hastening cone dark adaptation and improving responses to rapid scene changes in high-light conditions[73].

2.4 Signal transmission to the brain

Rods and cones signal photon capture through hyperpolarization. This relatively simple and intuitive signal in cones gets combined with those of other photoreceptors almost immediately by way interphotoreceptor gap junctions and horizontal cell synapses. By the point where one considers neurotransmitter release, the cone signal is not as much one reflecting light intensity as it is contrast. Photoreceptor signals branch out to influence many targets, and the significance of neuron activity becomes increasingly abstract with each subsequent synapse traversed. Eventually, the response of a neuron in the brain can be so refined as to signal the observation of a face or a potential collision. Such complex specificity relies on grandly sophisticated connections amongst series of neurons starting in the retina and stretching throughout the brain.

2.4.1 Retinal processing

Most neurons in the retina have gap junctions with other neurons of the same type, as well as some heterocellular gap junctions[88]. This tends to improve signal-to-noise ratios through averaging among similar neurons, while also enlarging receptive fields. Homocellular gap junctions are even seen with foveal cones[89], which bear the greatest demand for spatial specificity. At night, rod-cone gap junction conductivity also increases to permit rod signaling through cones[90]. The intracellular responses of rods and cones are thus thoroughly interconnected with those of their neighbors, and activity in photoreceptor spherules and pedicles reflect more than just the activity of their respective outer segments.

Aside from each other, rods and cones are also modulated by interplexiform cells and horizontal cells. Interplexiform cells receive input in the inner plexiform layer from bipolar cells and amacrine cells, and provide output to the outer plexiform layer. Interplexiform cells are not well understood, but seem to affect photoreceptor-horizontal cell-bipolar cell synapses through dopaminergic and glycinergic synapses. Dopaminergic innervation can decouple horizontal cells and alter the reception of glutamate by bipolar cells[88]. Glycinergic innervation seems to increase synaptic gain by photoreceptors and bipolar cells[91].

Horizontal cells receive excitatory input from rods and cones, and in turn impose lateral negative feedback on photoreceptor synapses. Humans have at least 2 types of horizontal cells[92], with some debate over a third type[93]. H1-type horizontal cells bear axons and associate with numerous rods and cones. Rods specifically innervate H1 axons, while cones innervate H1 dendrites. H1 cells connect to many medium- and long-wavelength cones (M and L cones), but relatively few short-wavelength cones (S cones). H2 horizontal cells also have axons, but instead

interconnect many S cones, with few M and L cones[94, 95]. Horizontal cells pool their inputs, with contributions from neighboring horizontal cells through gap junctions, and modulate the output of all their connecting photoreceptors.

Horizontal cell feedback is mediated by either ephaptic influence or synapse acidification. The ephaptic mechanism is invoked when the horizontal cell is hyperpolarized by a drop in photoreceptor input. Horizontal cell hyperpolarization draws cations into the horizontal cell dendrites from the extracellular space through hemichannels. Electrical resistance is high within the photoreceptor synaptic cleft, so the removal of extracellular cations produces a noticeable change in transmembrane potentials. The cone transmembrane potential rises to a less negative state as result, although the intracellular ion content is not immediately affected. Voltage-gated calcium channels in the cone respond by opening, which somewhat depolarizes the cone terminal and promotes more glutamate release from the cone. Hyperpolarization of cones and horizontal cells thereby leads to the effective depolarization of cones connected to the horizontal cell[95, 96].

Synapse acidification arises as a consequence of horizontal cell depolarization. When photoreceptors release glutamate, some protons are co-released from vesicles into the synaptic cleft[97]. Horizontal cells depolarize in response to glutamate release, which induces them to also release protons into the synapse. These protons flow out from numerous paths, such as horizontal cell hemichannels and outward-pumping vacuolar-type H^+-ATPase[98]. The rush of protons into the synapse exceeds the synapse's pH buffer capacity and increases the extracellular acidity. The excess protons bind to the exterior of photoreceptor voltage-gated calcium channels, reducing their conductance. The reduced flow of Ca^{2+} into the photoreceptor

terminal limits further release of glutamate. This negative feedback path causes photoreceptor and horizontal cell depolarization to lead to reduced photoreceptor glutamate release.

Horizontal cell negative feedback was once hypothesized to be mediated through GABAergic inhibition and disinhibition[99]. This has appeared progressively less likely, however, as the direct effects of horizontal cell GABA on cones are limited, and are not consistent with chloride-channel inhibition in primates[100]. Instead, GABA from horizontal cells seems to have a modulatory effect on lateral feedback. Horizontal cells do release GABA in response to photoreceptor input, but the main receptor for this GABA output appears to be on the horizontal cell itself. The horizontal cell GABA receptors open as channels for Cl^- and HCO_3^- ions[101]. Upon horizontal cell depolarization, the GABAergic channels would serve to shunt cation currents related to the ephaptic feedback mechanism[102]. Flow of HCO_3^- through GABA receptors may add to or counteract proton-mediated pH changes, depending on HCO_3^- concentrations and transmembrane potentials. HCO_3^- flux into the horizontal cell would further acidify the synapse, whereas HCO_3^- efflux from the horizontal cell would alkalize the synapse and disinhibit calcium channels[103].

Lateral feedback provided by horizontal cells contributes to center-surround antagonistic receptive fields[104]. Initially observed in retinal ganglion cells, center-surround receptive fields are present when activity related to a stimulus in the center of a neuron's receptive field is counteracted by responses to stimuli in the surrounding periphery of its receptive field. Photoreceptors intrinsically hyperpolarize to light, and depolarize when light is turned off. The part of the visual field stimulating a photoreceptor with photons, or doing the same to any other connected photoreceptors that share responses through gap junctions, would be considered the center of the photoreceptor's receptive field. The photoreceptor would specifically have an OFF

center, because it depolarizes and releases neurotransmitter in response to light being turned off.

The surround of the photoreceptor receptive field is provided by its horizontal cells. Photoreceptors not well connected to the central photoreceptor through gap junctions may still be connected through a horizontal cell. These more distant photoreceptors would detect stimuli in the central photoreceptor's receptive field surround. Just as in the center, darkness would depolarize the surround photoreceptors, and that depolarization would depolarize the common horizontal cell. Photoreceptors in the center would stimulate the horizontal cell, but that stimulation would not drive the horizontal cell enough to counteract depolarizing effects of the dark current. When enough of the surround is active, however, the horizontal cell can inhibit the central photoreceptor's glutamate release by acidifying the synapse. In contrast, light in the surround can reduce input to the horizontal cell. The horizontal cell hyperpolarization would promote more glutamate release from the central photoreceptor through ephaptic cation intake. Lights on in the surround would thus cause the central photoreceptor to release more neurotransmitter, giving it an OFF-center, ON-surround receptive field.

In uniform light conditions, the conflicting excitatory and inhibitory contributions of the center and surround would mostly cancel each other out. The photoreceptor would release glutamate at some baseline level, neither fully excited nor fully inhibited. Maximal neurotransmitter release would be observed when the center is in darkness, but the surround is bathed in light. The central photoreceptor would be driven by its own and neighboring dark currents, as well as by horizontal cell ephaptic cation intake. Conversely, bathing the receptive field center in light while leaving the surround in darkness would maximally inhibit the photoreceptor's glutamate release. Not only would the lack of dark current cause the photoreceptor transmembrane

potential to rise, but horizontals would further reduce calcium channel conductance through synapse acidification.

Across the outer plexiform layer, bipolar cells serve as the primary recipients of photoreceptor glutamate. Bipolar cells are so named because they have 2 projections: one receiving input from the outer plexiform layer, the other providing output to the inner plexiform layer. There are at least 10 types of bipolar cells in the primate retina[105]. One type receives rod input, while the others receive cone inputs, with little direct mixing of rod and cone inputs. Cone bipolar cells with flat post-synaptic regions can express AMPA or kainate receptors, and are therefore depolarized by cone glutamate release. These bipolar cells preserve the photoreceptor receptive field organization of OFF-center, ON-surround. Other cone bipolar cells have dendritic extensions that reach toward the cone, forming an invaginating post-synaptic region. Invaginating bipolar cells express mGluR6, which induces hyperpolarization in response to glutamate. These bipolar cells thus invert the signal of the cone, and have ON-center, OFF-surround receptive fields.

Midget bipolar cells receive input from only a few M or L cones, which preserves spatial specificity important for high-acuity vision. OFF-center midget bipolar cells also appear to have significant input from S cones, although that input is still weaker than M- or L-cone input[106]. In the fovea, midget bipolar cells receive input from only one M or L cone, and output to only one midget RGC. This creates some color opponency between the nearly pure input from a single cone, more importantly a single spectral tuning, for the center, with a mixed M and L cone surround[107]. Reliable color opponency of midget bipolar cells is lost farther in the periphery as they receive direct input from more nearby cones, likely with some mix of M and L cones, and

some small S-cone contribution. Spatial specificity is also somewhat reduced. Midget bipolar cells can be ON- or OFF-center, depending on glutamate receptor expression[108].

Color opponency is also a presumed feature of S-cone-selective bipolar cells. Midget bipolar cells achieve color opponency by having direct contact with only one cone, and necessarily only one type of cone, but S-cone-selective bipolar cells receive input from multiple cones. The dendrites of S-cone bipolar cells, however, travel longer distances than those of midget bipolar cells and selectively form synapses with S cones. S-cone bipolar cells express only metabotropic glutamate receptors, and therefore inherit short-wavelength ON responses from S cones. H2 horizontal cells confer an M- and L- cone-driven receptive field surround to S cones, which is similarly inverted by the S-cone bipolar cell. S-cone bipolar cells therefore have S-ON-center, M+L-OFF-surround receptive fields. These receptive fields are also described as having blue-ON centers and yellow-OFF surrounds[109]. The S-cone bipolar cells are thus an essential component for adding short-wavelength information to color discrimination and perception.

Many bipolar cell types synapse with cones indiscriminately within their dendritic field, regardless of cone spectral tuning. Such bipolar cells are categorized as diffuse bipolar cells. Dendritic fields of diffuse bipolar cells extend farther than those of midget bipolar cells, and can synapse with as few as 4 or as many as 15 cones. Lack of cone-type specificity in synapses implies that diffuse bipolar cell signals typically have no chromatic information. The longer reach of their dendrites also reduces their spatial specificity, limiting their usefulness for higher-acuity perception. Diffuse bipolar cells instead generally convey luminance information for their receptive fields, and may be able to respond to light-level changes more quickly than other bipolar cells. Because M- and L-cone density is about 20-times that of S cones in primates, the sensitivity of diffuse bipolar cells is shifted toward longer wavelengths. Like midget bipolar cells,

diffuse bipolar cells can express ionotropic or metabotropic glutamate receptors to have either ON or OFF receptive-field centers[108].

Midget, S-cone, and diffuse bipolar cells only receive input from cones. Primate retinas, however, are dominated by rods. Rod bipolar cells, the only bipolar cells that contact rods, are therefore the most abundant type of bipolar cell in the primate retina. These bipolar cells use only metabotropic glutamate receptors, and therefore have ON-center receptive fields. Rod bipolar cells do not appear to have distinct receptive-field surrounds. Horizontal cell release of GABA onto rod bipolar cells, however, does modulate bipolar cell activity, enhancing the effect of the center when there is contrast with surrounding light levels[110]. While cone bipolar cells synapse directly with RGCs, rod bipolar cells pass their signals through AII amacrine cells. AII amacrine cells depolarize following rod bipolar cell activation, and react by depolarizing OFF-center cone bipolar cells and hyperpolarizing ON-center cone bipolar cells. Rod bipolar cells thus utilize cone signal pathways to relay rod signals to the brain[105, 108].

Bipolar cells have traditionally been considered to operate based on graded potentials[111]. Under such a paradigm, the bipolar cell can operate over a continuous range of activity and neurotransmitter release is directly correlated with stimulation. Passive diffusion of ions through the bipolar cell, however, can impose limits on the frequency of signal changes from the photoreceptors to the RGCs. Midget RGCs, which receive their input from midget bipolar cells, have high spatial specificity, but relative low temporal specificity. Parasol RGCs conversely have low spatial resolution, but high temporal resolution. While midget bipolar cells do appear to operate using grade potentials, certain diffuse bipolar cells that innervate parasol RGCs have voltage-gated cation channels that allow them to form action potentials. Bipolar cells with

action potentials can generate signals more rapidly, and may be responsible for the faster timing profile of parasol RGCs.

While only some bipolar cells possess the capacity for generating action potentials, all RGCs use action potentials to conduct signals along the optic nerve into the brain. Four types of RGCs have been well-studied in primates: midget RGCs belonging to the parvocellular pathway, parasol RGCs of the magnocellular pathway, small bistratified RGCs of the koniocellular pathway, and ipRGCs that are important for non-image-forming vision. As mentioned above, midget RGCs receive input from midget bipolar cells. The midget RGCs inherit the receptive field dynamics of their bipolar cells, exhibiting high spatial specificity with M- and L-wavelength opponency. The signals of foveal midget RGCs are driven by only one cone and one bipolar cell[8], but bipolar cell-RGC ratios increase with eccentricity. At any given eccentricity, however, the dendritic range of midget RGCs is consistently shorter than that of parasol RGCs. Synaptic contacts with bipolar cells in the inner plexiform layer differ based on receptive field type, for not only midget but most RGCs. ON-center RGCs synapse with ON-center bipolar cells within the inner half of the inner plexiform layer, the half closer to the RGC somata and vitreous, and OFF-center RGCs and bipolar cells connect in the outer half of the inner plexiform layer. ON and OFF midget RGCs each constitute approximately 26% of the primate RGC population, and combined make up about 70% of RGCs in the fovea, where they are most densely present[5].

In comparison to the 26% of primate RGCs that belong to each midget RGC type, ON-center and OFF-center parasol RGCs each only comprise about 8% of the primate RGC population. As both midget and parasol RGCs cover the entire retina, this reflects the lower density of parasol somata and the wider reach of individual parasol dendritic fields. Parasol RGCs receive input from more bipolar cells than midget RGCs, and therefore their receptive field centers are

derived from more M and L cones. The likelihood of significant imbalance between M and L cones contributing the center is far less than that for midget RGCs, so parasol RGCs display roughly equal sensitivity to M and L spectral tunings in both the center and surround. Thus, instead of having receptive fields resembling M-ON center and M+L-OFF surround, parasol RGC receptive fields are simply ON-center, OFF-surround or OFF-center, ON-surround. Because they lack color opponency, and are therefore not affected difference between M and L cone responses, parasol RGCs reactive more robustly to achromatic contrast than midget RGCs. As discussed above, parasol RGCs also respond with greater temporal specificity, allowing them to appropriately respond to higher rates of flicker without saturating[5].

S-cone bipolar cells project to small bistratified RGCs. Small bistratified RGCs are similar to parasol RGCs in number and dendritic field extent. There is only one form of small bistratified RGC, and it constitutes about 6% of the primate RGC population[5]. Small bistratified dendrites cover areas about 10% wider than those of parasol RGCs, but their dendritic field sizes are considerably more variable. Dendrites receive input from two strata of the inner plexiform layer, the morphology for which these RGCs are named. S-cone bipolar cells with ON centers synapse with small bistratified RGCs in the inner half of the inner plexiform layer, and OFF-center diffuse bipolar cells form contacts in the outer half of the inner plexiform layer. It was originally believed that small bistratified RGCs had S-ON, M+L-OFF receptive fields without any center-surround structure. Under that conception, diffuse bipolar cells were assumed to provide the M+L-OFF signal to the small bistratified RGCs. More recently, however, the M+L-OFF component of the receptive field has been shown to have a larger radius than that of the S-ON component. The timing of the M+L-OFF response is also more consistent with primary contributions of horizontal cells acting on S-cones, rather than primary dependence on diffuse bipolar cells. Small

bistratified RGCs can thus be considered to inherit S-ON-center, M+L-OFF-surround receptive fields from S cones and S-cone bipolar cells[112].

The ipRGCs differ substantially from other RGC classes, most pronouncedly because of their ability to independently depolarize in response to light. These RGCs contain the photopigment melanopsin, which binds the same chromophore as rod and cone opsins, 11-*cis*-retinal. Unlike photopigments of rods and cones, however, melanopsin triggers the opening of cation channels through a phosphoinositide signaling cascade. Melanopsin therefore causes ipRGCs, essentially melanopsin photoreceptors, to depolarize in response to light, rather than hyperpolarize like rods and cones. After a photon isomerizes 11-*cis*-retinal to all-*trans*-retinal, another photon can reisomerize retinaldehyde to the 11-*cis* configuration while it remains bound to melanopsin[113]. The depolarization of ipRGCs caused by melanopsin is slow and consistent, and provides information on the overall brightness of the environment. ipRGCs have also been known as giant monostratified RGCs, as they possess extremely large dendritic trees. The dendritic fields are so broad that the neurons only possess trivial spatial specificity. The dendrites cover most of the retina, but ipRGCs only constitute about 0.2% of the RGC population[5]. In addition to their intrinsic photoreception, ipRGCs receive supplemental input from rods and cones. The cone-derived ipRGC receptive fields are M+L-ON and S-OFF, although the synaptic pathways establishing such responses are not known. Melanopsin, however, has a peak absorption wavelength of 480 nm, which is very close to that of S cones. Although the M+L-S opponency can be measured in laboratory environments, the melanopsin-driven depolarization in response to S wavelengths normally dominates and masks the cone-driven S-OFF response[114].

2.4.2 Retinothalamic pathway

RGC axons converge upon the optic disc and exit the eye through the lamina cribrosa as the optic nerve. This term is a misnomer, as RGCs are embryologically and physiologically part of the of the central nervous system, and nerves are axon bundles in the peripheral nervous system. The optic nerve is actually a tract, myelinated by oligodendrocytes beyond the lamina cribrosa. The optic nerve contains all visual information output for each eye, and thus each eye's entire monocular visual field. Optic nerves from each eye pass through the optic canal and converge on the inferomedial surface of the brain to form the optic chiasm.

Within the chiasm, RGC axons from the nasal, or medial, halves of each eye decussate to join axons from the temporal or lateral half of the opposite eye. The right visual field is captured by nasal right retina and temporal left retina, while the left visual field falls upon nasal left retina and temporal right retina. The decussation of nasal RGC axons in the chiasm places representations of each half of the visual field within the contralateral side of the brain. Animals with more lateral-facing eyes feature larger proportions of optic nerve axons that decussate through the chiasm. Axons projecting posteriorly from each side of the chiasm form the optic tracts. Although the optic nerve is technically a tract, these postchiasmatic bundles are the only RGC axon sections that are referred to as optic tracts. Rather than representing an entire monocular visual field, as each optic nerve does, each optic tract represents the contralateral visual field from both eyes.

Axons of ipRGCs are known to project to the suprachiasmatic nucleus and olivary pretectal nucleus, respectively for entrainment of circadian rhythms and driving the pupillary light reflex. Many RGCs provide input to the superior colliculus to guide eye movements. Most RGCs,

however, terminate in the lateral geniculate nucleus (LGN) of the thalamus. The retinothalamic pathway provides stimulation for conscious visual perception. Like the optic tracts that provide retinal input to thalamus, each LGN possesses representations of the contralateral visual field. More precisely, the 6 main layers of primate LGN each contain a monocular hemifield driven by either midget or parasol RGCs from the contralateral or ipsilateral eye. Thalamocortical neurons, those responsible for transmitting visual signals to cerebral cortex, receive only about 10% of their synaptic inputs from RGCs, but RGC connections are the only ones that can independently elicit thalamocortical responses. All other connections, such as those from interneurons, brain stem, or cortex, provide modulatory influences[115].

The 2 ventral-most main layers of LGN, layers 1 and 2, are the magnocellular layers. Cells of these layers are large and receive input from parasol RGCs. Layer 1, numbered ventral to dorsal, is driven by ipsilateral temporal retina, while layer 2 is innervated by contralateral nasal retina. Both layers receive mixtures of ON- and OFF-center parasol RGC contacts. Like their connected RGCs, neurons of the magnocellular layers display low spatial specificity, high temporal specificity, and are highly sensitive to achromatic contrast. Magnocellular neurons are also particularly sensitive to modulation by attention, through innervation from cortex and the brain stem. Retinotopy is maintained in magnocellular layers, as well as the other main LGN layers, such that columns perpendicular to layer borders represent the same region of the visual field in each layer[115]. The fovea is represented at the posterior pole, and the representation of central vision is magnified such that the volume dedicated to central 15° occupies approximately 80% of total LGN volume[116].

About 10% of LGN thalamocortical neurons occupy the magnocellular layers. In contrast, about 80% of LGN thalamocortical neurons exist in the parvocellular layers, layers 3–6. Midget RGCs

from contralateral nasal retina stimulate parvocellular neurons in layers 3 and 5, while midget RGCs of ipsilateral temporal retina project to layers 4 and 6. Unlike the magnocellular layers, parvocellular layers grossly segregate RGC axons based on receptive-field polarity. Each pair of layers responding to one eye has one layer dominated by ON-center midget RGCs, and most OFF-center midget RGCs connect to the other layer. ON-center midget RGCs are primarily received by the dorsal layers, layers 5 and 6, and the more ventral layers 3 and 4 synapse mostly with OFF-center midget RGCs. Parvocellular neurons are smaller than magnocellular neurons, and inherit the color opponency of midget RGCs. Parvocellular neurons exhibit high spatial specificity, poor temporal specificity with high frequency stimuli, and longer latencies than magnocellular neurons.

Ventral to each of the main 6 LGN layers lie the koniocellular layers. Koniocellular layers are thin and their constituent cells are smaller than those of the parvocellular layers. Koniocellular layers make up roughly the last 10% of thalamocortical neurons in LGN. Many koniocellular neurons receive input from small bistratified RGCs, particularly in the central and dorsal layers underlying the parvocellular layers. Koniocellular neurons thus display S-ON, M+L-OFF color opponency. Interestingly, a small number of other koniocellular neurons display S-OFF, M+L-ON responses, but the nature of these responses is not yet known. Unlike other LGN neurons, those in the ventral koniocellular layers beneath the magnocellular layers receive inputs from superior colliculus. While magnocellular and parvocellular neurons are mostly monocularly driven, many more koniocellular neurons display binocular sensitivity, which may be derived from collicular connections. Some koniocellular neurons also exhibit substantially more pronounced orientation selectivity than other LGN neurons. Koniocellular receptive field sizes are much larger, and response latencies longer, than those of magnocellular and parvocellular neurons.

Contrast sensitivity and operational temporal frequencies lie between those of magnocellular and parvocellular neurons.

2.4.3 Visual cortex

LGN thalamocortical neurons collectively exit the thalamus as the optic radiation on each side of the brain. The optic radiation forms part of the posterior limb of the internal capsule, and subsequently fans around the lateral ventricle, blanketing and extending past the lateral wall of the ventricle, ultimately to reach primary visual cortex in the occipital pole. The more lateral and inferior fibers of the optic radiation carry signals representing the superior half of the contralateral visual field. These fibers course anteriorly and laterally along the roof of the ventricle's inferior horn, some fibers proceeding beyond the anterior extent of the horn, before turning to travel posteriorly through the temporal lobe and along the ventricle's lateral wall. These lateral radiation fibers are known as Meyer's loop. Fibers representing the central contralateral visual field occupy the central majority of the optic radiation, and inferior visual field information is carried by fibers that travel more medially and superiorly. The superior radiation fibers follow the most direct course, passing through parietal lobe, laterally around the ventricle, and converge with the rest of the radiation upon primary visual cortex[117, 118]. The radiation fibers with paths through the parietal lobe are sometimes known as Baum's loop.

Myelin surrounding optic radiation axons creates a visible white line that passes through the gray matter along the borders of the calcarine sulcus of occipital cortex. This portion of occipital cortex is the target of the optic radiation, primary visual cortex (V1). Because of the line formed by the radiation axons, V1 is also known as striate cortex. Inferior radiation axons synapse within the lingual gyrus at the inferior edge the calcarine sulcus, while superior radiation axons

innervate the superior edge of the sulcus in the cuneus. V1 thus takes on a retinotopic map in which the inferior vertical meridian of the visual field is represented along the superior lip of the calcarine sulcus, the horizontal meridian occupies the bottom of the sulcus, and the inferior lip of the sulcus corresponds to the superior vertical meridian of the visual field. Thalamocortical axons with foveal representations synapse at posterior end of the sulcus, and projections for greater eccentricities terminate more anteriorly along the sulcus. Lines from lip to lip of the sulcus at a single distance from the occipital pole thus roughly correspond to semicircles in the visual field.

Cerebral cortex, including V1, consists of approximately 6 layers of interconnected neural networks. Layer 1 is the most superficial layer, and layer 6 is the farthest from the cortical surface. Most output from thalamus enters cortex by way of cortical layer 4. Consistently, most of the optic radiation axons synapse with neurons in layer 4C of V1. The more superficial half of layer 4C, layer 4Cα, is innervated by magnocellular neurons of LGN. The deeper layer 4Cβ receives input from LGN parvocellular neurons. Most other layers receive some connections from LGN, as well. Both magnocellular and parvocellular neurons send collaterals to layer 6 of V1. Parvocellular neurons innervate layer 4A. Koniocellular neurons do not innervate layer 4 or any deeper layers of V1, but do send axons to layers 1–3. Koniocellular neurons underlying the LGN magnocellular layers connect with V1 layer 1 and superficial layer 3. Koniocellular neurons situated under parvocellular layers project to the cytochrome oxidase blob regions of V1 layers 2 and 3[115]. Layers 4B and 5 of V1 are not innervated by LGN axons[119].

Columns in V1, extending down from the surface of cortex perpendicularly through the layers, contain neurons that all share similar biases for visual stimulation. Two prominent biases present in V1 columns are ocular dominance and orientation selectivity. Ocular dominance

presents as alternating stripes meandering along V1. Neurons in layer 4C are strictly monocular, only being driven by the contralateral or ipsilateral eye. Borders between stripes of ocular dominance columns thus have pronounced response differences in layer 4C. In more superficial layers, however, transitions between stripes are more gradual and some neurons have binocular receptive fields. Superimposed on ocular dominance stripes are strips of orientation selectivity. Neurons within orientation columns share biases toward stimuli, such as bars or edges, oriented at particular angles. Orientation selectivity changes gradually across V1, such that adjacent columns have similar orientation biases. Biases change and progress in circular patterns, such that strips of similar orientation columns form pinwheel patterns about points of convergence where all orientations are represented within small areas. Pinwheel centers tend to sit along the middle of ocular dominance stripes[119].

After initial processing and filtering, V1 neurons pass visual information along to other brain areas. Layers 2, 3, and 4B project to higher visual cortex. Layers 5 and 6 send axons to subcortical targets. LGN also makes direct connections to many visual cortex regions outside V1, primarily via koniocellular neurons[115]. Visual cortex beyond V1 is labeled roughly by sequential visual processing, and neurons of higher visual areas have increasingly abstract representations of stimuli as synaptic distance from V1 grows. Each visual area possesses its own map of visual space, and retinotopy is preserved in the earlier cortical regions. Secondary visual cortex (V2) is also labelled as prestriate cortex, and it maintains much of the tuning present in V1, with additional sensitivities such as binocular disparity. Areas farther from V1, such as V3, V4, and the middle temporal area (MT or V5), are considered extrastriate cortex. Extrastriate cortex is often broadly divided into dorsal and ventral streams, such that each stream processes visual information with a distinct contextual emphasis. The dorsal or parietal stream handles visual information relative to representations of self or one's environment, often termed the "where"

or “how” stream. Rather than just being sensitive to light in a particular part of the visual field, a neuron in a higher area of the dorsal stream might be most sensitive to visual stimuli than appear to be approaching the viewer. In contrast, the ventral or temporal stream is considered the “what” stream, and is more concerned with the identity or conceptual classification of visual stimuli. A neuron far along the ventral stream could particularly sensitive to when a face or a car is present within the visual field. As visual information moves through the dorsal and ventral streams and into higher association areas, neural responses to light evolve into a more complex understanding of the world that is used to guide motor commands, decision making, and abstract thinking.

2.5 Light localization

Humans rely on various information sources to localize light in space. Consider a single point source of light being viewed monocularly. Photons from that source would be absorbed by photoreceptors within a small area of the viewer’s retina. That location on the retina would determine where in the viewer’s visual field the source lies. To assign the source a location in space, the viewer must relate the visual field to the viewer’s environment and/or body position. Mapping the visual field to egocentric or allocentric coordinates requires additional information, which may be provided by the rest of the visual scene or input to other sensory modalities. When perceptual errors generate conflict among information sources, people can learn to adapt how the visual field information is used in combination with other senses and motor commands.

2.5.1 Perception of gaze direction

Determining the relationship between the visual field and the environment requires contextual information. This information can be contained within the visual field itself, if the field contains

sufficient detail to deduce the environment's relative orientation. Such an informative visual field is considered structured[120], and would not require any supplementary information to complete orientation-based tasks. Movies and video games often provide truncated structured visual fields, from which the viewer can deduce the orientation of the captured or virtual environment, while the image is completely independent of the viewer.

Some context may also be applied by relating the visual field to one's own body. Such a relationship can provide context from other sensory modalities, and would be critical for visual coordination tasks. Abstractly, one needs information about both the head and the viewing eye to relate the visual field to one's body position. The viewer must know the head's position and orientation relative to the body, and the eye's orientation within the head. The egocentric direction of gaze can then be resolved mathematically and thus place the visual field and any visual stimuli into the body's frame of reference.

More concretely, a viewer determines his or her relative direction of gaze through combinations of proprioception, corollary discharge, and vestibular inputs. Proprioception holds information on how the body is configured and the positions of body parts relative to each other. Proprioceptive information typically comes from muscle spindles, which signal changes in muscle lengths. Golgi tendon organs and cutaneous mechanoreceptors also provide proprioceptive information at extreme joint positions[121]. Extraocular muscles of mammals, however, do not consistently possess spindles or Golgi tendon organs. Instead, nervous end organs called palisade endings, or innervated myotendinous cylinders, found at the myotendinous junctions of extraocular muscle fibers, are thought to provide proprioception about the eyes.

The precise function of palisade endings, while assumed to involve proprioception, is not entirely clear. Palisade endings are cholinergic, and have origins within motor nuclei for the extraocular muscles. The structure and position of synapses, however, are more compatible with sensory than motor neurons. Both proprioceptive and motor functions thus seem possible, and a great deal of further investigation is required. In either case, palisade endings appear to be particularly helpful with fine eye movements used by mammals with forward-facing eyes. In such frontal-eyed mammals, palisade endings are found in greater densities than in lateral-eyed mammals, particularly along the medial rectus muscles. Frontal-eyed mammals make convergence and smooth pursuit motions more than lateral-eyed mammals, and palisade endings may be important for using those eye movements[122].

Corollary discharge, a higher-order variation of the related efference copy concept, modifies brain activity, particularly in sensory areas, in accordance with motor commands. Modifications can manifest as changes in sensory processing, such as distinguishing self-motion from environment motion, or as the basis for sensorimotor learning through the cerebellum[123]. For example, information that the head or eyes are moving rightward can be used to suppress the perception that a light in the visual field is moving to the left, as it is actually stationary. Similarly, a rightward eye motion combined with a stimulus that does not move in the visual field will be seen as the stimulus moving along with the eye. Many brain areas are known to alter receptive fields in anticipation of saccades, including V2, V3a, V4, the parietal reach area, superior colliculi, and frontal eye fields[124]. One particular pathway was mapped for oculomotor corollary discharge from the superior colliculus, through the mediodorsal thalamic nucleus, to the frontal eye field. Inactivating mediodorsal thalamus did not interfere with saccade production, but did prevent normal remapping of frontal eye field representations in anticipation of the saccade[125].

Vestibular sensation does not directly provide information on eye orientation within the head, but can be very important for determining head orientation. Embedded within the otolithic membranes of the utricle and saccule, calcium carbonate otoconia apply forces to hair cells corresponding to gravity and linear acceleration. The utricle is oriented to mostly respond to translations within the horizontal plane, and the saccule is sensitive to forces pushing forward and back or up and down. Consider sitting in a plane as it is gently ascending. Looking down the aisle would not provide much, if any, visual indication that the plane is sloped upward. Proprioceptive input from your neck would similarly lack that information, as looking upward with respect to your body would still direct your gaze to the ceiling. Pressure from your seat would be close enough to normal to suggest the plane might be level. Vestibular input from your otolith organs, mostly your saccules, would nevertheless inform your brain that aligning your head to look straight along the aisle, forward with respect to your body, is in fact looking upward with respect to gravity. That information allows you to ultimately perceive that the aisle is sloped and that the plane is aimed upward.

While the otolith organs provide information on gravity and linear acceleration, the semicircular canals supply the vestibular sense of angular velocity. As the head turns, endolymph fluid within the canals exhibits inertial resistance. Pressure from the endolymph pushes the cupula in a canal's ampulla in the opposite direction of the rotation, which in turn deflects hair cells attached to the cupula. While not obviously changing the context of visual information, stimulation of the semicircular canals is a major component of the vestibulo-ocular reflex. This reflex, with combined input from the otolith organs and semicircular canals, causes the eyes to rotate in the opposite direction of any head rotation, which typically aids in maintaining fixation through head movements. The otolith organs do trigger reflexive eye rotations in response to head translations, when the semicircular canals would be mostly quiet, but the gain of the reflex

is then heavily dependent on the distance from the focused visual target[126]. When eye orientation reaches an extreme and head movement continues, the eyes saccade back to center, which is considered physiological nystagmus.

The perception of head orientation, and interpreting visual stimuli within the context of head orientation, can depend on each of the above sensory inputs. Muscle spindles in the neck inform the brain of the orientation of the head with respect to the body. Corollary discharge for active head turns combines with proprioception to more accurately compensate for changes in orientation. Input from the semicircular canals, together with any oculomotor corollary discharge, can further specify the relationship between moving percepts and self-motion. When asking subjects to indicate their apparent direction of motion in a visual simulation, Crowell et al.[127] found that isolated proprioceptive or vestibular feedback was inadequate for accurate judgement of the visual stimuli. Subjects seemed to weigh feedback modalities differently, and performance improvements for combinations of 2 feedback modalities were inconsistent across subjects. Maximal performance was achieved when all 3 sources of feedback were available.

Along with when making active head turns, Crowell et al. saw that subjects displayed accurate performance when they made active eye movements. Because vestibular feedback does not provide any information on the eyes beyond the orientation of the head, accounting for eye movements relies solely on proprioception from extraocular muscles and oculomotor corollary discharge. Bridgeman and Stark[128] attempted to separate these feedback sources by having subjects gently press on either their uncovered viewing eye or their occluded eye. Pressing on the viewing eye triggered compensatory muscle activity, and thus invoked corollary discharge, without changing the viewing eye's orientation. The occluded eye, however, did rotate in response to the contralateral eye press, and the combined proprioceptive signal for the

cyclopean eye was altered. Pressing on the occluded eye did not trigger muscle activity, but did alter the combined proprioceptive signal. While pressing on an eye, subjects were asked to describe the location of or point to a visual target in an unstructured visual field. By comparing degrees of eye rotation to localization responses, the investigators determined that proprioception altered visual perception with a gain of close to ¼. Corollary discharge, after removing the contribution of proprioception, modified perception with a gain of about ⅝. The authors therefore concluded that corollary discharge dominates proprioception in accounting for eye movements.

Interestingly, Bridgeman and Stark's calculations imply that extraretinal feedback only modifies perception by a total of ⅞ of the actual eye movement. This calculated underestimation of eye eccentricity is remarkably consistent with earlier work by Morgan[129] and later experiments by Lewald and Ehrenstein[130]. Both of these other studies asked subjects to fixate on LEDs at various eccentricities and aim a hidden pointer as accurately as possible at the fixation target. Bridgeman and Stark calculated an eccentricity underestimation of 0.13 in Morgan's data, and Lewald and Ehrenstein observed an underestimation factor of 0.12. Additionally, Lewald and Ehrenstein found that perception of head-turn eccentricity was influenced by eccentric eye orientations, such that head-turn eccentricity, when the head was aimed forward, was perceived to be approximately 0.23 of eye eccentricity. While extraretinal feedback on eye and head orientations provides enough information to map the visual field into egocentric space sufficiently for daily life, that information is incomplete, and mapping accuracy deteriorates with eye eccentricity.

2.5.2 Adaptation to visual misalignments

Various studies have examined the effects of inconsistencies between visual input and extraretinal localization cues. Normal visual localization assumes that light enters the eye without any refraction prior to interacting with the cornea. When the path of light is altered before it reaches the eye, the brain incorrectly localizes the light as though no refraction had occurred. Mislocalization such as this transpires when one looks at something submerged underwater from above the water's surface. Because of the refraction at the water-air interface, the depth and location of objects can be inaccurately perceived. Lenses placed in front a viewer's eyes can introduce similar dissonance between percepts and true object locations for all visual input. Depending on how the viewer interacts with the visual scene, one can quite rapidly adapt to succeed at visuomotor tasks even with severe alterations to visual input.

In 1896, George Stratton investigated the importance of having properly oriented visual input by wearing inverting lenses over his eyes[131]. A series of lenses inverted a 45° section of the visual scene just before the light reached his eyes, with the rest of the visual field obscured by a mask. One eye was always occluded to avoid the need for variably converging lens axes. Initially, Stratton reported that the world appeared upside-down, and visuomotor guidance was completely unreliable. Visual changes associated with head movements occasionally made Stratton nauseous. Coordination, however, improved over time, as did his ability to imagine the world about him in terms of the inverted visual field. Stratton originally only wore the lenses for 21.5 hours spread across 3 days, with a blindfold blocking vision otherwise. He did not report any aftereffects following removal of the lenses at the end of the experiment, but subsequently wore the lenses for 87 hours over 8 days in a longer experiment[132]. While the world without the lenses did not appear upside down after the second experiment, Stratton found himself making

marked visuomotor errors corresponding to what would have been appropriate when viewing through the inverting lenses. Normal coordination mostly returned by the second day after completion of the experiment.

Many researchers, including Hermann von Helmholtz[133], Margaret Wooster[134], and James Gibson[135] similarly investigated visual perturbations by examining the effects of prism glasses on viewers' behaviors. Prism glasses created less drastic changes to visual input than those studied by Stratton, and were limited to shifting vision along the horizontal axis. Subjects were able to view scenes binocularly with a wider field of view than afforded by Stratton's device. Gibson used optometrist's trial prisms with angles of 26°, creating rightward visual shifts of approximately 15°. Subjects wearing the prism glasses were initially reaching and pointing too far to the right, and walking into obstacles on their right sides. Because visual shifts above and below the center of the prisms were more dramatic than that at the center, vertical lines also appeared curved to the left. After about an hour with coordination exercises, visuomotor errors were absent or greatly diminished. Impressions of curvature were also reduced over the experiment's hour, and vertical lines appeared to curve rightward by the same extent as the adaptation after prisms were removed.

Wooster noted that subjects did not improve visuomotor coordination unless reaching movements were made while fixating on an object. Held and Hein[136] further investigated this dependence in 1958. Subjects were asked to mark intersections of lines, shown on an oblique mirror, without being able to view their hands or pens. Subjects then viewed their hands through a glass prism for 3 minutes in 1 of 3 conditions: no motion, passive motion, or self-produced motion. In the first condition, the subject simply scanned the hand through the prism without moving it. In the motion conditions, the elbow was stationary and the forearm moved

left and right in tune with a metronome set to 1 Hz. Passive motion called for the subject to relax as the experimenter moved the hand, while self-produced motion was controlled by the subject. After viewing, the subject once again marked line intersections seen on an oblique mirror. The authors found that hand-eye coordination only changed when subjects moved their arms themselves, seen as a significant horizontal shift in markings. Refining Wooster's observation, they concluded that reafference is required for visuomotor adaptations to shifts induced by prisms. Without the conflict between corollary discharge, specifically that of the arm in their experiments, and visual input, no changes are made to account for shifts in visual perception.

While viewing through prisms without reafference fails to change visuomotor coordination, there is some suggestion that it can still alter eye orientation for straight-ahead fixation[137]. When subjects' gazes were maintained in lateral positions through prism glass viewing, but no pointing error feedback was provided, lateral shifts in straight-ahead fixation were observed after the experiment. Such changes in straight-ahead fixation have been observed in adaptation to prisms and considered an important part of the adaptation process[138]. Because fixation shifts can be observed without coordination adaption, however, more investigation may be required to define the relationship between visuomotor adaptation and changes in straight-ahead eye fixation.

2.6 Blindness

Failure in any part of the visual system can lead to some visual deficits. Depending on the location and severity of damage, deficits might go unnoticed or render an individual completely blind. Visual system deficits can be broadly categorized based on whether the damage is optical

or neurological in nature. Damage to the optical visual system, any part anterior to the retina, dominates globally as the more common cause of blindness. Optical problems with the eye are relatively easy to correct with surgical adjustments and replacement optics, but treatment can be difficult to obtain in developing parts of the world. Lens cataracts are by far the most common cause of blindness, and can be treated by removing the lens and implanting an artificial replacement. Corneal opacity, such as that caused by trachoma infections, can addressed by removing the causes of corneal irritation, keratectomy, or in more extreme cases, corneal transplants. Damage to the neural visual system, from the retina through the brain, is much more complicated, and often persists even in developed communities as forms of untreatable vision loss.

2.6.1 Photoreceptor degeneration

Neural activity in the visual system starts in the photoreceptors, and these neurons can be destroyed or incapacitated by various mechanisms. Because the inner and outer segments are dependent on the choroidal blood supply, interruption or separation from that blood supply can lead to the deterioration of the photoreceptor's transductive compartments. Photoreceptor synaptic terminals and cell bodies, at least partially sustained by inner retinal blood flow[139], can survive without choroidal vasculature, but lack any phototransduction capabilities. Alternatively, inherited diseases, such as retinitis pigmentosa, can cause critical failures specifically in photoreceptors or the RPE, eventually destroying most photosensitivity in the retina. Removal of photoreceptor function deprives the visual system of sensory input, and can lead to pathologic reorganization of the surviving retina.

The most direct and complete disruption of interaction between photoreceptors and the RPE is a retinal detachment. The neural retina is physically separated from the RPE, making RPE support of the retina difficult or impossible. Most retinal detachments occur following a tear in the retina, through which liquified vitreous humor can flow and push the retina away from the RPE. Such tears often occur as the vitreous body loses elasticity and progressively separates from the retina. Acute tension on the remaining vitreoretinal attachment, particularly along the superior-inferior equator of the eye, pulls the thin connected retina away from adjacent retina. The retina can also be pulled away from the RPE by scar tissue on its vitreal surface, or pushed off the RPE by the pressure of subretinal bleeding or edema. When treated with scleral buckling or vitreous surgery, the retina can be reattached to the RPE with much vision restored[140]. Postoperative visual acuities decline with detachment duration, and best results are achieved with treatment within 3 days of initial detachment[141, 142]. Left untreated, detachments can progress throughout the eye, and if detachment occurred in both eyes, can lead to complete vision loss[143, 144].

Age-related macular degeneration (AMD) is a common disease in which changes about the RPE in the macula cause damage to the nearby photoreceptors[145]. Age is the major risk factor for AMD, as suggested by its name, along with correlations to smoking, exercise, and dietary habits. Rod density in the macula declines with age, and the RPE accumulates retinoid-derived lipofuscin granules over time. Lipofuscin consists of A2E (N-retinylidene-N-retinylethanolamine), which itself is metabolized from A2PE-H_2 (dihydro-N-retinylidene-N-retinylphosphatidyl-ethanolamine), a diretinal compound formed in within rod discs[146]. Lipofuscin A2E is believed to create photooxidative damage to RPE cells. Extracellular aggregates of lipids and proteins called drusen, along with thickness changes in Bruch's membrane, can also interfere with the proper diffusion of RPE nutrients and metabolites. These disruptions to the RPE ultimately lead to

photoreceptor degeneration in the macula, leaving the afflicted with scotomas in place of their high-acuity central vision. AMD is considered dry without any neovascularization, and late stages of dry AMD are marked by geographic atrophy in photoreceptors, RPE, and choroidal capillaries. When neovascularization occurs, edema and bleeding can cause sudden and severe damage to not only the RPE and photoreceptors, but also the inner retina[147].

A far rarer disease that also interferes with RPE maintenance of photoreceptors, Bietti crystalline dystrophy, relates to biallelic mutations of the *CYP4V2* gene[148]. *CYP4V2* codes for an enzyme involved in the hydrolysis of omega-3 polyunsaturated fatty acids. While the specific functions of this enzyme are not well understood, its dysfunction may disrupt the proper metabolism of photoreceptor discs. The disease is characterized by reflective crystalline deposits around the RPE and Bruch's membrane, which may be products of unhealthy RPE activity[149]. Crystals may also manifest in the paralimbal cornea. Disappearance of subretinal crystals coincides with local atrophy of the RPE and photoreceptor inner and outer segments. Visual deficits start manifesting in the central and pericentral visual fields, and patients typically present with reduced visual acuity and nyctalopia. As the disease progresses, larger sections of the visual field can become compromised, leaving those afflicted legally or completely blind in the fifth or sixth decade of life.

Other diseases that present with similar visual deficits as Bietti crystalline dystrophy, but without subretinal crystals, are collectively called retinitis pigmentosa (RP). Inflammation of the retina is not a prominent factor, despite the implication of the name, but clumps of hyperpigmentation from disrupted RPE, called "bone spicules" by analogy, can be seen in fundus examinations[150]. RP diseases all involve degeneration of the photoreceptors and the RPE. Rod-cone dystrophies are most common, in which rod deterioration progresses first, and cones are

not affected until later disease stages. Cones are not primarily affected in rod-cone dystrophies, but secondarily succumb to degeneration after most rods have fallen apart[151]. Cone-rod dystrophies, however, affect both photoreceptor types simultaneously, and visual acuity loss occurs earlier in the disease progression. Many forms of RP are known to stem from mutations of specific genes, with various inheritance patterns[152]. RP disease onsets and progressions vary, but end-stages can leave patients without any light perception.

Any of these or other diseases that destroy photoreceptors cause surviving neurons and glia to undergo many deleterious changes. After a majority of photoreceptors die, inner retinal neurons are deprived of their normal synaptic input. As a putative consequence of abnormal retinal activity, neurons begin sprouting new connections to form synapses in unusual locations. Glycinergic amacrine cells have been shown to atypically project into the outer plexiform layer in RP. Although surviving cones in rod-cone dystrophies still provide some input to the retinal network, neurons also seem to display changes in gene expression, such that many rod bipolar cells switch from having metabotropic ON pharmacology to appearing as OFF bipolar cells. Müller cells also show metabolic changes as photoreceptors become silenced, such as initial upregulation of glutamine synthetase expression, followed by its pronounced lack of expression[153].

Once the oblation of rods and cones is locally complete, changes within the remaining retina become more dramatic. The Müller cells increase in size and combine with neighbor Müller cells to form large columns between the vitreous and remaining RPE. The columns spread out externally to form a partial glial seal that separates the neural retina from the RPE and choroid. As time passes, Müller cell columns, as well as pigment clusters in RP, act as conduits for neuron migration. Some neurons from the inner nuclear layer move as far as into the subretinal space,

and neuritogenesis continues amidst relocation. Gross locations in the retina may not change, but stratification and vertical organization of cells disappears. Müller cell metabolic changes progress chaotically, with such variation that the glia likely do not retain the ability to sustain neural activity. Neurons steadily die over time, leaving the Müller cells as the last cells to survive in the retina[153].

2.6.2 Damage downstream of phototransduction

Failures farther along neural visual pathways can result from other sources of damage, and create particular visual field defects. The next step from the photoreceptors that is prominently vulnerable is the optic nerve, and glaucoma is most common form of optic neuropathy. Glaucoma manifests as progressive damage to the RGCs with accompanying losses in the visual field, without symptoms implying non-glaucomatous diseases such as anterior ischemic optic neuropathy or demyelinating optic neuritis. Many forms of glaucoma result from high intraocular pressure related to insufficient drainage of aqueous humor from the anterior chamber. Pressure builds back through the posterior chamber and is communicated to the retina and the optic disc. Resulting chronic deformations of the lamina cribrosa apply strain to and critically damage RGC axons[154]. Instabilities in ocular blood flow[155] or genes predisposing lamina cribrosa deformation and RGC apoptosis[156] can also lead to glaucoma development with normal intraocular pressure. As the axons are damaged and RGCs die, visual field deficits can appear in sectors or arcs projecting from the physiologic blind spot. Field defects need not be symmetrical between eyes, so vision loss can go undetected until later stages of the disease. Most instances of glaucoma affect both eyes, and can ultimately lead to the complete loss of light perception.

Injuries to the RGC axons outside the orbit can result from infarcts, trauma, and pressure from abnormal growths or edema[157, 158]. Demyelinating optic neuritis, such as in connection with multiple sclerosis, can also interrupt RGC signaling by attacking axon myelin. Acute myelin damage can cause vision loss that often begins to improve within a month[159], but permanent visual deficits and blindness can follow related damage to RGC axons[160, 161]. Such optic nerve damage can have visual field manifestations similar to those of glaucoma. Severe optic nerve damage can destroy the entire monocular visual field, although such loss can be masked by sight from the fellow eye if damage was not bilateral. Damage to optic chiasm, where signals from each eye get divided between cerebral hemispheres, results in bitemporal hemianopia[162]. Because the chiasm contains axons from the nasal half of each retina, the temporal half of each monocular visual field would be blind. Binocular viewing would still lack perception in the temporal monocular crescents, the lateral extremes of the visual field only visible to the nasal retina of the ipsilateral eye. Posterior to the chiasm, injury to the optic tract can produce homonymous hemianopia, in which vision is lost in the contralateral visual field of both eyes. If the left optic tract suffers damage from an infarction, the right side of both monocular visual fields would be blind. As with damage at any level of the visual system, bilateral complications can destroy all visual perception.

Homonymous hemianopia similarly appears with lesions that obliterate unilateral visual processing anywhere beyond the optic tract. Complete destruction of the LGN would therefore have a visual field presentation much the same as a severed optic tract. Incomplete hemianopias present as more congruous the farther from the chiasm damage is inflicted. In addition to hemianopias, LGN lesions can also result in sectoranopias. LGN is supplied by both anterior and lateral choroidal arteries, and loss of blood flow to one or the other creates identifiable visual field defects. Anterior choroidal artery blockage disturbs blood flow to the

medial and lateral horns of LGN, corresponding to defects in the contralateral inferior and superior visual field. Lateral choroidal artery disruption creates defects in the contralateral central visual field[163]. Permanent bilateral damage to the entirety of LGN is very rare, but has caused nearly-complete vision loss following infarcts related to preeclampsia[164] and diarrheic episodes[165, 166].

The optic radiation projecting from LGN can be damaged to form deficits in either the superior or inferior contralateral visual fields, or complete homonymous hemianopia. Homonymous hemianopia results from damage to the axons in the internal capsule before paths divide, or deep posterior lesions in temporal or parietal lobe. Damage to superior parietal lobe may only affect the parietal optic radiation fascicle, and present with contralateral homonymous inferior quadrantanopia. Internal capsule and parietal injuries to the optic radiations typically develop as part of infarcts or from neoplasia. Damage to the temporal fascicle, or Meyer's loop, is a common complication of anterior temporal lobectomy. Depending on the extent of resection, destruction of the temporal fascicle can cause homonymous superior quadrantanopia[163]. If joined radiation fascicles are damaged bilaterally, or the radiations are affected by widespread damage, as in cerebral palsy[167], the afflicted individual can become completely blind. Blindness from bilateral lesions beyond the LGN is classified as cerebral blindness.

Visual cortex is commonly damaged by disruption of the posterior cerebral artery flow and traumatic insults. The posterior cerebral artery supplies the entire visual cortex of one hemisphere, but the posterior pole of occipital cortex also receives blood from a branch of the middle cerebral artery. As the posterior pole holds the representations of the macular visual field, ischemic infarcts about the posterior cerebral artery can produce homonymous hemianopias that spare the macula. The parts of visual cortex that rely entirely on the posterior

cerebral artery fail, blinding most of the contralateral visual field, but the supply from the middle cerebral artery can maintain vision from the contralateral macula. Conversely, because the neural volume devoted to the macula is disproportionately large, a traumatic injury to the posterior pole can produce a contralateral macular scotoma without necessarily affecting peripheral vision. If the basilar artery is occluded, which supplies both posterior cerebral arteries, extramacular visual fields can be blinded bilaterally. More general hypotension affecting the entirety of the primary visual cortex, or any other large lesion spanning V1 bilaterally, can destroy all visual perception. This form of cerebral blindness is labeled as cortical blindness, and those afflicted are often unaware of their lack of vision[163, 168].

3 Visual prostheses

Various methods have been explored aimed at restoring vision to the blind, particularly those blinded by damage to neural tissue. Stem cell treatments attempt to replace damaged neurons with new cells that can perform the same tasks. Optogenetic approaches bestow light sensitivity upon surviving neurons so that they can be selectively activated by artificial light signals[169]. Visual prostheses, the focus of the research described in subsequent chapters, are designed to act in place of damaged neural tissue by electrically or magnetically stimulating surviving neurons. The targets of visual prostheses would have normally received signals from more distal neurons along the visual pathway, but damage has prevented the conduction of that signal. Prostheses activate their targets so that they can provide neural input to the proximal visual system. The visual stimuli generated by prostheses, as the visual system is no longer directly responding to light, are called phosphenes. Although current prostheses do not generate signals similar to normal retinal output, simple stimulation techniques have restored some visual perception to blind individuals.

3.1 Visual prosthesis targets

In principle, visual prostheses can be designed to target any part of the visual system. For any particular individual, a prosthesis will only be effective if it interfaces with tissue proximal to the blindness-causing damage. For example, if a patient suffers from glaucoma with severe optic nerve damage, the only viable prosthesis targets would be neurons that would have directly or indirectly received signals from the optic nerve. Stimulating more distal neurons or optic nerve sections in the retina, such as bipolar cells or surviving RGC somata, would not make sense. Those cells would likely still be receiving proper visual information from photoreceptors, but

their signals are not being successfully conducted to more central parts of the visual system. The normal optic nerve target for visual information is the LGN, so a prosthesis targeting the thalamus or visual cortex might be able to replace some of the missing input from the optic nerve to generate light perception in this example.

Aside from recipient pathology, factors surrounding surgical implantation and neural representations of vision also influence choices of prosthesis targets. The retina is a popular target because it is more accessible than intracranial targets, and the spatial configuration of neurons corresponds directly to their representations of the visual field. Retinal stimulation, however, is the most limited in applicable pathologies. The optic nerve is similarly accessible for stimulation, but relationships between electrodes and the visual field need to be determined experimentally for each individual after implantation. LGN and cortical prostheses are appealing for their broader range of applicability, but both require more complex surgery and post-surgery visual field mapping.

3.1.1 Retina

When a prosthesis is designed to target the retina, it must ultimately elicit signals from RGCs. Intended targets may be more distal along the visual pathway, such as bipolar and amacrine cells in the inner nuclear layer. The only forms of blindness that are addressable by retinal prostheses are those in which photoreceptors have died or otherwise ceased to function. Reasonable numbers of RGCs must be spared from damage, such that there is still a functioning optic nerve. Optic nerve function is typically verified through visual percept responses to either bright light or electrical stimulation of the eye. Stimulating arrays can be placed in numerous locations, typically above or below the retina, or between the choroid and sclera.

Epiretinal visual prostheses are placed between the vitreous body and the nerve fiber layer of the retina. One or more retinal tacks secure the array to the sclera beneath the retina, maintaining the array's position and ensuring contact between the electrodes and the retina. Epiretinal electrodes should ideally stimulate the RGC somata directly beneath them, which would create percepts localized to that area of the visual field. Given that the nerve fibers between the electrodes and the RGC somata are unmyelinated and have voltage-sensitive ion channels, there is considerable risk of activating axons unintentionally. Axonal stimulation would create percepts at locations corresponding to somata more distant from the optic disc, which might appear as a line or arc when all involved RGCs are considered[170].

Instead of implanting a stimulator above the retina, subretinal prostheses are placed in the space once inhabited by photoreceptor outer segments between the outer limiting membrane and the RPE. Fluid is used to induce a retinal detachment at the desired site of implantation, and reattachment of the retina to the RPE holds the device in place without a retinal tack. In concept, stimulation from the outer side of the retina might allow for utilization of some inner retinal circuitry and processing. Even if a subretinal prosthesis can preferentially stimulate bipolar and amacrine cells instead of RGCs, remodeling of synaptic organization in the degenerated retina would likely destroy the desired fine lateral processing. In practice, performance has not differed greatly between epiretinal and subretinal prosthesis recipients, so the benefits of subretinal stimulation may not be as dramatic as originally hoped. With design improvements, however, the distance left between subretinal electrodes and RGC axons may prove advantageous, assuming that RGC axon stimulation hinders epiretinal prosthesis performance.

Limiting surgical entry into the eye, suprachoroidal prostheses are placed closer the eye's exterior, either between the choroid and sclera or within a pocket sliced midway through the depth of the sclera. The latter placement is sometimes referred to as intrascleral. Similarly to subretinal arrays, a suprachoroidal stimulator is mostly held in place by the apposition of the choroid and the sclera, or for intrascleral placements, the walls of its scleral pocket. A scleral anchor patch can prevent connected wires from applying force to and moving the stimulating array. While suprachoroidal placement may pose less risk to the interior of the eye and offer stable positioning of stimulator, the increased distance from the neural retina makes eliciting percepts more difficult. In order to properly activate surviving RGCs, more charge would be required from a suprachoroidal prosthesis than from an epi- or subretinal prosthesis. The added distance from the retina may also degrade the spatial specificity of stimulation, reducing the visual acuity one might expect for a prosthesis user.

Regardless of stimulator position relative to the layers of the eye, each still benefits from the predictable representation of the visual field in the retina. Even if axons are stimulated or charge spreads more than desired, electrode placements farther from the macula can be expected to elicit percepts farther into the periphery in corresponding directions. Accordingly, stimulators are typically placed centered with respect to the macula. Such placement provides the best chance that percepts will appear in the center of the user's visual field. While peripheral placements can be used to generate percepts, users may not always be comfortable getting percepts only in an eccentric portion of the visual field. Aside from comfort, RGCs normally exist in higher densities closer to the macula, disproportionately large areas of cortex are devoted to macular processing, and many higher visual functions preferentially rely on information from the macula. If the prosthesis is successful in eliciting percepts, those percepts

may be more useful when they are generated through the greater neural machinery devoted to central vision.

Attached to the intraocular stimulator, extraocular components provide power and interface with outside systems. Early retinal prostheses were primarily derived from successful cochlear prostheses, and some retinal prostheses maintain cochlear prosthesis design concepts for placement of support electronics. Cochlear prostheses typically have connected electronics secured to the temporal bone behind the ear. Instead of routing cables toward the cochlea, retinal prosthesis cables were extended forward to reach the eye and pass through the sclera. Early designs have called for percutaneous connections to supply power and stimulation instructions, which increased risks of infection and restricted options for device use. More developed designs allow for wireless communication of power and data between subcutaneous and external transceiver coils. Some compact designs have also simplified implantation by moving the extraocular implant components from the temporal bone to just outside the sclera. Connected using a scleral band or direct suturing, episcleral electronics move with the eye and only require a short transscleral cable to support the stimulator.

Visual information that drives stimulation can be derived from an external camera, or depend entirely on light that naturally enters the eye. External cameras, those not directly connected to the stimulator, can be easily combined with direct computer control to provide flexibility in how the prosthesis is used. Designing the prosthesis to depend on externally-processed visual information also simplifies image processing revisions, processor repairs, or upgrades, as no implanted components would need to be modified. As this simplifies the design of the implanted components, most retinal prostheses are developed for use with external cameras.

Epiretinal and subretinal designs do, however, have the special opportunity to directly use incident light as the basis for stimulation. Because such stimulators encounter the light that would have normally triggered responses in their targets, directly or indirectly following photoreceptor transduction, photosensitive elements with amplifiers can be expected to intrinsically provide light-appropriate stimulation patterns. Unlike with an external camera, eye movements are naturally taken into account as photodiode-driven stimulation is provided. Suprachoroidal stimulators are not able to use incident light as well, simply because that light is blocked by absorption in the choroid. Without alternative routes of triggering stimulation, though, use of photovoltaic circuits limits the potential for outside control of the device and opportunities for upgrading how the prosthesis operates.

3.1.2 Optic nerve

Rather than target RGC somata or retinal neurons in the inner nuclear layer, visual prostheses can stimulate RGC axons after they have passed through the optic disc. Because the ultimate targets of stimulation are still retinal cells, most of the addressable pathologies for optic nerve prostheses are the same as those for retinal prostheses. Optic nerve stimulation might be preferred over retinal stimulation in the case of a severe retinal detachment, when the retina may not be expected to tolerate or respond well to retinal stimulators. The entire monocular visual field is also compactly represented in the optic nerve, which might be attractive if the scope of stimulation is favored over specificity. Optic nerve stimulators do not require any surgical penetration of the sclera, but implantation can involve lateral orbitotomy[171] or pterional craniotomy[172].

Optic nerve prostheses have been designed with spiral cuff surface electrodes and penetrating electrodes. Spiral nerve cuffs were previously used in a variety of peripheral nerve applications, particularly for eliciting motor responses[173]. Spread out axially along the nerve, electrodes at different azimuthal positions can be used to stimulate axons on different sides of the nerve. Penetrating electrodes, alternatively, can target various azimuths and depths of a nerve at one axial location. Using more focal stimulation, penetrating electrodes may offer lower stimulation thresholds and greater spatial specificity than electrodes resting on the nerve surface.

Unlike the visual field representation across the retina, relationships between sectors of the optic nerve and the visual field are not well structured. Some basic spatial relationships are preserved, such as the location of foveal representations. The fovea is temporal to the optic disc, and axons of foveal RGCs are accordingly concentrated on the temporal side of the optic nerve. More specific relationships, however, such as degree of eccentricity, are not strictly observed in the optic nerve[174]. Although some order in percept locations might be expected based on electrode positions, specific electrode–visual field relationships would need to be determined postoperatively through user responses.

3.1.3 Thalamus

As mentioned above, blindness caused by damage to the optic nerve and tract would only be addressable by targeting more proximal neurons along the visual pathway. Glaucoma or traumatic optic neuropathy can prevent visual information from being relayed from the retina to the LGN of the thalamus. Although the rest of the visual system might function properly, it would be deprived of visual input and leave the patient blind. Most visual information passes through the LGN, and if the LGN is undamaged and the rest of the central visual system is

healthy, direct stimulation of the LGN could elicit visual percepts much like more distal stimulation.

Deep brain stimulation (DBS) has been used extensively for the treatment of movement disorders, and has seen increasing interest for use with epilepsy and psychiatric disorders. One or more electrode leads are implanted into the brain to provide stimulation to the thalamus or basal ganglia. Typical targets for DBS stimulation include the ventral intermediate nucleus of the thalamus, globus pallidus, and subthalamic nucleus[175]. The LGN is not far from these common DBS targets, and microstimulation applied to primate LGN has been shown to generate visual percepts[176].

Aside from bypassing the optic nerve and tract, there are other benefits to LGN stimulation related to neural organization. Axons originate from the fovea with greater density than other parts of the retina, and that bias in visual field representation is amplified even further in the LGN[177]. An electrode that elicits percepts in the central field could thus be expected to generate smaller and more precise percepts. Neurons in the LGN also display fine laminar organization. Parvocellular layers receive input from midget RGCs, and thus carry red/green color information. Koniocellular layers are innervated by small bistratified RGCs, and thus activate in association with short-wavelength light. Further, parvocellular layers are divided into layers with ON-center dominance and OFF-center dominance[178]. LGN stimulation therefore offers not only spatial specificity for generating percepts, but also some opportunity for color tuning and ON vs. OFF pathway selection. Realizing the benefits of stimulating different layers, however, would likely require using combinations of electrodes beyond the capabilities of early LGN devices.

LGN also possesses a distinct retinotopic organization. Given the method of electrode implantation, however, there is considerable uncertainty precisely where electrodes will fall

within that retinotopic map. Thus, like with optic nerve stimulation, the visual field positions of percepts would not be known until after the LGN prosthesis is implanted and the user can respond to stimulation. Regardless of where electrodes are implanted in a single LGN, though, all percepts would be localized to the contralateral visual field. While retinal and optic nerve stimulators have potential access to most of the visual field with only unilateral implantation, comparable coverage with an LGN prosthesis would require bilateral implantation into both sides of the thalamus.

3.1.4 Visual cortex

Prostheses targeting visual cortex have the widest applicability of all visual prostheses. They can serve patients blind from photoreceptor degenerations or glaucoma, just as LGN prostheses could. In cases of damage to the thalamus or optic radiations, cortical implant capabilities would persist whereas LGN prostheses would be ineffective. Vision loss from LGN damage is rare, but does occur in association with vascular lesions[164, 179]. Extensive optic radiation damage is also not common, but can be a complication of temporal lobectomy for epilepsy[180] or cranio-cerebral trauma[181]. Multiple sclerosis can have wide-spread effects on the optic radiations, but such lesions are often transient and would not necessitate a visual prosthesis[182]. Persistent vision loss from non-cortical cerebral damage, while addressable by cortical prostheses, may however only have consequences for one side or specific sections of the visual field. As such, situations specifically requiring a cortical visual prosthesis may be limited to less common cases in which non-cortical cerebral damage affects the entire visual field, or prosthetic stimulation is desired to fill in the patient's scotomas. Like for LGN prostheses, any attempt to generate phosphenes on both sides of the visual field would require bilateral implantation of stimulators.

Primary visual cortex (V1) would be an ideal target of stimulation for cortical visual prostheses, in terms of information transfer. Because it is the major direct recipient of LGN projections, V1 is the site of earliest cortical processing in the visual system, and has the simplest representations of vision in the cortex. The chances of mismatch between intended and elicited percepts would likely become greater as synaptic distance from V1 grows and the complexity of visual patterns that normally trigger activation increases. Unfortunately, the topography of V1 makes stimulation from an implanted device impractically difficult. The horizontal meridian of the visual field, peripheral to eccentricities 1°–4° from the center of vision, is buried beneath the medial surface of the occipital lobe along the bottom the calcarine fissure.

As the fissure would not accommodate stimulators easily, the foveal projection of V1 at the occipital pole would need to be supplemented with targets in higher visual areas. V2 and V3 have foveal representations next to that of V1 on the ventrolateral cortical surface next to the occipital pole. From this confluent foveal representation, V2 and V3 circumscribe V1, curving out along the dorsolateral and ventral occipital surfaces before they extend anteriorly across the medial occipital wall[183]. The dorsolateral and medial exposures of V2 and V3 are much more accessible than extrafoveal V1. Thus, instead of targeting only V1, a stimulator could target a larger area along the medial occipital cortex surface. Those portions of the visual field not stimulated through V1 could be stimulated in V2 or V3. Alternatively, stimulation of V2 and V3 along dorsolateral occipital cortex, within approximately 4 cm of the occipital pole, has also been shown to generate simple visual percepts in response to stimulation. As only the lower quadrant of the contralateral hemifield is represented in the dorsal arms of V2 and V3, percepts from dorsolateral stimulation would be confined to the lower half of the chosen hemifield[184, 185]. Multiple complete hemifield maps are exposed more anteriorly along the dorsolateral surface of cerebral cortex, but these areas have more complex response biases, such as toward motion in

areas V3A and MT, or objects in areas LO1 and LO2, and may not produce simple percepts in response to direct stimulation[186].

There is considerable variation in cortical area boundaries among individuals. Sizes of cortical areas can differ by factors of 2 or 3 between individuals. As such, there will be some uncertainty regarding what cortical areas are being stimulated by a prosthesis. That uncertainty eclipses finer questions regarding visual field mapping, as each area will have its own visual field map. Particularly for stimulators spanning V1–V3, retinotopy with respect to the stimulator will be disturbed whenever area boundaries are encountered. Efforts to use anatomical landmarks[187] and features, such as myelination[188], have produced atlases of cerebral cortex that might one day be helpful for predicting an individual's personal visual field representations. Until any such atlas method is perfected, however, mapping electrode stimulation to percept locations will need to be done based on subject responses.

3.2 Devices

Transforming prosthetic vision efforts from promising concepts to usable devices is difficult, and attempts have been made with many approaches and various levels of success. One of the first visual prostheses, applying stimulation based on visual information, was developed by William H. Dobelle. Subjects were implanted with subdural arrays of 64 platinum electrodes on the medial surface of occipital cortex during 1978 and earlier[189]. The arrays relied on percutaneous connections for power and stimulation commands. The stimulator of one subject was connected to a camera to form a visual prosthesis in 2000[190], followed by other subjects with implanted cortical electrodes. While providing some visual information, these devices were ultimately plagued by infections, seizures, and system failures[191].

Shortly after the Dobelle implants were connected to cameras, the artificial silicon retina (ASR) was developed by the Optobionics Corporation as the first retinal prosthesis[192]. Implanted subretinally, the ASR used about 5000 microphotodiodes to stimulate surviving retina. The ASR did not have an external power source, and relied entirely on incident light to generate stimulation. Implantations were performed extramacularly, at about 20° eccentricity. Although the ASR did stimulate the retina with electrical current, current levels were below the threshold necessary to generate visual percepts. The main benefit of the ASR was neurotrophic in nature, slowing or even reversing the progress of photoreceptor degeneration[193]. The ASR failed as a visual prosthesis, but nevertheless aided subsequent attempts in creating more functional devices.

3.2.1 Devices approved for marketing

Following the problems encountered with early visual prostheses, several retinal prostheses have demonstrated sufficient stability and safety to earn regulatory approval for marketing. Retina Implant AG's Alpha IMS and Alpha AMS, as well Pixium Vision's IRIS II, have CE marking for use in Europe. Second Sight Medical Products, Inc.'s (SSMP's) Argus II also has CE marking, as well as Food and Drug Administration (FDA) approval under a humanitarian device exemption (HDE) for use in the United States. The Argus II further has approval for use in Canada and certain Asian countries. Although none of these devices offer high-quality vision restoration, the proliferation and wide-spread use of visual prostheses helps accelerate the development of more effective prosthetic technologies.

3.2.1.1 Argus II

SSMP began a feasibility study for the Argus II in 2006, following a small study with its predecessor, the Argus I. While the Argus I had an epiretinal array of 16 electrodes and a transceiver implanted in the temporal bone, the Argus II was designed with 60 electrodes and a periocular transceiver. Special glasses and a video processing unit (VPU) provide power and stimulation commands to create visual percepts. During the feasibility study, implanted subjects demonstrated benefit for some simple visual functions, such as light localization[194], hand-motion guidance[195], direction of motion detection[196], following lines on the floor[197], or even letter recognition[198].

The electrode array of the Argus II consists of 10 × 6 flat platinum electrodes. Each electrode is 200 µm in diameter, with 300–600 µm spacing between electrodes. Arrays implanted earlier in the feasibility study, called slotted arrays, had center-to-center distances of 575 µm horizontally, and 725 µm vertically. Revised array designs, dual-metal arrays, were smaller with 525 µm electrode distances both horizontally and vertically. Slotted arrays, not including the silicone perimeter, thus measured 5.75 mm × 4.35 mm, and dual-metal arrays were 5.25 mm × 3.15 mm. Expected extents of visual field coverage were approximately 20° × 15° and 18° × 11°.

Arrays were placed, for most recipients, over the macula, such that the nasal edge of the array sat near the optic disc. Each array was secured in place using a retinal tack, which passed through the temporal side of the array and through the sclera. A compressed spring on the vitreal side of the tack maintains pressure to fix array against the retina. If necessary, an additional tack could be inserted through the nasal side of the array to improve array-retina apposition. Temporal to the primary tack, conductors combine within a ribbon cable that passes through the sclera to connect the array to extraocular electronics.

On the other side of the pars plana incision through the sclera, a cylindrical electronics case sits sutured to the sclera between the superior and lateral rectus muscles. The case is hermetically sealed and contains the circuitry necessary for translating stimulation commands into current pulses for each electrode. The ribbon cable that attaches the electronics case to the electrodes projects from the case anteriorly. The case's position is further maintained by a scleral band, which extends past the inferior and superior borders of the case. The silicone band passes under the extraocular muscles and is pulled tight around the sclera. The periocular antenna is also attached to the scleral band, centered underneath the lateral rectus muscle. Coils of wires encased in silicone allow the electronics case to communicate and receive power from external equipment. The inferior edge of the coil is sutured to sclera, while other sutures on the nasal side of the eye to hold the rest of the scleral band in place.

The wireless connection to the periocular electronics is very short range, and requires the external antenna to be within a few centimeters. The external antenna coil is encased in plastic with an RF circuit board, both attached to a long cable. In surgery, the positions of the periocular and external coils are adjusted until the VPU registers a usable connection. For daily use, the external transceiver package, including the antenna coils, is attached to the arm of a glasses frame on the side of the implanted eye. The VPU beeps to alert the user whenever there is no connection, prompting the user to either adjust eye orientation or the position of the glasses. The position of the antenna with respect to the glasses can also be moved along the anterior-posterior axis to improve link strength.

The opposite end of the transceiver's cable connects directly to the VPU. In standalone mode, the VPU is the most distant part of the system from the implanted array. The battery pack attached to the VPU supplies all power used by the processor and the implant. Any commands

for stimulation must be generated in the VPU, either based on video input or direct computer control. When the VPU is connected to a computer with special software, custom pulse trains can be configured for selected electrodes based on current amplitude, phase order (cathodic or anodic pulse phase first), phase duration, interphase gap time, pulse frequency, and overall stimulation duration. The programming computer can also be used to check electrode impedances and change settings within the VPU's firmware. When an electrode's impedance exceeds a specified threshold, usually around 100 kΩ, the electrode is considered faulty and the VPU automatically disables the electrode. A disabled electrode cannot be used for video-based or direct stimulation, and can only be activated for further impedance checks.

The VPU receives video input from a camera seated in the glasses. The camera is positioned just above the center of the nosepiece, and can be physically rotated to match the orientation of the electrode array within the eye. Captured video is sent to the VPU along the same cable as communication with the glasses-mounted transceiver. The VPU filters and downsamples a portion of the image into 10 × 6 pixels, and each pixel is given a brightness value within 0–31. A video configuration file (VCF) is then used by the VPU to translate brightness values into current amplitudes for each electrode. VCFs are different for each user, and are based on electrical thresholds for generating visual percepts. Although pulses from different electrodes are ideally spaced out within intervals between repetitions, electrodes and can be combined into shared timing groups to reduce the necessary current output per electrode. Electrodes are commonly grouped into quads, 2 × 2 groups of 4, when perceptual thresholds are unusably high.

While the feasibility study in Mexico and the United States started in 2006, the clinical trial for European use started in 2008. The Argus II subsequently received CE marking in 2011, making it the first visual prosthesis approved for marketing. This approval, however, was limited to

Europe, and was still subject to country-specific regulations. The feasibility study in the United States continued through 2013, when the FDA granted an HDE for commercial use of the Argus II. The HDE required demonstration of the device's safety, although significant efficacy was not necessary. Combining investigative implantations and those performed under regulatory approval, over 250 individuals have been implanted with the Argus II. Post-approval studies continue to monitor device safety and collect efficacy data.

3.2.1.2 Alpha IMS and AMS

Retina Implant AG has developed several generations of subretinal photovoltaic visual prostheses, similar in concept to the ASR. Unlike the ASR, however, their prostheses do not depend entirely on incident light for power. Instead, Retina Implant coupled photodiodes with externally-powered amplifiers, and thus ensured that the stimulators could generate enough current to depolarize retinal neurons. Early prototypes were tested in humans starting in 2005, and the short-lived Alpha IMS was given CE marking in 2013. The Alpha AMS replaced the Alpha IMS and was CE marked in 2016.

The intraocular stimulator of the Alpha IMS was placed subretinally beneath the macula. The array spanned 3 × 3 mm and was 0.1 mm thick. Each of 1520 independent stimulation units consisted of a photodiode, amplifier circuit, and a 50 × 50 µm iridium electrode[199]. A polyimide foil conductor sat below the stimulators on the choroidal side of the implant and supplied power and control signals for the amplifiers. Current was returned through a reference electrode implanted beneath the temporalis muscle. Electrodes provided 1 ms pulses at rates of 1–20 Hz, typically set to 5 Hz[200].

The polyimide foil passed through the choroid and sclera about the eye's equator and connected to an extraocular cable. That cable traveled subdermally around the orbital rim and proceeded under the temporalis muscle to behind the ear. The cable terminated in its connection with the receiver coil, which was protected by a ceramic casing and embedded in the retroauricular temporal bone. Although earlier devices relied on percutaneous connections, all power and signals were sent to the Alpha IMS wirelessly, just as for the Argus II. The external coil communicating with the subdermal receiver coil was held in place by a magnet incorporated in the implanted case.

Similarly to how the Argus II depends on the VPU, each recipient of the Alpha IMS was given a handheld control unit that supported the implanted electronics. The control unit housed the battery for the system and connected to the external transmitter coil through a cable. Users could turn the device on or off, and had 2 knobs on the unit for controlling stimulation brightness and contrast.

Clinical trials for the Alpha IMS began in 2010, data from which allowed CE marking in 2013. More than 50 individuals in Europe and Asia were Alpha IMS recipients. Users demonstrated visual function improvements similar to those of Argus II recipients. Unfortunately, device functionality did not persist, and the Alpha IMS systems failed less than 2 years after implantation. One prominent cause of these failures was erosion of electrical insulation.

Retina Implant AG began clinical trials for their next generation implant, the Alpha AMS, in 2014. The updated design increased the number of stimulation units from 1520 to 1600, increasing the footprint of the array to 4 × 3 mm. The polyimide foil was widened to accommodate the new chip, and stimulation was changed from using monophasic to biphasic pulses. The Alpha AMS was CE marked in 2016, and at least 6 individuals have been implanted since the clinical

trial began[201]. The expected lifetime of the Alpha AMS is projected to be 5–6 years, and it has so far been stable through 38 months.

3.2.1.3 IRIS II

Pixium Vision obtained CE marking for its Intelligent Retinal Implant System (IRIS) in 2016. The IRIS devices are similar to Argus prostheses in that they provide epiretinal stimulation based on information from an external camera. Pixium Vision took over the development of the IRIS when it purchased Intelligent Medical Implants (IMI) in 2012. Early IRIS implantations were performed as early as 2005, and Pixium altered the design to develop the IRIS II in 2014. A total of 10 subjects have been implanted with IRIS II since the start of its clinical trial in 2016.

Following acute retinal stimulation with iridium oxide electrodes in 2003–2004[202], IMI first implanted its Learning Retinal Implant System into 4 subjects with photoreceptor degeneration in 2005–2006[203]. IMI soon after renamed its device as the IRIS in 2007. 49 stimulating electrodes were placed over the nerve fiber layer of the macula, secured the with a retinal tack on the temporal side of the array. A polymer film connected the array to extraocular electronics through a pars plana incision. Like the Argus II, the electronics case was sutured to the sclera between the lateral and superior rectus muscles. The case contained an RF transceiver, which received power for stimulation, but not stimulation commands. Stimulation commands were instead received via an intraocular microphotodiode, seated on the film between the tack and the scleral incision[204].

The IRIS Visual Interface was a pair of glasses that incorporated a camera, transmitter coil, and infrared transmitter. The transmitter coil was positioned within the glasses frame on the side of the implanted eye and produced a high-frequency alternating electromagnetic field. The

periocular coil drew power from the external coil's activity and supplied electricity for stimulation circuitry and electrodes. The infrared transmitter was placed in front of the eye and optically communicated stimulation commands to the intraocular microphotodiode. This optical approach was favored to allow for high data transmission rates, as well as to allow the user to stop stimulation by closing the eyes. The glasses-mounted camera sent visual information through a cable to the IRIS Pocket Processor, similar to the Argus VPU, which used the same cable to send power and stimulation commands to the interface and implant[204].

The IRIS was originally named the Learning Retinal Implant System in reference to its learning retinal encoder. The encoder was intended to optimize retinal stimulation by modeling how the visual system translates stimulation input into experienced percepts[205]. It was implemented as a large set of spatiotemporal filters that determined stimulation patterns with respect to captured video. During programming, subjects would be presented with several random encodings and provide feedback on which appeared most like a specified shape. The process would be repeated, and the encoder would use the responses to determine an optimum combination of filter settings for the implantee. Simulations with sighted subjects estimated that fewer than 100 iterations would be required to program an encoder.

By the time Pixium developed its IRIS II version of the prosthesis, several enhancements were added to the design[206]. Instead of 49 electrodes, the IRIS II contains 150 stimulating electrodes, each less than 100 µm in diameter. The connection between the array and the retinal tack was also changed so that the tack itself would not need to be removed to explant the stimulator. The IRIS II's tack does not directly hold down the polymer film, but the film is placed such that a hole circumscribes the tack, and a silicone retainer ring locks the film in position. The retainer ring

can be removed without removing the tack, allowing the film and array to be explanted while leaving the retina tacked to the sclera[207].

Outside the eye, the IRIS II uses an asynchronous time-based image sensor (ATIS) instead of a conventional camera to capture visual information[208]. The ATIS was designed to capture the image-change information associated with the magnocellular visual pathway, with less focus on static parts of a scene. Each pixel of the sensor autonomously signals changes in its receptive field, with temporal resolutions finer than 2 ms. Unchanging, sustained visual information is also collected in parallel so that retinal stimulation can be based on changes or static luminance[206].

Pixium completed a clinical trial with the 49-electrode IRIS between 2013 and 2017. The company filed for CE marking of the IRIS II in December 2015, and began the IRIS II clinical trial January 2016. CE marking was obtained in July 2016, and recruitment for the clinical trial was completed with 10 implanted subjects in January 2017. One additional recipient was commercially implanted in September 2017. In October 2017, Pixium announced that IRIS II devices were ceasing to operate properly 9–12 months after implantation, and that eye movements may be causing microfractures in electrode wiring. Moisture accumulation may have also created shorts within the device. Pixium has proposed explanting and re-implanting devices for their clinical trial subjects after application of additional sealant to vulnerable parts of the system. New IRIS II implantations have been halted until design modifications can be completed, perhaps until 2019 or later[209].

3.2.2 Devices starting clinical trial testing

Although only the Argus II, Alpha IMS and AMS, and IRIS II have received regulatory approval for marketing, many other visual prostheses are being developed with the hopes of future

commercialization. SSMP intends to expand its client base by producing the Orion visual cortex prosthesis. In tandem with its testing of the IRIS II, Pixium has started clinical trials for its subretinal PRIMA (Photovoltaic Retinal IMplAnt) device. NIDEK Co., Ltd. and Osaka University are collaboratively developing an intrascleral suprachoroidal prosthesis and have performed several in-human studies since before 2007. Bionic Vision Australia also conducted a clinical trial for a prototype 24-channel suprachoroidal retinal prosthesis in 2012–2014, and successor technologies are being separately developed by the Bionic Vision Technologies (BVT) company and a collaboration between the University of Sydney and University of New South Wales. A collaboration headed by the Illinois Institute of Technology (IIT) is also preparing a clinical trial for an intracortical visual prosthesis targeting dorsolateral occipital cortex.

SSMP implanted a NeuroPace Responsive Neurostimulation System onto the right medial surface of the occipital lobe of one blind subject in October 2016. As a chronic proof-of-concept test, the company confirmed that surface stimulation of medial visual cortex could elicit reproducible phosphenes through at least a 7-month period, and no serious adverse events were reported within 11 months following implantation[210]. SSMP's Orion Cortical Visual Prosthesis System is designed to stimulate medial visual cortex with flat surface electrodes, attached to a nearby transceiver coil similar to that of the Argus II. The FDA provided final approval for SSMP to implant 5 blind subjects with Orion devices, as well as Expedited Access Pathway designation for the clinical trial, in November 2017. The first Orion implantation was performed at University of California at Los Angeles (UCLA) in January 2018, and the remaining implantations are expected to be performed at UCLA and Baylor College of Medicine in Houston[211].

After acquiring IMI in 2012, Pixium began collaborating with Stanford University's Daniel Palanker in 2013. Palanker has worked on the development of a photovoltaic retinal prosthesis since before 2003[212]. Rather than relying on natural light like the ASR, subretinal photodiodes were designed to respond to high-intensity near-infrared light that would be projected in pulses by user-worn goggles. Infrared wavelengths were chosen to avoid phototoxicity, and pulsed stimulation was desired to permit charge balancing on the electrodes. Stimulation units each consist of a central electrode coated with sputtered iridium oxide, which is surrounded by 2 or 3 photodiodes linked in series, and a circumferential return electrode at the unit's perimeter. Each stimulation unit is 70 µm wide. 142 units in a hexagonal matrix form one implantable array, approximately 1 mm in diameter and 30 µm thick. Multiple arrays can be implanted to cover larger areas the visual field[213]. Pixium is commercializing the design as the PRIMA, and was given French regulatory approval to begin a clinical trial with dry-AMD patients in October 2017.

Leading the development of suprachoroidal retinal prostheses, NIDEK and Osaka University published results of acute tests of retinal stimulation from within a scleral pocket in 2007. Arrays of 9 200 µm platinum electrodes were inserted into scleral pockets of two subjects with RP. A platinum wire was inserted through the pars plana into the vitreous cavity as the system's return electrode. The researchers termed this configuration suprachoroidal-transretinal stimulation (STS). Both subjects were able to detect phosphenes in response to STS[214]. The NIDEK-Osaka group later built upon their acute study with a 1-month semi-chronic trial with 2 subjects in 2010[215], and a 1-year trial with 3 subjects in 2014–2015. The prostheses were all explanted after each trial concluded. Devices in both of the later trials had 49 500 µm electrodes, although only 9 were active in the 2010 version. The STS systems received video input from an external camera, and an external processor conveyed stimulation instructions and power through a transceiver coil implanted in the temporal bone. Subjects occasionally

performed better in visual tasks with STS than with the system turned off. The subjects that demonstrated the most benefit from STS had the stimulator implanted close to the fovea. In the 1-year trial, the investigators noticed that functional performance did not vary between returning current through the electrode in the vitreous cavity or the system's subdermal decoder. In light of repeated iridocyclitis, possibly related to prolonged electrical stimulation, future designs will only have a return electrode on the subdermal decoder. NIDEK and Osaka University also plan to expand the stimulated field of view by implanting 2 adjacent electrode arrays with future designs[216]. The developers hope to begin clinical trials for market approval in 2018[217].

BVA first implanted a subject with its 24-channel prototype suprachoroidal retinal prosthesis in May 2012. In each of 3 subjects, a 19 × 8 mm array was inserted into the suprachoroidal space, targeting the macula. Arrays included 33 400–600 μm platinum electrodes and two 2 mm return electrodes. Electrodes along the perimeter of the array were connected to jointly act as a return, along with the 2 large return electrodes and a separate return electrode behind the ear. The remaining 20 electrodes were used for stimulation, and were arranged in a 3.5 × 3.5 mm hexagonal matrix. The stimulating array was powered and controlled through a retroauricular, percutaneous connector, and its use was therefore limited to laboratory testing. Implanted subjects achieved visual function similar to that seen with other retinal prostheses, and the prototypes were explanted 2 years after implantation. BVT is currently developing a 44-channel prosthesis with 1 mm electrodes to follow up on BVA's device. The University of Sydney and University of New South Wales, formerly affiliated with BVA, is also conducting preclinical tests of a 99-channel suprachoroidal prosthesis, the Phoenix-99. Both successor projects to BVA's prosthesis will rely on wireless power and communication[218].

IIT's Philip Troyk began taking interest in cortical visual prostheses in 1983. After some cortical stimulation work with the National Institutes of Health (NIH) terminated, Troyk developed a team across multiple institutions to form the Intracortical Visual Prosthesis (ICVP) Project in 2000. The team decided to avoid the calcarine fissure and medial occipital lobe, and instead opted for stimulation of dorsolateral occipital cortex. Unlike Dobelle's implants or the Orion, the ICVP will use small iridium electrodes that penetrate into visual cortex. Eighteen electrodes are clustered together on each Wireless Floating Microelectrode Arrays (WFMA), along with stimulation circuitry and a transceiver coil. The WFMAs are expected to "float" on the arachnoid and pia mater in the subdural space, without any wire connections, only anchored to the arachnoid and pia, and the underlying cortex, by the electrodes. Each WFMA is approximately 5 × 5 mm, and multiple WFMAs can be implanted to expand the visual field coverage of the ICVP. Implanted WFMAs would be wirelessly and independently activated through a telemetry controller worn in a headband. A camera on the headband would provide source information for a belt-worn video processor[219]. The ICVP Project received $11.8M from NIH in April 2017 to develop and perform a clinical trial with their device, and human implantations may be expected in 2019 or later.

4 Camera alignment

4.1 Field of view

Normal human monocular visual fields span about 160° horizontally and 135° vertically[220]. The central 10° about the fovea, however, have the finest acuity and are afforded the greatest processing power in visual cortex[221]. To fully examine a scene, one moves the head and eyes so that photons from regions of interest land upon the fovea. Fine details only exposed to the retinal periphery are either not detected or lost amongst signal averaging in neural circuitry. Visual details of interest spanning large areas must be captured by the retina in series and subsequently stitched together using location information discussed in Section 2.5.

Visual prostheses, because of technical constraints, are only able to stimulate small subsections of the visual field. Retinal devices tend to stimulate within concentrated areas of the visual field, typically 20°–30° wide, but may stimulate larger arcs with axonal stimulation[170]. These devices are usually centered with the fovea to take advantage of the greater cortical representation and conform to more natural viewing. Optic nerve, thalamic, and cortical prostheses access larger areas of the visual field, but the points of stimulation are less concentrated because of difficulties with electrode placement and how the field is represented at these sites.

As discussed in Chapter 3, all visual prostheses need to capture visual information in some way. Ideally, the prosthesis would capture the same visual information that retina would have received under normal conditions. Natural-light photovoltaic prostheses, such as the Retina Implant Alpha devices, accomplish this by detecting visible-light photons inside the eye and triggering stimulation at the same location. The field of view (FOV) of the electrode array is

simply determined by its size and position on the retina, and where the eye is pointed. This is the same FOV that the targeted retina would be expected to process for the brain.

Other types of retinal prostheses, including the photovoltaic PRIMA, and all intracranial prostheses, instead depend on cameras separated from stimulators for capturing visual information. This affords great flexibility in system expansion and control, as the stimulators always rely on external commands, but the acquisition of information tends to be distanced from its natural configuration. The closest a camera can come to viewing the exact same visual information the eye should receive would require it to be placed on or inside the eye. The eye and the camera could then be made to share the same optical axis, and the two would therefore always point in the same direction. It would not be possible for the eye to point in a different direction and provide incorrect localization cues. As of right now, however, the only remote cameras used or planned for protheses are separated from the eye, such as being fixed to a pair of glasses.

Even if a remote camera were optically coaxial with eye, however, the system would still require some translation of the camera image to electrode stimulation. The prosthesis stimulator, always stimulating the same patch of neurons, would elicit phosphenes in specific regions of the visual field. The shape and expanse of the stimulated visual field would only correspond to a subsection of the camera's FOV. Portions of the camera's FOV would therefore need to be selected for processing and used with the stimulator, while the rest of the FOV would be ignored. The processed portion of the camera's FOV is thus considered the stimulator's FOV.

In order for the stimulator to generate percepts properly, the stimulator's FOV should match what a healthy retina would have viewed at the stimulated visual field locations. This requires consideration of the position of the camera's FOV relative to what normally would have filled the implantee's visual field, as well as what parts of the visual field are being stimulated. For example, because the fovea does not lie on the optical axis of the eye, a camera with the same optical axis as the eye would not be coaxial with the visual axis of the eye. The optical center of the camera's FOV would not correspond to what should be the center of implantee's visual field. Instead of determining the stimulator's FOV by assuming that the camera's optical center matches the center of visual field, the stimulator's FOV should be based on the eccentric point of the camera's FOV that captures what the fovea should receive. Drawings of the stimulator's FOV and an example of eye-camera misalignment are shown in Figure 3.

Figure 3 – Eye-camera misalignment

Drawings show how eye-camera misalignment might manifest. **Top:** A light bulb is shown within the camera's FOV. The dashed lines indicate the stimulator's FOV within the camera's FOV. The + symbol indicates the CAP, which is in this case set to the optical center of the camera. The bright blob represents how a corresponding phosphene may appear to the prosthesis user, with its location indicated with respect to the camera's FOV. The misalignment is evident in how the locations of the bulb and the phosphene do not overlap. The user would incorrectly point to the phosphene when trying to point to the light bulb. **Bottom:** A schematic with eye and camera geometry corresponding to the misalignment seen above. In this case, the implanted eye, shown as a blue ball inside the white ovoid head, is aimed below and to the right of the camera's optical center and the CAP. The light from the bulb is captured by the camera for stimulation (black arrow), but the phosphene is perceived in space where the eye is pointing (blue arrow).

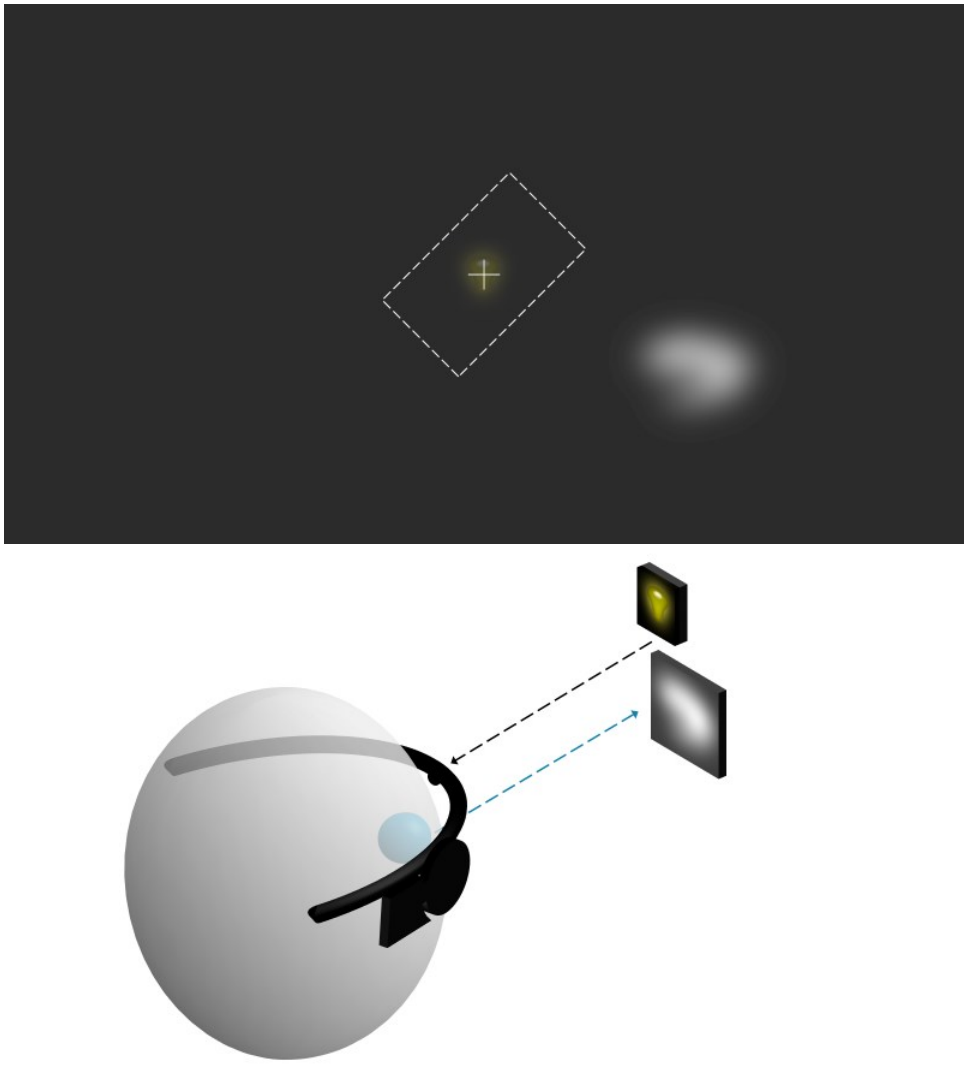

Aligning the stimulator's FOV from a remote camera to match what the corresponding retina would have captured in the visual field, or at least adjusting the stimulator FOV to compensate for localization errors, is called camera alignment. The offset between the camera's FOV center and the center of the stimulator's FOV is referred to as the camera alignment position (CAP). The CAP would not necessarily match the effective visual center of the camera's FOV, which would be the camera point driving the center of vision. Differences between the CAP and camera visual center would occur when stimulation is not symmetric about the center of vision, and the center of the stimulator FOV still drives the physical center of a stimulating array. An example of such asymmetry would be retinal stimulating arrays not centered with the fovea. Less structured stimulators for which retinotopy is not assumed, such as optic nerve or DBS-based thalamic stimulators, would have stimulation locations postoperatively defined relative to the center of vision. Such devices may therefore have the geometry of the stimulator FOV defined to be centered in the visual field, even if phosphenes are not generated symmetrically about the center. If the chosen center of stimulation locations matches the center of the visual field, the CAP would then also be the visual center of the camera.

As a CAP is merely a device setting, it need not be appropriate for coordination tasks and could be completely arbitrary. The CAP that would minimize localization errors in hand-camera coordination is referred to as the optimal CAP (OCAP). The OCAP is an ideal offset between camera and stimulator FOV centers, and may have non-implementable values outside a device camera's FOV.

4.2 Alignment techniques

In the case of a camera optically coaxial with the eye, camera alignment would mostly depend on the angular separation of eye's optical and visual axes. Strictly focusing on correct input to visual pathways, without trying to compensate for nonvisual sources of pointing biases, the CAP should be set so that the camera's visual center is where the eye's visual axis intersects the camera. The OCAP may differ from this CAP depending on complications such as systematic errors in motor commands.

With remote cameras separated from the eye, such that the two have separate optical axes, camera alignment can only be performed for a particular viewing distance. The objective would still be to have the camera's visual center view the same target as the fovea would under normal conditions. Let the line joining the camera's visual center and its viewed target be the camera's visual axis. Physical separation between the eye and camera implies that they cannot share the same visual axis, so the axes can only intersect at one point. A CAP can be defined so the visual axes intersect at a particular distance, such as at arm's length, but deviations between the what camera views and what the eye should be viewing would increase as viewing distances change from that point. Substantially different viewing distances, particularly involving ones close to the user, would require different CAPs for visual axes to align properly.

Lacking direct methods to align camera and eye visual axes, and potential complications with nonvisual pointing biases, camera alignment is typically performed by asking users to perform pointing tasks. The specific contributions of visual and nonvisual sources of pointing bias are not considered as important as the overall pointing bias, which must be corrected to improve performance in pointing and reaching tasks. After either direct stimulation or target localization

point tasks, the CAP is set to match the average location where the user points to indicate stimulus positions.

4.2.1 Direct stimulation

One simple method of aligning the camera's visual axis to where a user claims to see phosphenes is to directly stimulate the target tissue without involving the camera. Direct stimulation issues commands to the stimulator, indicating which electrodes should be driven with what waveforms. This completely bypasses the camera and all video processing. If one can stimulate the center of vision or the center of the stimulating array, the user can point to where that phosphene appeared. Where the pointed location appears in the camera's FOV can then be chosen as the CAP. In concept, this ensures that the stimulator's FOV includes the location from which phosphenes appear to originate.

Ideally, setting the CAP using a direct stimulation procedure would involve both eye tracking and head-motion sensors. Movement of the eyes and/or head would naturally alter where phosphenes from direct stimulation would be perceived in egocentric space, as discussed in Section 2.5. If the camera is head-mounted, the effects of head motion at the time of stimulation would be less important, because head motion would affect the camera FOV as expected. In order to acquire the gestured location in the camera's FOV meaningfully, however, that location would still need to be adjusted for any movements of the camera that occur between time of stimulation and time of percept location measurement. Differences in eye orientation between trials at times of stimulation would need to be considered to account to for their effect on perceived phosphene locations. If the prosthesis system does not automatically

adjust the CAP in response to eye movements, indicated percept locations should be adjusted so the chosen CAP is most appropriate for the average eye orientation.

With Argus II system programming, neither eye nor head movement is accounted for in direct-stimulation camera alignment. Instead, users are asked to keep both head and eyes as still as possible during the alignment task. Unfortunately, one can still expect reflexive eye movements toward percept locations if phosphenes are not perfectly presented about the center of vision, and the eyes may not return to the exact same location afterward. Additionally, without any visual feedback regarding head movement, users can find keeping a constant head position difficult, particularly while responding or waiting for the operator. These eye and head movements can add substantial error to CAP estimatation.

In the current Argus II procedure, the user is seated in front a touchscreen monitor at 30–38 cm distance. The central 4 electrodes are activated with the same settings as though a visual stimulus appeared in that part of the stimulator FOV. The brightness of the simulated target, and therefore the amplitude of the current pulses, is set to be the minimum that is distinctly visible to the user. Larger configurations with more central electrodes are used if maximum current from the central 4 electrodes is insufficient. The user is asked to keep head and eyes as still as possible, and point on the touchscreen where each phosphene appears. Eight stimuli are presented, after which the user is asked to continue holding still. On the screen, the centroid of the 8 responses is highlighted and the programmer adjusts the CAP until the highlighted spot is centered in the stimulator's FOV.

Aside from error introduced by head and eye movements, the Argus II direct-stimulation alignment procedure is further hampered by the boundaries of the screen. If a phosphene appears to originate outside the area of the screen, the user is forced to respond at the screen's

edge. The centroid would appear at the screen's edge as well, and the stimulator FOV would be brought to center upon that. The CAP would more appropriately be adjusted to bring the FOV beyond the edge of the screen, but this cannot be done with any accuracy because of the lack of data.

4.2.2 Target localization

An alternative to direct-stimulation camera alignment asks users to point at actual visual targets. Seated in front a 37.47 × 30.00 cm touchscreen, the users are asked to scan the screen with the camera and touch the single spot of light on the display, as shown in Figure 4. Room lights are turned off so that the only bright source of light is the target on the screen. The visual target is a 2.9 cm wide square or circle, spanning about 4° of the user's visual field. Users are instructed to use only one hand for a trial run, usually the dominant hand, to avoid confounding data with different pointing bias from the other hand. Camera-screen distance is initially set to 36–38 cm, and users are asked to avoid moving their heads closer to or farther from the screen. Camera-screen distance was set before each trial run, and the same distance was maintained within-subject across all trial runs.

Figure 4 – Target localization task

An Argus II user uses his right hand to indicate the perceived location of the white circle. Here, the user perceived the circle to be lower than its actual position. Errors such as these can be used to guide selection of an appropriate camera alignment position.

Visual targets appear at random locations anywhere on the screen, save a margin extending 5.85 cm from each screen edge. This margin ensures that there will always be at least some room, 3x the radius of the target, for the user to make measurable errors in any direction, regardless of the target location. When the user touches the screen, the current target is removed and a new one appears elsewhere on the screen. The user is not usually given any feedback on the accuracy of the response, and knows to immediately start searching for the next target. Runs can consist of as few as 10 or as many as 100 trials, but 20 trials is typical.

Users may complete 5–10 trial runs in order to determine an appropriate CAP if no OCAP data were previously recorded.

The localization errors made within each trial run can be used to measure a user's average pointing bias. As shown in Figure 5, the errors can be collectively displayed relative to a normalized target. The error centroid relative to the target can be considered the average pointing bias. To correct for this bias, the CAP would be adjusted in accordance with the vector pointing from the target to the error centroid. For example, if the average pointing bias is 2° right and 5° up, the CAP would be adjusted 2° right and 5° up from its current setting. The new adjusted CAP, provided that the adjustment is not too eccentric for the system to use, would be the estimated OCAP based on these data. Bias measurement may not be exact, particularly with few trials, but biases can be expected to be reduced with 1°–2° after a few iterations.

Figure 5 – Data from target localization

The circle in the center is a normalized representation of all 20 targets presented in a sample run of the target localization task. Each green square represents a user's touch response location relative to the displayed target. The dashed ellipse indicates standard deviations along the horizontal and vertical axes. Units are in pixels, which were displayed with a resolution of about 34 pixels/cm. The user was decently accurate in this trial run, and any CAP adjustment would be slightly down and to the right.

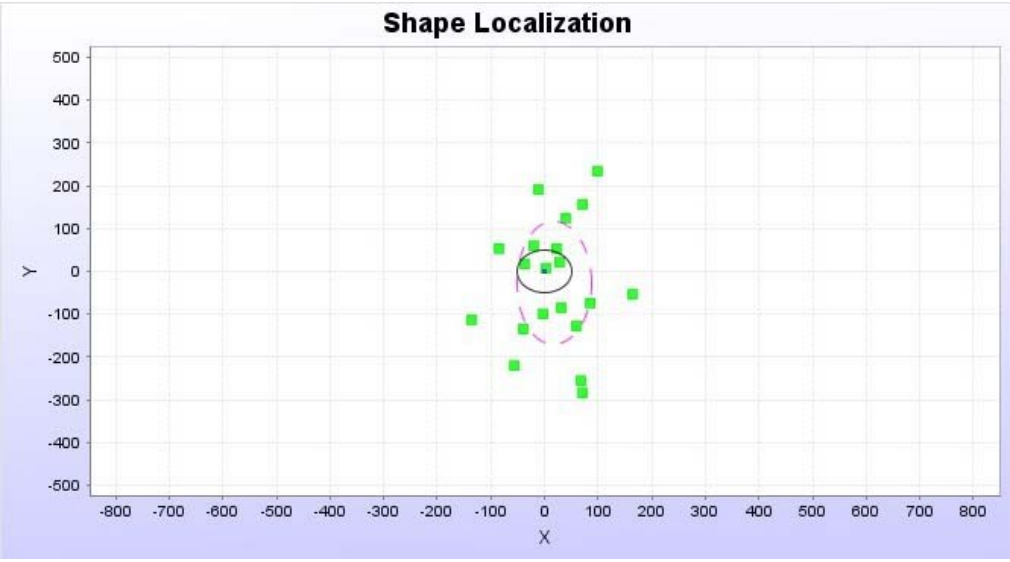

Unlike for direct-stimulation alignment, head and eye movements do not severely impede OCAP estimation. Head movement is in fact necessary for scanning the screen and finding the target. Eye movements do decrease localization precision, but can be expected to have symmetrically distributed effects, and thus not to severely impact accuracy. Any systematic effects of these movements may also be expected in everyday attempts at visual localization, and so should be incorporated in OCAP estimation.

Both camera alignment techniques suffer from limitations imposed by the size the screen. If the target ever appears to be outside the boundaries of the screen, the user's response is relocated to the edge of the screen. Just as this limits what CAPs can be set with direct stimulation, truncated localization errors interfere with OCAP estimation from target localization. Suppose an OCAP is 10° above a current CAP setting. Any target presented within the top quarter of the screen would be expected to have a smaller response error than those lower on the screen, because users would likely perceive higher targets above the screen and be forced to touch lower. The alignment procedure would still indicate the CAP needs to be adjusted upward, but the estimated OCAP would be lower than the true OCAP. Subsequent trial runs with smaller CAP-OCAP deviations would be necessary to more accurately estimate the user's OCAP.

4.3 Misalignment

Users of prostheses with remote cameras can expect some degree of misalignment before the initial camera alignment is performed. With default settings, such as mapping the optical center of the camera FOV to the center of the stimulating array, misalignment would likely manifest as a clear pointing bias in any localization task. Misaligned users may also incorrectly describe target locations, such as claiming that a light directly in front of the user is on his or her left side.

Camera alignment works to reduce these errors, but does not remove all localization error. Direct-stimulation and target-localization camera alignment techniques allow compensation for an average sustained misalignment, but provide no help regarding transient misalignments. Further, the dynamics of sustained misalignments and user OCAPs over time are largely unknown, and may present further challenges for optimizing prosthetic hand-camera coordination.

4.3.1 Transient misalignments

Most current visual prosthesis do not track eye movement. Thus, even if the camera is aligned with the eye when the eye is pointing in one direction, the eye and camera visual axes would be misaligned whenever the eye points in a different direction. The best configuration possible for a prosthesis with a static camera alignment setting would be to align the camera with the eye when the eye is in its neutral orientation. If the user can maintain the eye's neutral orientation during a task, the camera alignment will be maximally effective. Eye movements away from the neutral position might be expected to occur symmetrically over time, so average error should remain near 0. Error for any particular localization attempt could be high, however, depending on how the much the eye is deviated at the time of stimulation.

To avoid effects of transient misalignments, prosthesis users are instructed to move their eyes as little as possible. Unfortunately, devices with head-mounted cameras require head movement to scan scenes for visual information. Lacking any kind of fixation target that matches head movements, users still reflexively make eye movements that are driven by the vestibular system. Such misalignments of head and gaze directions are known to cause corresponding localization errors[222]. Users cannot avoid this complication by stopping head movements,

because static images conveyed from the camera to the stimulated tissue do not elicit persistent percepts. Instead, the percepts fade and are not refreshed until the scene changes.

Aside from discrete localization errors, transient misalignments can also induce illusions of visual target movement. Because camera input does not change unless the camera itself moves, any elicited percept remains unchanged in the visual field, except for fading effects, during an eye movement. Without movement across the visual field that is consistent with the eye movement, the percept seems to have moved with the eye in the direction of the eye movement[223]. After head scanning, for example, a user may stop the camera while the target is in view, but the saccade to recenter the eyes would induce the illusion of target motion. Illusory movement can also be perceived during head scanning, as reflexive eye movements may interfere with relating percept motion in the visual field to head motion.

4.3.2 Passive adaptation

Prior to the experiments detailed in Chapter 5, it was unknown whether users of visual prostheses could adapt to camera misalignments. In concept, one might expect users to realize that sources of light are consistently not where they appear to be. If a user reaches for a white cup on a black tablecloth, but reaches too far to the right, he or she should notice the grasping motion fails to touch the cup. The user could then tactilely search for the cup by moving the hand around table. Once found, the user should know that the cup was more left than he or she expected. Repeated reaching attempts such as this could be expected to impress upon the user that objects are always more to the left than they appear. Reaching movements could be adjusted accordingly, and camera alignment would be functionally unnecessary.

Such passive adaptation, in contrast to active training to compensate for pointing and reaching biases, could make prosthesis programming much simpler. If a user can adapt to the default camera alignment, there would be no need to train the user on dedicated camera alignment tasks and spend time for those tasks at programming centers. Users would only need to be made aware of the misalignment, and that it may take time for them to adjust. If needed, rehabilitation specialists could incorporate corrections to localization errors into their standard regimens.

Based on the work with prism glasses discussed in Section 2.5, one might even expect adaptation to occur quite rapidly. During training periods for Gibson's[135] experiments, visual percepts were shifted 15° to the right in subjects' visual fields, and subjects were able correct localization errors within an hour. Initially, subjects would point or reach to the right of target objects, or walk too far to the right and hit objects or door frames. As they practiced, however, errors gradually diminished until subjects appeared to perform normally.

Expecting such adaptation from users of visual prostheses, however, assumes a much greater quality of vision than prostheses currently generate. The visual percepts experienced by users tend to be very crude, and seldom if ever provide enough detail to create a structured visual field[224]. Orientation using current visual prostheses typically requires some other contextual knowledge, such as the user knowing that he or she is in a room with a bright window. There would not be enough visual information for a user to readily detect conflicts with other sensory inputs. A user may feel that visual input does not make sense give other available information, but that would mean very little if the visual percepts don't make sense on their own anyway.

Reafference is also limited, as there is usually insufficient detail and contrast for the user to recognize parts of his or her own body, such as the hand. Luo et al. attempted to circumvent this

problem by applying a flashing LED to a user's finger, but did not find any benefit from this in their testing scheme[225]. Spatial and temporal dynamics of stimulation may not allow percepts to move with clear relationships to motor commands, so the link between motor output and visual input is lost. Without the ability to see their own movements through the perceptual shift created by camera misalignment, intuitive adaptation to misalignment may not be possible for prosthesis users.

Much of how users perceive stimulation from visual prostheses is not well known, and the potential for adaptation to camera misalignments rests soundly on the quality of those percepts. It is understood that prosthetic vision is crude, but could it convey necessary morsels of information to correct localization errors? Three Argus II users were given camera misalignments of 15° or more to investigate this possibility. The misaligned users were observed for 12 months and localization data were analyzed for indications of adaptation. Method details and results are discussed in Chapter 5.

4.3.3 Constancy of optimal camera alignment position

Complementary to the question of whether a visual prosthesis user can adapt to camera misalignment, one can ask how long a proper camera alignment remains valid. Current programming and rehabilitation procedures call for camera alignment whenever new equipment is issued, but do not involve checking alignment suitably or realigning the camera afterward. This as assumes that camera alignment is only needed once for any given set of equipment, or at least that realignment should only be necessary if the user notices an unexpected misalignment.

This assumption of a constant OCAP, however, was never investigated or confirmed. OCAPs are not generally tracked, or even measured beyond the what is necessary to set a user's CAP. If an OCAP is outside the range of the camera, the closest possible CAP is all that is set and recorded. It is possible that OCAPs, without consideration for transient eye movements, do not change. Once an appropriate CAP is set, one might not expect any change in OCAPs if users are able to adapt to misalignments. Adaptation would effectively anchor the OCAP to the current CAP. If OCAPs are indeed quite stable, the policy of setting the CAP once and never remeasuring the OCAP would be appropriate and efficient.

At the same time, it is possible that neutral, or resting, eye orientations may change over time. Proprioceptive relations between the eye and the body's trunk or motor pointing biases may also display time-dependent variation. If users cannot adapt to misalignments, any such changes could interfere with the appropriate mapping of percepts in egocentric space and/or conducting accurate movements based on visual stimuli. The OCAP would therefore change, and a user's CAP would need to be adjusted accordingly.

A number of studies emphasize the potential use of visual prostheses for hand-eye or hand-camera coordination[194, 195, 225-233]. Under the right circumstances, prosthetic vision can be an effective tool for tasks that require visual guidance of motion. In a laboratory setting, camera misalignments can be readily detected and corrected to optimize user performance. At home or work, however, camera misalignments can limit the utility of a prosthesis for coordination tasks, and realignment would not be possible until the next visit to a programming center. If OCAPs are not constant and users require regular realignment, real-world benefit for prostheses for hand-camera coordination would be less than is apparent.

As the ability to accurately localize light sources, obstacles, and objects of interest is critical for wayfinding and hand-camera coordination, investigating OCAP dynamics will be important for the functional success of visual prosthesis users. Toward this end, OCAPs of 2 Argus II users were measured regularly for 5 years and examined for significant changes over time. Method detail and results are given in Chapter 6.

5 Passive adaptation to misalignment

As discussed in Section 4.3.2, it is not well known if and how visual prosthesis users might adapt to camera misalignments. This is mostly an issue for devices that use cameras or sensors physically separated from the stimulator. If users can passively adapt to misalignment, such that the OCAP moves to the default or other current CAP over time without specific training, camera alignment tasks may not be necessary for standard programming. If passive adaptation does not occur, however, camera alignment is the only way to improve intuitive hand-camera coordination. Lack of adaptation would also leave open the possibility of OCAP changes over time that are not anchored to the current CAP, which would necessitate regular camera realignment.

Users of visual prostheses, such as the Argus II, likely do not receive the necessary visual information to view their bodies or their actions. Knowing that reafference is required for adaptation to prism-shifted vision, one can expect users with misaligned cameras to not display the same kind of adaptation known to occur with prism glasses. In the experiment in detailed in this chapter, 3 Argus II users were given misaligned cameras for over a year. Users were tested using target localization both with and without auditory feedback. Assuming that users had no visual indications of misalignment, feedback-OFF tests were almost entirely open-loop. Open-loop testing, in which subjects provide responses to stimuli but have no feedback regarding the accuracy of their responses, allowed OCAPs to be measured without any opportunity for learning or correction. Implicit feedback could still occur, for example, if a subject perceived an object on a kitchen table at home, reached for the object, and realized that it was not in the perceived location. If, however, no aspects of home use provided the information necessary to alert subjects of localization errors, and thus drive adaptation to misalignment, open-loop

feedback-OFF tests should show users' lack of adaptation most clearly. The analyses in this chapter tested the hypothesis that users could not demonstrate significant reduction in localization errors within 5–7-month timespans. The results indicate that that users of current visual prostheses, at least when not given constant auditory feedback, indeed cannot passively adapt to camera misalignments. When OCAP movements are observed, related to adaptation or not, OCAP changes between days are biased toward the directions of intraday OCAP trends.

5.1 Methods

5.1.1 Subjects and equipment

Three Argus II users participated in this experiment. All 3 subjects were enrolled in the Argus II Feasibility Study. S1 was female and right-hand dominant. S2 was male and also right-hand dominant. S3 was male and left-hand dominant. All 3 subjects were implanted in the right eye.

Two subjects, S1 and S2, were implanted in June 2007. S3 was implanted in June 2009. S3 was implanted with an updated electrode array design that was slightly different from those implanted in S1 and S2. The slotted array used for S1 and S2 covered a retinal area spanning 19.62° × 14.85°. The dual-metal array used for S3 had the same number of electrodes, but was slightly smaller, covering 17.92° × 10.75°.

Each subject's electrode array was rotated relative to the retina however the surgeon felt appropriate for implantation. S1's array was rotated +14°, where positive rotation is seen as counterclockwise from the fundoscopic perspective. S2's array was not rotated away from 0°, while S3's array was rotated −44°. Subject's cameras were rotated accordingly so apparent directions of motion were accurate.

Argus II cameras used for this study, version A2E8, had FOVs spanning 66° × 49°. CAPs could be deviated from the camera center at (0°, 0°) by step sizes of approximately 0.27°. Steps could be made along the camera's horizontal and vertical axes, which were rotated to match the electrode array's orientation. The camera's horizontal and vertical axes were thus only parallel to those of the visual field when the array rotation was 0°. CAPs for both the slotted and dual-metal arrays could be adjusted ± 24.18° along the camera horizontal axis, and ± 18.29° along the vertical axis.

5.1.2 Study procedures

OCAPs were estimated using target localization as described in Section 4.2.2. S1 and S2 were seated so the touchscreen was always 36 cm away from the camera before each trial run, while S3's camera-screen distance was always 38 cm. Targets were predominantly 4°-wide squares. No margins were employed in the test software, so targets could appear anywhere on the screen. Subjects were permitted to use either hand interchangeably during this study, and use of only one per trial run was not enforced. Prior to and through the first 6 months of the study, corrective auditory feedback was provided for every target localization trial. If the subject touched the target, the computer announced, "Correct." If the subject touched outside the border of the target, but within about 1 cm, the computer would say, "Close," followed by the relative direction of the target from the subject's touch. For example, if the subject touched about 1 cm above the target, the computer would say, "Close, it was lower." If the subject touched farther from the target, the feedback would only provide the relative direction of the target. For example, "It was lower." Subjects performed as many trials per session as time allowed. Session trial counts ranged 7–880, with a median of 240.

Starting in December 2012 (S1 and S2) and January 2013 (S3), subjects were given CAPs deviated by 13°–40° from their OCAPs. Session initial OCAPs (see Section 5.1.3) measured just before misalignment and misaligned CAPs are shown in Table 1. OCAP estimates were made using aligned CAPs, and misaligned CAPs were applied later in the same session (S1 and S3) or at the beginning of the next session (S2). Once misaligned settings were applied, they were not changed until after the completion of the experiment in 2014. Misaligned CAPs where therefore used both in lab and home for entire duration of the study. Subjects were not explicitly made aware of misalignments, but were told that some settings were changed and that objects may not be where they appear to be. Subject were asked to notify testers if anything was problematic or settings need to be changed. Subjects ultimately did not report any related problems or abnormal observations, and misalignments were apparently not noticed.

Table 1 – Camera misalignments

Subject CAPs were misaligned from OCAPs so any adaptation effects could be observed. CAPs and OCAPs are shown after converting rotated-array coordinates to standard, unrotated coordinates.

Subject	OCAP (°)		Misaligned CAP (°)		Total misalignment (°)
	Horizontal	Vertical	Horizontal	Vertical	
S1	−20.37	−16.16	19.89	−13.19	40.37
S2	−18.46	9.56	0.00	7.66	18.56
S3	−2.13	−2.17	7.17	7.17	13.18

Subjects attended testing sessions every 1–2 weeks as they were able. Each session consisted of up to 18 runs with medians of 40–80 trials per run. Corrective auditory feedback was provided for each trial until July 2013. Afterward, to assess the impact of in-lab auditory on adaptation, subjects were asked to perform target localization as they had previously, but without any corrective feedback. Testing continued until December 2013 (S1) and January 2014 (S2 and S3).

Subjects thus performed target localization with auditory feedback for 6–7 months, and without feedback for 5–6 months.

5.1.3 Analysis of adaptation

Estimated OCAPs appeared to gradually shift within sessions as time progressed. Linear models were computed with respect to time using ordinary least squares (OLS), performed independently along the horizontal and vertical axes. The model ordinate values at the session start time were considered the estimates for that day's initial OCAP horizontal and vertical coordinates. Session initial OCAPs, as well as distances of error centroids from the normalized target, were tracked across days to investigate adaptation to misalignment. Initial OCAP estimates were favored over session averages to reduce the influence of trial count on estimates.

To visualize OCAP movement, initial OCAP estimates were plotted on the visual field along the subject's misaligned CAP. Progression of OCAPs was indicated with arrows, and each marked point represented the average of initial OCAPs within some time bin. Small time bins with 0–2 sessions per bin were first used to demonstrate interday variability of OCAPs. Large time bins with up to 5 or 7 sessions then showed OCAP movement with less noise. 95% confidence regions for time-bin OCAP averages were determined by bootstrap resampling all OCAP estimates within date-subject strata 10^4 times and subsequently recalculating intraday linear models and time-bin means. The centers of bootstrap means were calculated for each time bin, and the most-distant 5% of bootstrap means from each center were discarded. The convex hull around the remaining 95% of bootstrap means was considered the point's 95% confidence region.

The precision of subject touch responses was separately evaluated between February 2016 and January 2017 with S2 and S3. S1 could not be scheduled for these tests. In a separate setup, subjects rested their heads on a chinrest without the Argus II system or any visual stimuli. When told to start, subjects performed a nose localization task by moving a finger between the nose and the point on a touchscreen directly in front of the nose. Subjects repeated this process 20 times per trial run. Each session involved about 2 runs with each hand, and subjects attended sessions every 1–2 weeks as they were able. Means of within-run variances of touch-responses were used to estimate limits of subject horizontal and vertical touch precision. S1's touch precision was estimated relative to the those of S2 and S3 considering relative touch-response variances in target localization.

For analysis purposes, feedback-ON and feedback-OFF analyzed separately, minimal adaptation was considered as reduction in the average distance between the error centroid and the normalized target by $1.96 \times SD_{touch}$, where SD_{touch} is the standard deviation (SD) of the subject's touch responses. This minimum was chosen assuming that random touch errors are normally distributed, and the probability of the error centroid distance falling by that much would be less than 5%. Using this error-reduction minimum, a criterion slope for error centroid distances plotted over time was calculated:

$$slope_{criterion} = \frac{-1.96 \times SD_{touch}}{day_{end} - day_{start}}$$

The line with the criterion slope passing through the centroid of plotted points was considered a regression line that would demonstrate minimal adaptation. Residuals from this hypothetical adaptation regression line were calculated, and then resampled with replacement 10^4 times to create bootstrap data distributions about the line of minimal adaptation. OLS regression lines

were determined for each bootstrap distribution. This provided a sampling distribution for regression line slopes, assuming that minimal adaptation was actually observed. The slope of the OLS regression line for the actual data was compared to this distribution. An observed slope greater than the top 5th percentile, thus having a probability less than 5% of being observed by chance, was considered a significant lack of adaptation.

Power to distinguish the observed slope as being greater than the criterion slope was determined by bootstrapping residuals along the observed regression line. As for the bootstrap analysis of the minimal adaptation line, OLS regression lines were calculated for each bootstrap distribution about the observed line. The percentage of bootstrap slopes that were greater than the criterion slope was taken as the test power.

5.1.4 Intraday-interday OCAP directional analyses

Intraday models of OCAP changes over time were also used to investigate the relationship between intraday OCAP changes and OCAP changes between sessions. Each intraday OCAP trend was compared to that of the immediately previous session using 3 vectors in the horizontal/vertical (H/V) plane: (1) the projection of the prior session's intraday model on H/V plane, (2) an interday baseline vector from the beginning of the prior intraday model to the beginning of the current, and (3) an interday reset vector from the end of the prior intraday model to the beginning of the current. Relationships among these vectors were calculated within subject, and then pooled for analysis.

Directions of the 3 vectors were considered independently of magnitudes. For each analysis, consider paired datasets X and Y that each contain n vectors. Let $\boldsymbol{x}_i \; \forall \; i \in [1, n]$, be a vector in X, and $\boldsymbol{x}_{i\theta}$ be the direction of $\boldsymbol{x}_i$. Similarly, let X_θ be the set of vector directions belonging to X.

Differences between directions $\boldsymbol{x}_{i\theta} - \boldsymbol{y}_{i\theta}$ were plotted on unit circles. Each plot point can be considered the end of a unit vector $\boldsymbol{d}_i = (1, \boldsymbol{x}_{i\theta} - \boldsymbol{y}_{i\theta})$ that starts at the circle center. Let D be the set $\{\boldsymbol{d}_1, \dots, \boldsymbol{d}_n\}$. The circular mean of D_θ is calculated by computing the mean vector $\overline{\boldsymbol{d}}$ of all $\boldsymbol{d}_i$ in D and taking its direction:

$$\overline{\boldsymbol{d}} = \frac{1}{n}\sum_{i=1}^{n} \boldsymbol{d}_i$$

$$mean(D_\theta) = \overline{\boldsymbol{d}}_\theta$$

Circular correlation coefficients were calculated by using an analog of Pearson's correlation coefficient based on sines of deviations:

$$corr(X_\theta, Y_\theta) = \frac{\sum_{i=1}^{n}(\sin(\boldsymbol{x}_{i\theta} - \overline{\boldsymbol{x}}_\theta)\sin(\boldsymbol{y}_{i\theta} - \overline{\boldsymbol{y}}_\theta))}{\sqrt{\sum_{i=1}^{n}\sin(\boldsymbol{x}_{i\theta} - \overline{\boldsymbol{x}}_\theta)^2}\sqrt{\sum_{i=1}^{n}\sin(\boldsymbol{y}_{i\theta} - \overline{\boldsymbol{y}}_\theta)^2}}$$

$\overline{\boldsymbol{d}}_\theta$ confidence intervals and probabilities of observing given variances and correlation coefficients were determined using at least 10^3 bootstrap samples. $\boldsymbol{d}_{i\theta}$ were resampled with replacement to construct each bootstrap sample for $\overline{\boldsymbol{d}}_\theta$'s confidence interval. For variances and correlation coefficients, each $\boldsymbol{x}_{i\theta}$ was randomly paired with some $\boldsymbol{y}_{j\theta}, j \in [1, n]$, to create bootstrap samples with no correlation between X_θ and Y_θ.

Let $\boldsymbol{D}_\theta$ be a random variable of which each $\boldsymbol{d}_{i\theta}$ is a sample. The distribution of $\boldsymbol{D}_\theta$ was tested for circular uniformity using Watson's one-sample test. Watson's test was also used to determine the probability that the underlying distribution of $\boldsymbol{D}_\theta$ was a von Mises distribution. Note that a von Mises distribution with a concentration parameter κ of 0 is equivalent to a uniform distribution.

Nonuniform distributions of $\boldsymbol{D}_\theta$ with $\overline{\boldsymbol{d}}_\theta$ confidence intervals confined to the 0° semicircle were considered to indicate significant similarity of X_θ and Y_θ. Nonuniform distributions with $\overline{\boldsymbol{d}}_\theta$ confidence intervals confined to the 180° semicircle were taken to indicate a reset effect, where the progress of the first intraday trend is somewhat reversed before the start of the second intraday trend.

5.2 Results

5.2.1 OCAP changes with camera misalignments

In the plots below, points on the visual field represent initial OCAP averages within a given timeframe, and arrows indicate the sequence of OCAPs over time. OCAP positions for timeframes with no data were interpolated and are shown as small black triangles. Dotted arrows connect two periods that have no data, and simply preserve the representation of time. Sessions with auditory feedback are shown in red, and those without in blue. Sessions without feedback took place after all sessions with feedback were completed. The constant, misaligned CAPs are shown as solid green circles.

5.2.1.1 Small time bins

Data were first plotted showing initial OCAP averages within small time bins, bin sizes equal to each subject's typical time between test sessions. Time bins thus only contained only 0–2 test sessions. OCAP variability was typically too high to see clear trends after averaging within such small time bins.

Figure 6 – OCAP changes, misaligned camera, S1, 7-day time bins

Points represent average initial OCAPs within 7-day time bins. Each time bin contained 0–2 testing sessions. Variability was very high with these small time bins.

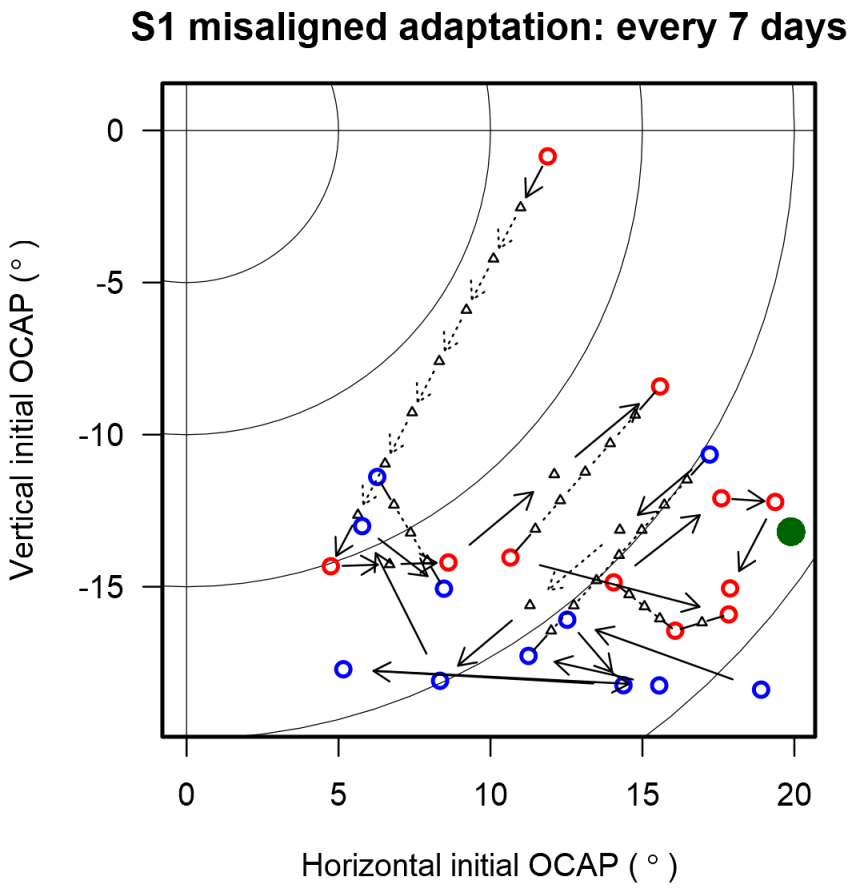

Figure 7 – OCAP changes, misaligned camera, S2, 14-day time bins

Points represent average initial OCAPs within 14-day time bins. Each time bin contained 0–2 testing sessions. Variability was very high with these small time bins.

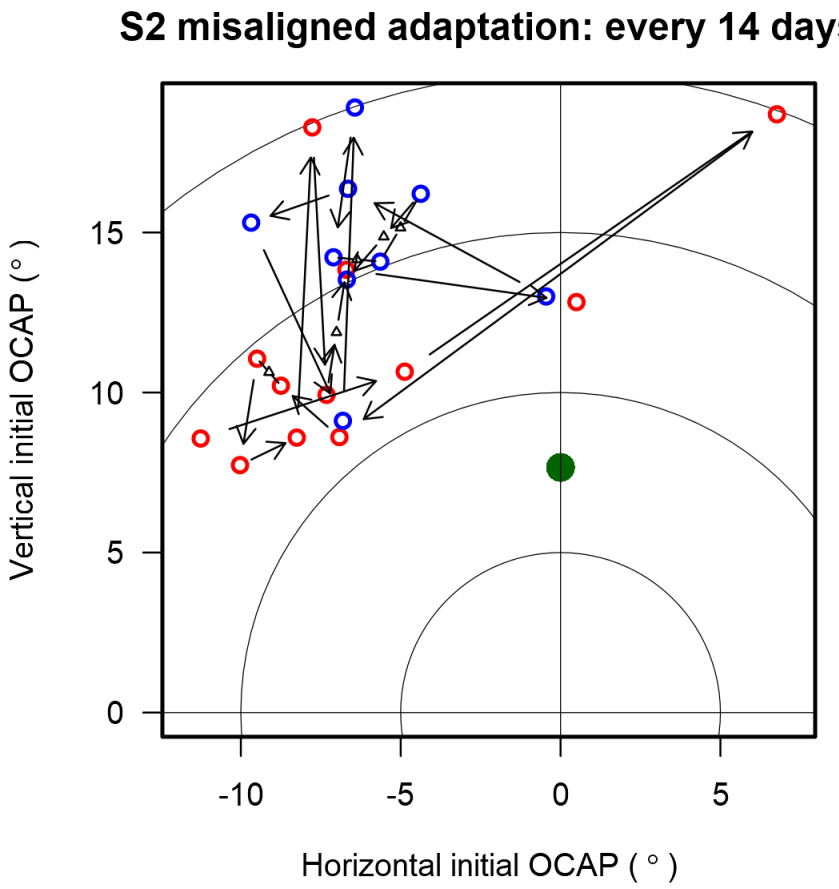

Figure 8 – OCAP changes, misaligned camera, S3, 7-day time bins

Points represent average initial OCAPs within 7-day time bins. Each time bin contained 0–2 testing sessions. Variability was very high with these small time bins.

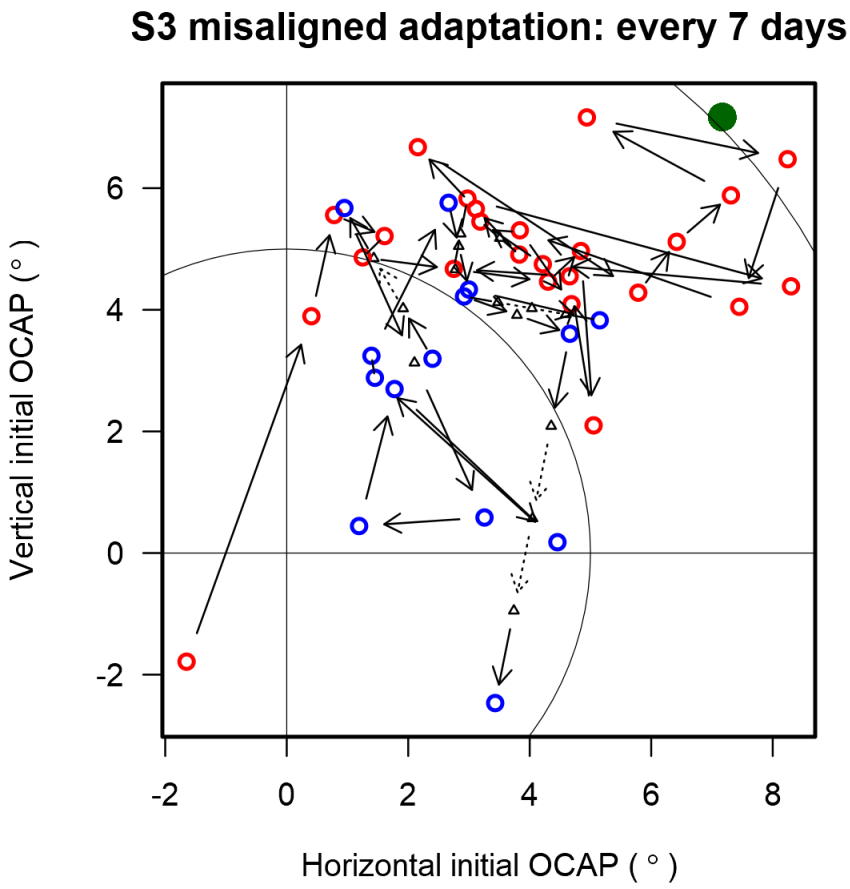

5.2.1.2 Large time bins

Larger time bins for averaging initial OCAPs were used to reduce noise. For S1 and S3, noise appeared sufficiently reduced after increasing time-bin size by a factor of 5. S2, whose responses were inherently more variable and who participated in less frequent testing, required increasing time-bin sizes by a factor of 7. 95% confidence regions for time-bin OCAP averages were drawn using dashed lines for each time bin.

Figure 9 – OCAP changes, misaligned camera, S1, 35-day time bins

Points represent average initial OCAPs within 35-day time bins. Each time bin contained 1–5 testing sessions. S1's OCAPs gradually moved towards the constant, misaligned CAP when auditory feedback was provided. Without auditory feedback, OCAPs gradually moved back away from the constant CAP.

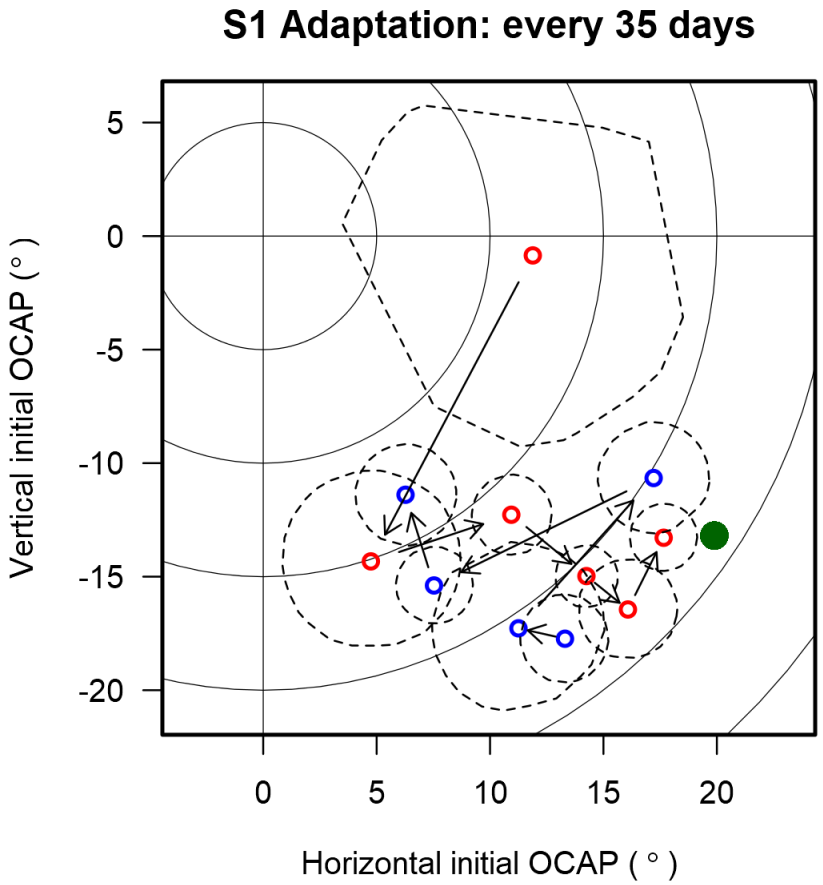

Figure 10 – OCAP changes, misaligned camera, S2, 98-day time bins

Points represent average initial OCAPs within 98-day time bins. Each time bin contained 4–7 testing sessions. S2's OCAPs did not appear to move toward or away from the misaligned CAP with feedback conditions. Feedback-on OCAPs were generally lower and closer to the misaligned CAP feedback-off OCAPs. Feedback may have had a constant, but not additive, effect on S2's responses.

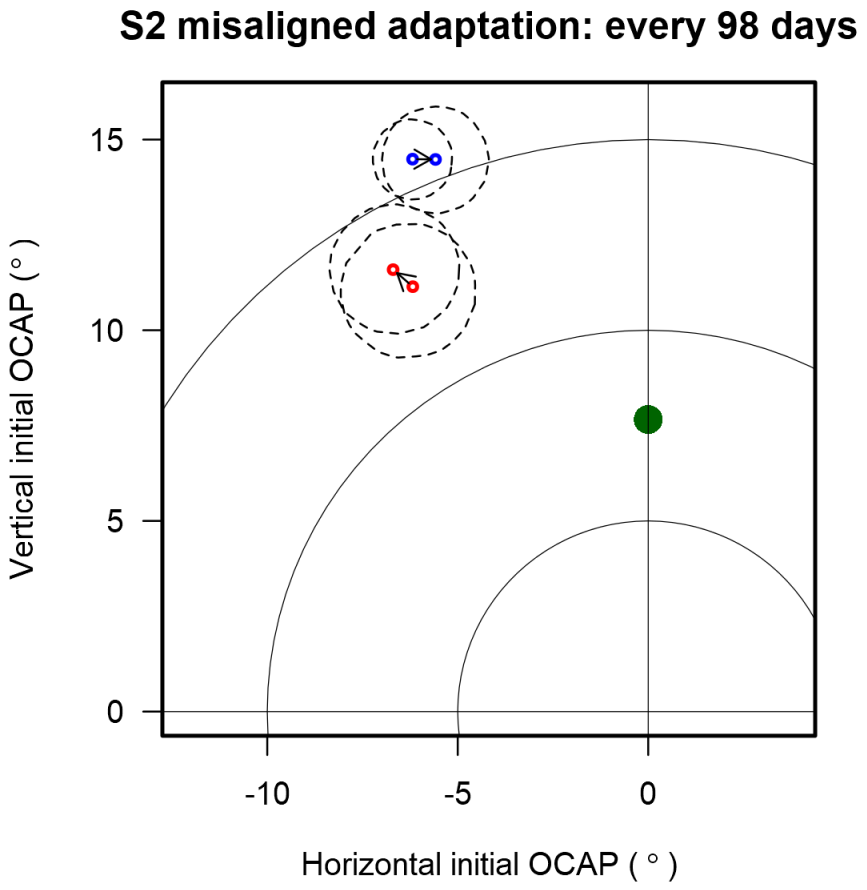

Figure 11 – OCAP changes, misaligned camera, S3, 35-day time bins

Points represent average initial OCAPs within 35-day time bins. Each time bin contained 1–6 testing sessions. S3's OCAPs gradually moved towards the constant, misaligned CAP when auditory feedback was provided. Without auditory feedback, OCAPs almost immediately returned to pre-adaptation positions, and then moved a pattern consistent with aligned-CAP measurements, with no noticeable effect of misalignment.

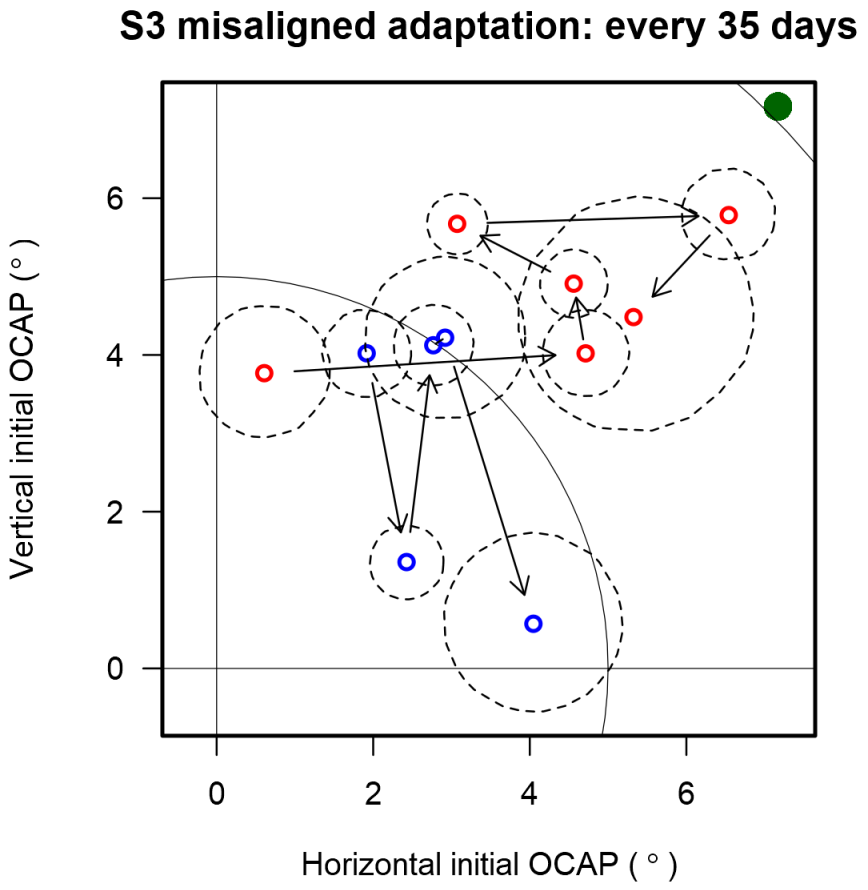

5.2.2 Touch precision

SDs of touch responses along horizontal and vertical dimensions, with and without head motion and visual targets, are shown in Table 2. Some visual target localization and all nose localization runs were performed with each hand separately, but one hand was never consistently more precise than the other. Data for both hands were therefore pooled within-subject for analysis.

Because S1 could not perform the nose localization task, that subject's SD of touch responses without visual targets was estimated based on subjects' relative performance with visual targets. In the normal target localization task, as used for monitoring OCAP changes with aligned and misaligned targets, S1's SDs of localization errors relative to the normalized target were between those for S2 and S3. S1's horizontal and vertical SDs of localization errors were each within 0.2° of the mean of those for S2 and S3. Accordingly, the SDs of touch responses in nose localization for S1 were estimated as the mean of those for S2 and S3.

Table 2 – Touch-response standard deviations

Standard deviations for touch responses when localizing visual targets, and when touching a screen in front of the nose without any visual stimulation. For visual targets, response SD was calculated by taking the SD of localization errors with respect to the normalized target. *S1 did not perform nose localization, and the displayed SDs were estimated by taking the mean of those for S2 and S3.

Subject	Standard deviation of touch responses (°)				
	Visual target localization		Nose localization		
	Horizontal	Vertical	Horizontal	Vertical	Mean
S1	7.0	6.7	1.6*	2.0*	1.8*
S2	8.7	9.0	2.2	2.7	2.5
S3	5.0	4.8	1.0	1.3	1.2

5.2.3 Localization error reduction

Improvement in accuracy was the focus of these analyses, so distances of error centroids were favored over individual error magnitudes for tracking error reduction. Tables show the observed OLS regression line slope of error centroid distances over time and the adaptation criterion for each subject. In plots, points each represent the error centroid distance from the origin for one trial run. OLS regression lines for centroid distances vs. experiment day, starting with each subject's first day with the camera misalignment, are shown in black.

Mean nose localization response SDs in Table 2 were used to determine necessary slopes for minimal adaptation. Visual target localization SDs could have been used, but these would have imposed stricter requirements for minimal adaptation, and made significant lacks of adaptation more likely to be observed. Solid blue lines in plots show the criterion rates of improvement and pass through dataset centroids. Dashed blue lines mark the upper extent of the criterion line's 95% slope CI.

5.2.3.1 Auditory feedback enabled

During the first 6–7 months, all subjects showed some reduction in error centroid distances. S2, however, only had a slightly negative trend and still demonstrated significant lack of adaptation. S1 and S3 reduced centroid distances more quickly than required for minimal adaptation, so adaptation could not be ruled out. Table 3 provides numerical details for each subject, and the following figures show the progression of centroid distances over time.

Table 3 – Feedback-ON error reductions and adaptation criteria

S2 demonstrated significant lack of adaptation when auditory feedback was enabled, whereas S1 and S3 did not.

Subject	Slope (error °/day)		Probability observed ≤ criterion	Test power
	Observed	Adaptation criterion		
S1	−0.05	−0.02	0.99	$< 10^{-4}$
S2	−0.004	−0.03	10^{-3}	0.94
S3	−0.02	−0.01	1.0	$< 10^{-4}$

Figure 12 – Localization error changes, misaligned camera, auditory feedback ON, S1

Points each represent the error centroid distance from the origin for one trial run. The OLS regression line slope is more negative than the criterion slope, and therefore also its upper 95% confidence limit, so adaptation cannot be ruled out.

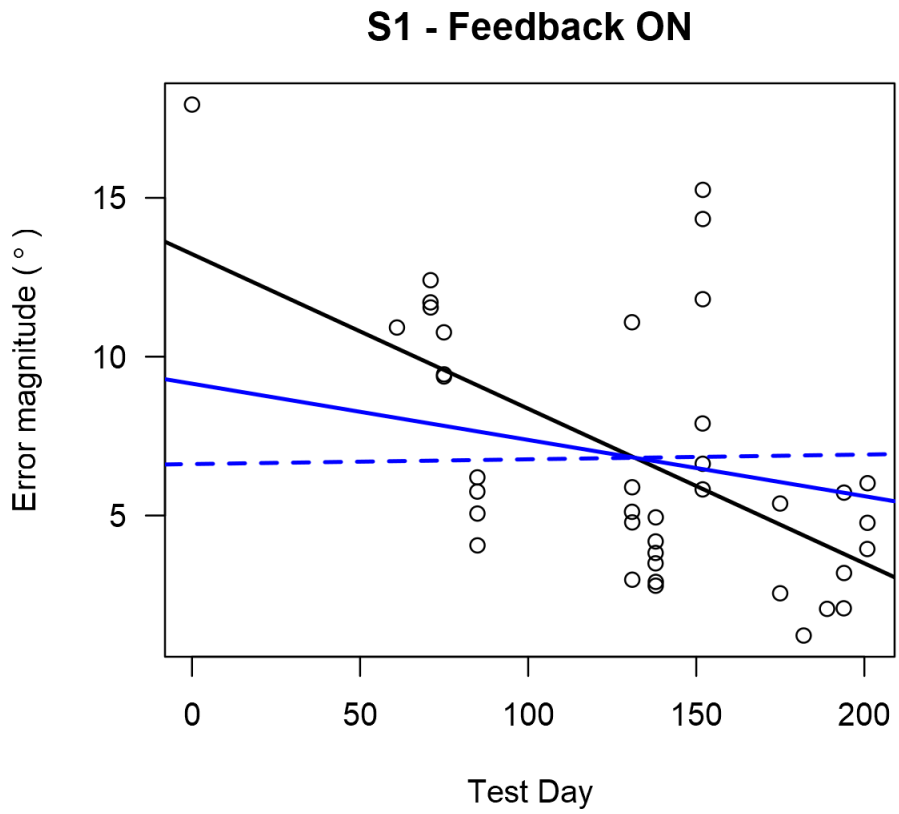

Figure 13 – Localization error changes, misaligned camera, auditory feedback ON, S2

Points each represent the error centroid distance from the origin for one trial run. The OLS regression line slope is less negative than both the criterion slope and the criterion slope's upper 95% confidence limit. S2 displayed a significant lack of adaptation to the camera misalignment during this period.

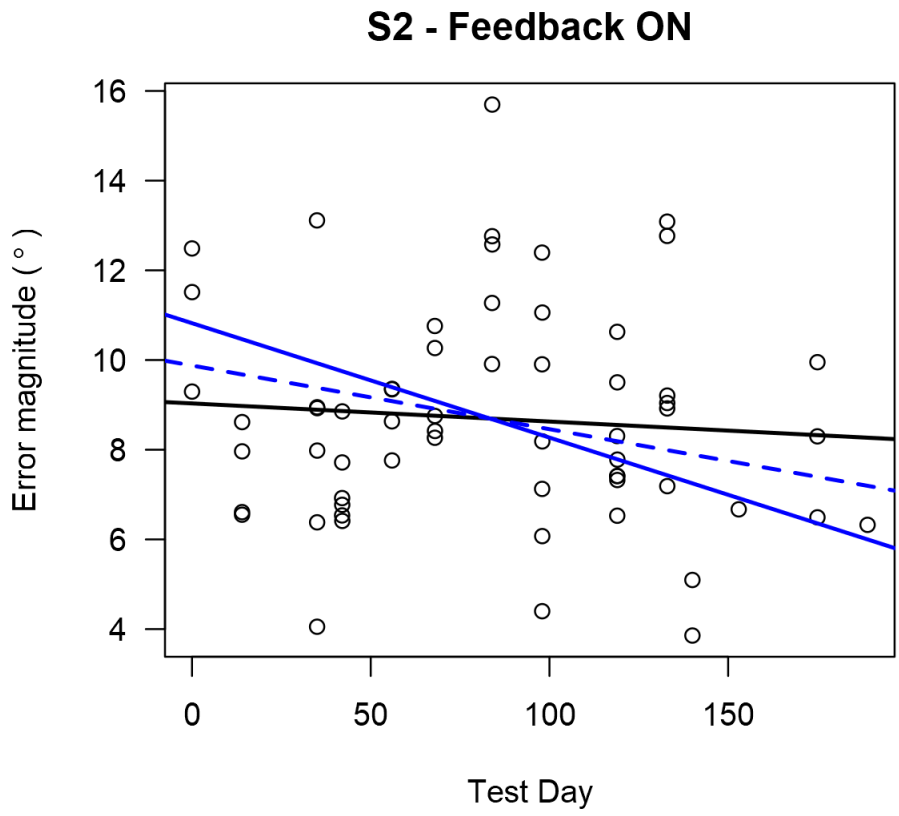

Figure 14 – Localization error changes, misaligned camera, auditory feedback ON, S3

Points each represent the error centroid distance from the origin for one trial run. The OLS regression line slope is more negative than the criterion slope, and therefore also its upper 95% confidence limit, so adaptation cannot be ruled out.

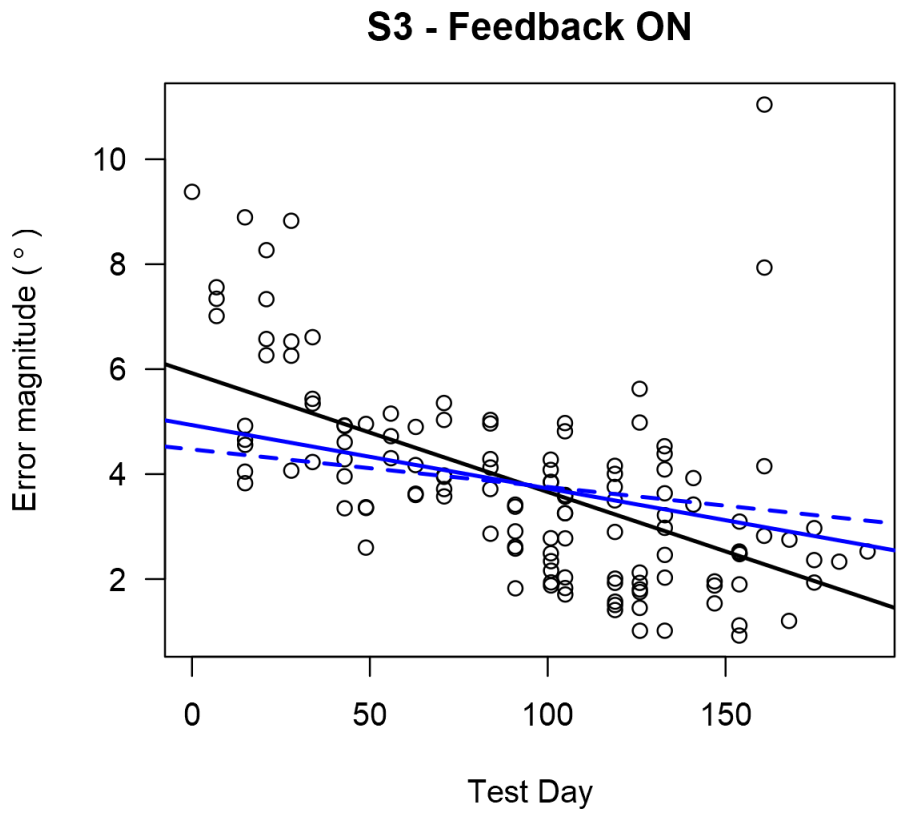

5.2.3.2 *Auditory feedback disabled*

Unlike in the feedback-ON period, all subjects showed significant lack of adaptation in the feedback-OFF period. Only S3 demonstrated some modest reduction in error centroid distances, but this was not large enough to exceed the bounds of the subject's touch-response CI.

Table 4 shows rate details and results of statistical tests. Test day numbering in plots is continuous with numbering in the corresponding plots of Section 5.2.3.1.

Table 4 – Feedback-OFF error reductions and adaptation criteria

All subjects demonstrated significant lack of adaptation when auditory feedback was disabled.

Subject	Slope (error °/day)		Probability observed ≤ criterion	Test power
	Observed	Adaptation criterion		
S1	0.02	−0.02	2×10^{-4}	0.99
S2	0.01	−0.03	$< 10^{-4}$	1.0
S3	−0.007	−0.01	0.03	0.56

Along with rates of error reduction, centroid distances in the feedback-OFF period were compared to final values of the linear models in the feedback-ON period. The ordinate value of the OLS regression line on last day of the feedback-ON period for the same subject is shown in each of the following plots with an orange horizontal line. The 95% CI of the final model point, as determined by bootstrapping observed residuals and recalculating the feedback-ON model, is marked by dashed orange lines. Models in the feedback-OFF period were generally higher than the feedback-ON endpoint CIs. S3, who barely showed a significant lack of adaptation, notably never had any feedback-OFF centroid distance data points within the feedback-ON endpoint CI.

Figure 15 – Localization error changes, misaligned camera, auditory feedback OFF, S1

Points each represent the error centroid distance from the origin for one trial run. The OLS regression line slope is positive, and thus necessarily greater than the criterion slope and its upper 95% confidence limit. S1 showed significant lack of adaptation after auditory feedback was removed. Further, error centroid distances, even at the beginning of the feedback-OFF period, were significantly greater than those at the end of the feedback-ON period.

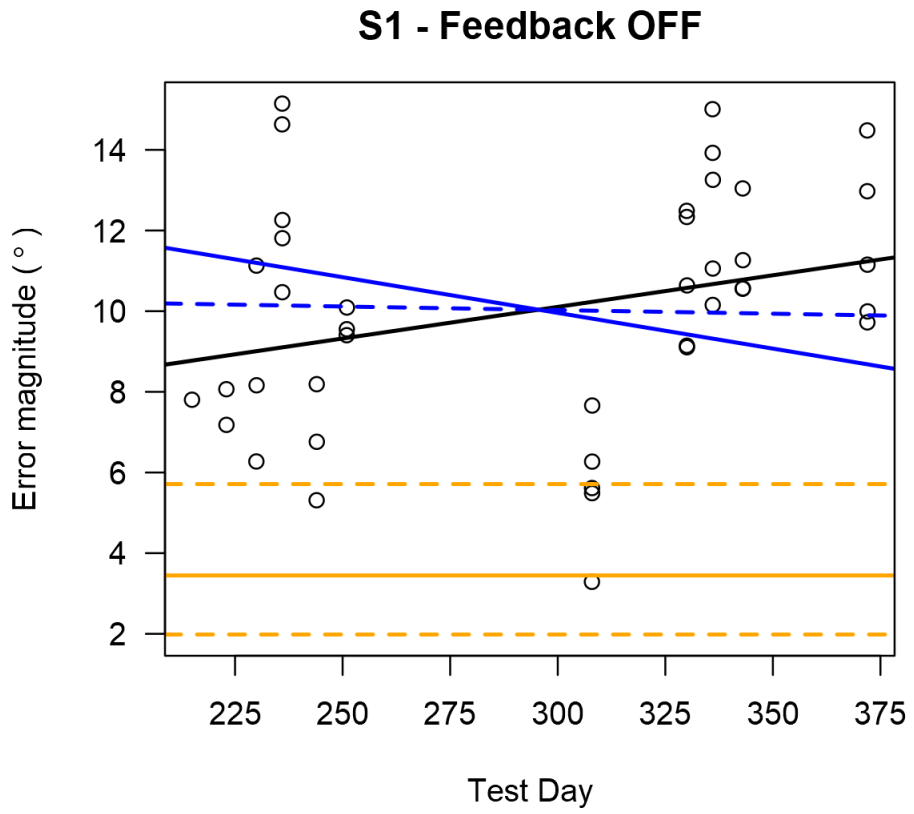

Figure 16 – Localization error changes, misaligned camera, auditory feedback OFF, S2

Points each represent the error centroid distance from the origin for one trial run. The OLS regression line slope is positive, and thus necessarily greater than the criterion slope and its upper 95% confidence limit. S2 showed significant lack of adaptation after auditory feedback was removed. Error centroid distances at the beginning of the feedback-OFF period were within the 95% CI of those at the end of the feedback-ON period.

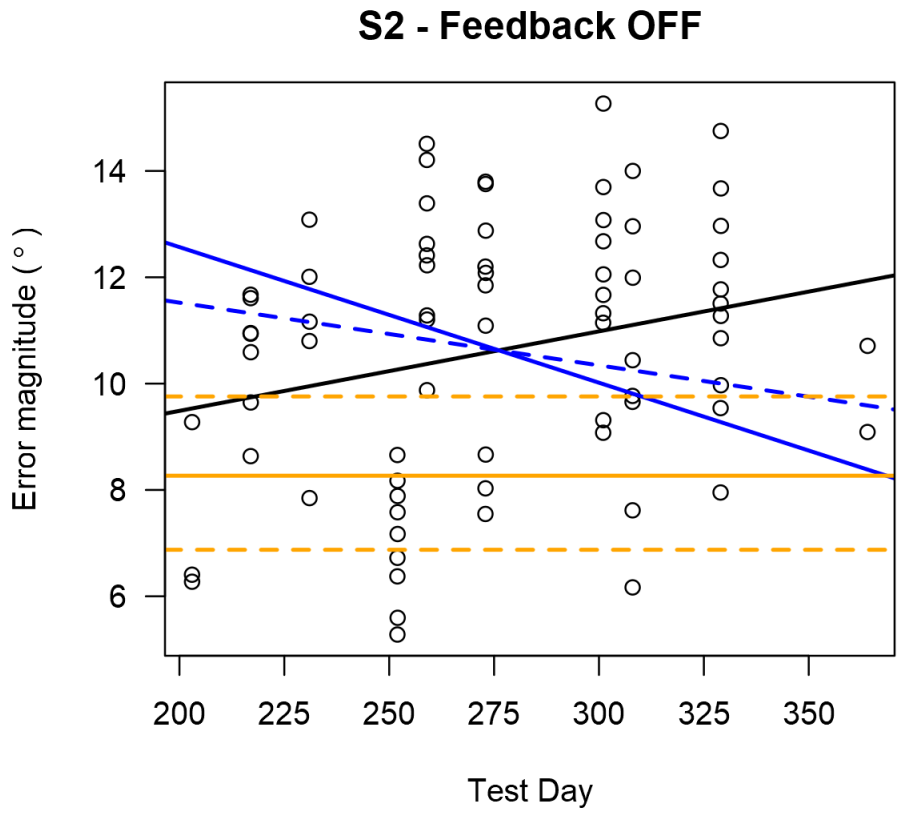

Figure 17 – Localization error changes, misaligned camera, auditory feedback OFF, S3

Points each represent the error centroid distance from the origin for one trial run. The OLS regression line slope is greater than the criterion slope, and just barely greater than the criterion upper 95% confidence limit. S3 thus showed significant lack of adaptation after auditory feedback was removed. Error centroid distances throughout the feedback-OFF period were outside the 95% CI of those at the end of the feedback-ON period.

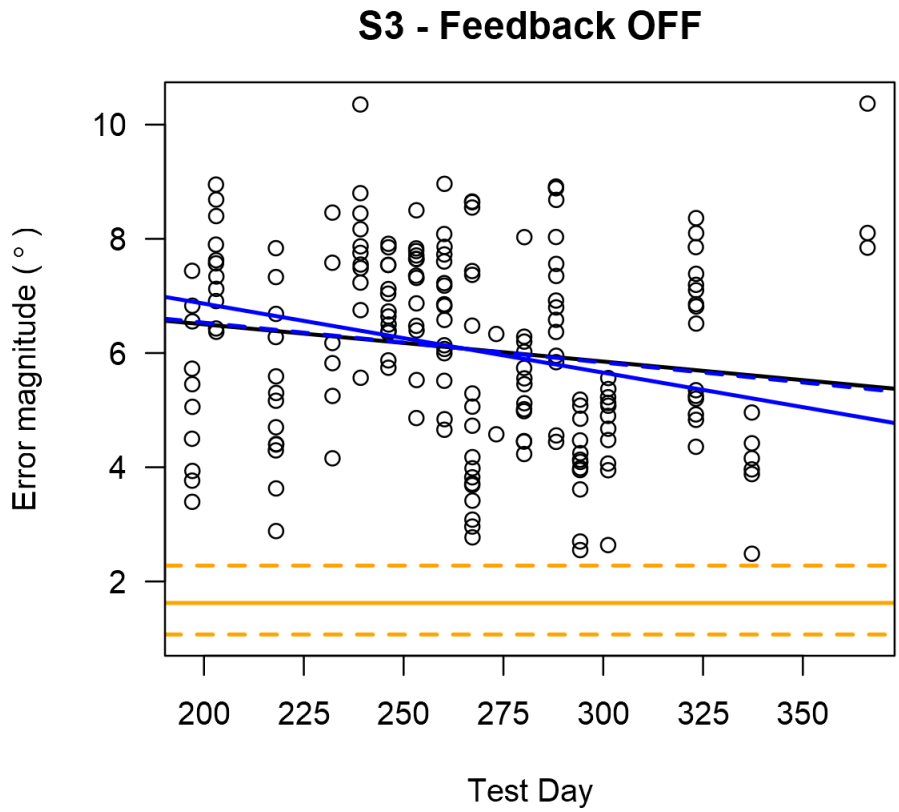

5.2.4 Intraday-interday OCAP directional relationships

An intuitive picture of how the vectors in Section 5.1.4 relate might have large intraday trend and reset vectors, similar in magnitude with opposite directions, and a small interday baseline vector pointing roughly in the same direction as the intraday trend. This was not typically the case, however, as all three vectors had average magnitudes within 4.4°–4.6°. Trends in vector lengths, some consistently being larger than others, only emerged when small windows of direction differences were isolated. The vectors generally had triangular, rather than collinear, configurations. As suggested by Figure 6, Figure 7 and Figure 8, only after considerable averaging does a clear pattern emerge among the intraday trends.

Data from all 3 subjects and both condition periods (feedback-ON, feedback-OFF) were pooled within the misaligned-CAP experiment. Although OCAPs moved in separate directions in feedback-ON and feedback-OFF conditions, the relationships between intra- and interday vectors within conditions were similar. Data from the two conditions were therefore combined to increase analysis power. Despite differences among subjects and feedback conditions, relationships between intraday trends and interday vectors remained consistent.

The similarity between intraday and interday baseline vectors was assessed by subtracting the prior intraday vector direction from that of the interday vector. Direction differences were plotted with stacks within 1° bins around the circle. The differences were significantly nonuniform, and the two vectors tended to point in the same direction.

Figure 18 – OCAP inter- and intraday vector directions, misaligned cameras, S1, S2, and S3

Distribution of direction differences between interday and intraday vectors. Direction differences had a mean confined to the 0° quadrant and were significantly nonuniform. Note distribution bias toward and stack at 0° difference.

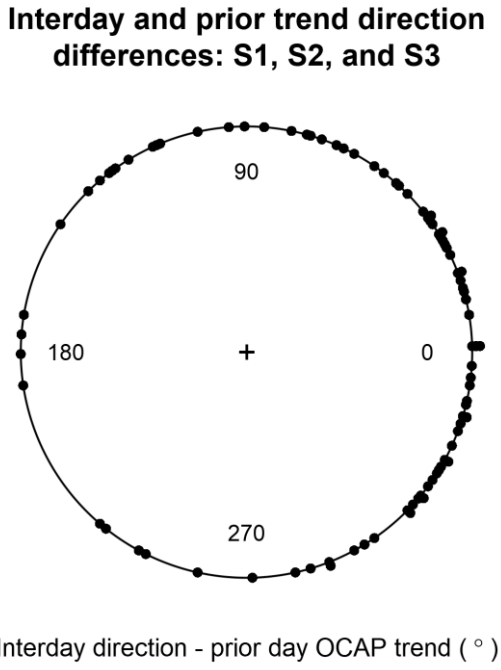

- Difference circular mean angle = 14°
  - Bootstrap 95% CI: −5°–34°
- Difference circular variance = 0.58
  - Probability of observing variance as low or lower assuming no correlation < 10^{-6}
- Interday change, prior-day trend direction circular correlation coefficient = 0.41
  - No positive correlation probability < 7×10^{-5}
- Fit probability:
  - uniform: < 0.01
  - von Mises: > 0.1

For each day, excluding the first day, the direction of the prior test day's intraday trend (prior OCAP trend) was subtracted from the direction of the vector pointing from the end of the prior day's intraday linear model to the beginning of the current day's model (reset). Intraday trends and interday reset vectors tended to point roughly in opposite directions.

Figure 19 – OCAP interday reset, misaligned cameras, S1, S2, and S3

Distribution of direction differences between interday reset and previous intraday trend vectors. Direction differences had a mean confined to the 180° quadrant and were significantly nonuniform. Note the distribution bias toward 180° difference.

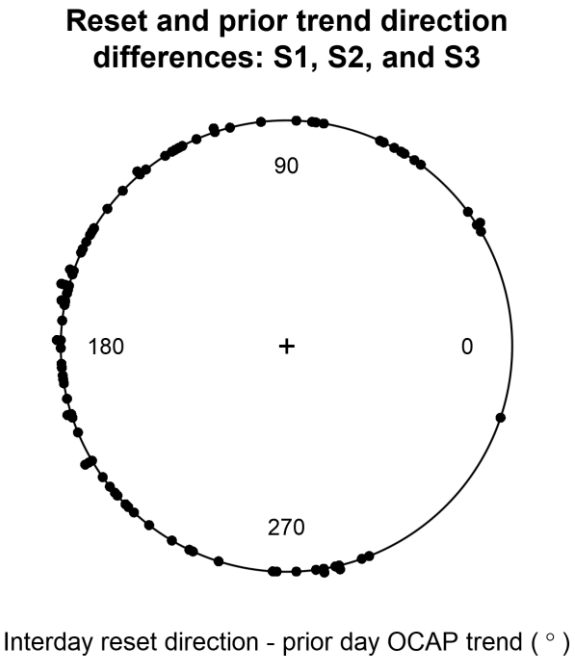

- Difference circular mean angle = 169°
  - Bootstrap 95% CI: 149°–189°
- Difference circular variance = 0.57
  - Probability of observing variance as low or lower assuming no correlation < 4×10^{-5}
- Interday reset, prior-day trend direction circular correlation coefficient = 0.36
  - No positive correlation probability < 5×10^{-4}
- Fit probability:
  - uniform: < 0.01
  - von Mises: > 0.1

5.3 Discussion

When auditory feedback was provided, S1 and S3 seemed to adapt slowly to camera misalignment. Figures in Section 5.2.1 show that OCAPs for S1 and S3 grossly moved toward the misaligned CAP during the feedback-ON period. Early OCAP estimates in these figures do not agree with the OCAPs and misalignments indicated in Table 1 because of limitations imposed by the boundaries of the screen, as discussed in Section 4.2.2. Nevertheless, the direction and magnitude of registered localization errors still change reliably with CAP or OCAP movement. The apparent trend of OCAP movement to the misaligned CAP is thus not invalidated by screen-imposed bias on OCAP estimation.

Corresponding with the figures in Section 5.2.1, data in Section 5.2.3.1 show that error centroid distances for S1 and S3 dropped over the months that auditory feedback was enabled. These very slow rates of centroid distance reduction, −0.05°/day and −0.02°/day, represented average total changes too large to be explained by noise in touch responses. Both rates exceeded those necessary to qualify for minimal adaptation, so the possibility of adaptation to camera misalignment could not be rejected.

Unlike S1 and S3, S2 showed a significant lack of adaptation during the feedback-ON period. Figure 10 shows that S2's OCAPs did not change substantially over the first 6 months, although there was considerable session-to-session variability. Figure 13 supports this observation, as error centroid distances showed little change over this period. Although centroid distances did fall minutely, it is very unlikely that such a modest decline would be observed if at least minimal adaptation were occurring.

When auditory feedback was removed, all 3 subjects exhibited increases in localization error. Feedback-OFF OCAPs plotted in Section 5.2.1 showed a variety of trajectories relative to the CAP: S1's OCAP steadily moved away the CAP, S2's OCAP acutely jumped upward away from the CAP, and S3's OCAP took a tangential course around the CAP. As shown in Section 0, all subjects had significant lack of adaptation apparent in error centroid distances, with those of S1 and S2 increasing over time.

S3's error centroid distances did decrease during the feedback-OFF period, but only after a marked increase in centroid distances relative to the end of the feedback-ON period. Although S3's centroid distance reduction was only barely outside the 95% confidence limit for minimal adaptation, no error centroid distance for any individual trial run was within the 95% CI of the feedback-ON endpoint. If auditory feedback was not the dominant, or only, factor driving and

maintaining adaptation for S3, one might expect greater retention of error correction. Although less pronouncedly, feedback-OFF centroid distances and linear models for S1 and S2 were also predominantly above the feedback-ON endpoint CI.

One could argue that feedback-OFF models should start as anchored to the feedback-ON endpoint, and then be fitted as best as possible from there. This anchoring, however, would assume that the start of the feedback-OFF period is indistinguishable from the end feedback-ON period. This would be true if removing feedback had no immediate effect. As best seen with S1 and S3, though, error centroid distances markedly increased as soon as auditory feedback was removed. Figure 10 also suggests a distinct difference between the 2 periods for S2, even though OCAPs did not seem to change consistently within periods. Because target localization with auditory feedback differed so clearly from that without auditory feedback, linking linear models for the 2 periods by anchoring would not be appropriate.

Overall, subjects did not seem to passively adapt to camera misalignment through the daily use of their devices. Camera misalignments were maintained out of lab, and at least S2 and S3 consistently used the device at home. The only indications of adaptation, however, were linked to auditory feedback provided with in-lab testing. As soon as auditory feedback was removed, localization error increased for all subjects, and all subjects significantly lacked signs of adaptation to misalignment. If any other factors contributed to adaptation, they were not substantial enough to observe in the feedback-OFF period, and left the subjects functionally unable to adapt.

Subjects claimed to not be actively adjusting pointing strategies based on feedback, but may have subconsciously registered corrections that were consistently repeated over time. Along with slowly pushing OCAPs toward the CAP for S1 and S3, auditory feedback may have also

maintained learned corrections or prevented error increases by continuously reminding subjects of nature of their errors. This reminder effect would be consistent with the jump in errors seen across subjects when feedback was removed. Without the program reminding subjects of their errors, the subjects quite rapidly changed pointing behavior to what may have been more natural to them. Only S1 seemed to show a gradual reversal of pointing corrections, but even then, only after a marked increase from the feedback-ON endpoint. S2 may have naturally been developing a greater upward pointing bias over time, and feedback may have suppressed this and delayed its manifestation until feedback was removed. The apparent reduction in error for S3 in the feedback-OFF period may only be a reflection of the subject returning to normal OCAP shifts, with any reduction in error simply being coincidental.

OCAPs certainly changed over time in both the feedback-ON and feedback-OFF periods. Plots in Section 0 showed that the interday changes were related to faster intraday changes observed within each session. Between sessions, changes within the previous session were grossly reversed, so intraday shifts were not simply building on top of one another to create the observed interday shifts. There was considerable variability in OCAPs across days, and similar variability was seen in how intraday OCAP models related to each other. Nevertheless, the consistency of relationships between intraday trends and interday changes increases confidence that observed OCAP changes were not random, but were measurably directed, albeit noisy, shifts in OCAPs over time. When auditory feedback was provided, shifts in S1 and S3 seemed guided toward the constant CAP. Without feedback, subject OCAPs appeared to move in their own directions, although how those directions were determined or what caused OCAPs to shift is unknown.

6 Optimal camera alignment position stability

Users of visual prostheses with the camera separated from the eye require camera alignment to ensure pointing and reaching accuracy. Current programming methods implicitly assume that OCAPs do not change over time, and camera alignment is only required once. If this is not true, however, a user's accuracy may be expected to decline as the difference between the set CAP and OCAP grows. In such cases, regular camera realignment may be required to maintain accuracy.

Based on observed OCAP trends seen in Chapter 5, at least for S1 and S3, OCAPs did not appear to remain constant over time. S2 only showed significant OCAP changes between feedback-ON and feedback-OFF conditions, but no significant changes within condition periods. This chapter provides details on OCAP measurements taken for the same 3 Argus II users over 3.6–5.3-year timespans. Aside from the adaptation experiment of Chapter 5, CAPs were set as close to OCAPs as possible during each testing session. Feedback-OFF, open-loop testing was continued for nearly all tests after the adaptation experiment concluded. Analyses tested the hypothesis that differences seen in OCAP estimations were not simply the product of random measurement noise, but reflected significant changes in OCAPs with time. OCAP trends within days and within trial runs were examined to confirm that consistent OCAP changes could be observed on shorter timescales. The results suggest that user OCAPs are not constant, even within days or trial runs, and average OCAPs slowly drift to new locations in the camera's FOV over months and years.

6.1 Methods

6.1.1 Subjects and equipment

The same subjects as described in Section 5.1.1 participated in this study. Data collection spanned 2012–2017, excluding periods of misalignment detailed in Section 5.1.2. During this time, Argus II equipment was upgraded for all subjects from version A2E8 to A2E14. The most significant change between these versions, for the purposes of this study, was that the A2E14 camera had a smaller FOV. While the A2E8 camera captured a 66° × 49° FOV, the A2E14 camera only captured 49° × 38°. Subjects with OCAPs outside the A2E14 camera FOV could only be given CAPs approaching the boundary of the camera FOV, with no way to resolve the remaining misalignment. As discussed in Section 4.2.2, such misalignments impeded accurate OCAP estimation beyond the range of the camera, and exacerbated errors and misalignments with OCAPs already beyond the A2E8 camera FOV.

S1 was upgraded on August 3, 2015, S2 was upgraded on July 8, 2015, and S3 was upgraded on July 1, 2015. While A2E8 equipment did not consider array size for camera alignment options, A2E14 equipment catered more to dual-metal and similar stimulating arrays than the older slotted arrays. Most CAP step sizes along each dimension were still 0.27°. Possible CAPs for the slotted arrays in S1 and S2 ranged ±14.70° along the horizontal dimension, and ±11.90° along the vertical dimension. S3's dual-metal array covered a smaller area, and therefore could be given CAPs ranging ±15.52° horizontally and ±13.83° vertically. As with A2E8, CAPs were adjusted along intrinsic camera axes, and cameras were physically rotated to match array orientations.

6.1.2 Study procedures

Starting in July 2012 (S1 and S2) and September 2012 (S3), subjects attended test sessions approximately every 1–2 weeks. Subjects performed target localization as described in Section 4.2.2. Each session involved 20–1000 trials, with median trial counts of 85 (S1), 200 (S2), and 360 (S3). Testing with S2 and S3 continued regularly through January 2017, with some follow-up testing in July, August, and November 2017. S1, because of difficulties unrelated to this study or the visual prosthesis, was only able to attend a few sessions after the misalignment adaptation study: 7 sessions between November 2014 and February 2016. S1 thus participated in a total of 14 sessions outside of the misalignment study, whereas S2 participated in 63 sessions and S3 in 98 sessions.

After the misalignment adaptation study was completed, CAPs were altered to either further investigate misalignment or track OCAP changes. Localization tests with new sustained camera misalignments were performed until April (S3) and July (S2) 2014 to confirm the lack adaptation effects. S1 retained the original camera misalignment until November 2014, which was the next visit for this subject when localization testing could be performed. Afterward, until the transition between A2E8 and A2E14 in mid-2015, mixtures of test runs with proper CAPs and runs with misaligned CAPs were used to evaluate the effects of misalignment on OCAP estimation.

Once A2E14 equipment was issued, target localization tests were dedicated to evaluating direct-stimulation OCAP estimates. OCAP estimates were first obtained by using direct stimulation at the beginning of the test session. Afterward, target localization was used to determine subject accuracy with the recommended CAP. Further target localization was used to test more accurately aligned CAPs determined by inspection. Starting in April 2016, target-localization

alignment was tested independently of direct stimulation by starting sessions with default CAPs of (0°, 0°) and iteratively testing resultant OCAP estimates.

Parameters for target localization tests were modified to optimize OCAP estimation as testing progressed. Testing in 2012, before the misalignment study commenced, always included auditory feedback. Based on the effects of auditory feedback seen in Section 5.2.3, auditory feedback was disabled for nearly all tests after the misalignment study. Auditory feedback was only reenabled when target localization was used for another protocol that required feedback.

When new camera misalignments were tested without auditory feedback, it was suggested that subjects could still be receiving corrective feedback from the limits of the screen area. To compensate for this, margins of the form described in Section 4.2.2 were implemented in April 2014. These margins reduced feedback from the screen's borders, and also improved OCAP estimation by reducing limitations on subject responses. Margins were initially implemented with custom software, which allowed margins of any width to be specified. In November 2015, this custom software was replaced by software from SSMP. SSMP's modified software greatly simplified OCAP estimation, but limited margin widths to 3× the radius of the target.

Subjects were allowed to use both hands interchangeably, even within runs, for all target localization tests prior to October 14, 2015. Afterward, subjects were occasionally instructed to use only one specific hand for each trial run. Because substantial OCAP differences were observed depending on which hand was used, subjects were always told to use only one hand per run for all tests after December 2, 2015.

6.1.3 Analysis of OCAP-time relationships

Although target localization tests were used in a variety of experiments with different conditions, each trial run still provided some estimate of the subject's current OCAP. So long as cameras were not drastically misaligned, and the hand used was taken into consideration, all such OCAP estimates could be used track OCAP changes over time. Accordingly, primary analyses only considered trial runs with proper camera alignments. Proper camera alignments were considered those for which the distance between OCAPs and CAPs were estimated to be less than 3°, or the alignment was constrained by the camera FOV. Misalignments imposed by the camera FOV were undesirable, but under such conditions, no other system configuration would have provided better estimates of the subject's OCAP. When investigating only directional changes among OCAPs, but not absolute OCAPs, misaligned-camera data were included in the analysis. The filtered data were then used to determine OCAP dynamics across years, within sessions, and within trial runs.

For convenience, OCAP estimates were first analyzed along horizontal and vertical dimensions separately. Intraday OCAP model start positions for each dimension were plotted over time. When different hands were specified for trial runs, points from each hand were plotted together with different colors. These plots permitted seeing variations over time more clearly than combined horizontal-vertical plots, and also provided insight on how using a different hand changed OCAP estimation.

Based on patterns in horizontal and vertical OCAP plots, data were fitted to linear or nonlinear models. When substantial differences were apparent between dominant- and nondominant-hand responses, models were generated only considering both-hands and dominant-hand

responses, excluding purely nondominant-hand responses. Otherwise, all data points were used for model generation. Linear models were determined using OLS. When data appeared to have a pronounced change in relationship to time, data were fitted to sine curves using nonlinear least squares (NLS). Sine curves were preferred over parabolas as models simply because the infinite extent of a parabola along the vertical axis would be applicable to OCAPs. Linear models would also not extend to infinity, but more data would be required to know where maxima or minima occur. Because OCAP behavior was not consistent among subjects, and more time would be required to make conclusions regarding the OCAP-time relationship, statistical tests were not applied to these models. The models drawn in the plots were simply provided as visual aids.

Following separate horizontal and vertical plots with time, 2-dimensional OCAP estimates were averaged within time bins and plotted with connecting arrows, as in Section 5.2.1. Plots were constructed as described in Section 5.1.3. Because data were collected over more time than in the misalignment adaptation study, larger time bins could be used to further reduce noise in plots. Thus in addition to single-session time bins, data were plotted in time bins with up to 9 and up to 18 testing sessions.

Beyond visualization, time bins were also used to quantify the how likely an OCAP measured in one month would be same as that of another month. If OCAPs are assumed to be constant, with some large random errors involved in OCAP estimation, averaging over larger time periods than runs or days should simply provide a more accurate estimate of the underlying constant OCAP. For constant OCAPS, 95% confidence regions should be approximately the same. Randomly choosing a confidence region and another 30-day time bin's OCAP should therefore have that OCAP fall within the confidence region about 95% of the time.

30-day time bins with OCAP averages and confidence regions were determined for all 3 subjects. Empty time-bins were excluded, and no extrapolations were considered. For each time bin, the number of OCAP averages from other time bins within its confidence region were counted. The sum of counts across all time bins within a subject was compared to the number of possible ordered combinations of time bins: $n(n-1)$ for n time bins. The number of possible pairings, $\binom{n}{2}$, was not used because confidence regions were not equally-sized, and one point being in another's confidence region did not necessary imply they were both within each other's confidence regions.

Relationships between intraday OCAP trends and interday OCAP changes were analyzed as described in Section 5.1.4. Because only OCAP changes were considered in the directional analyses, and not absolute OCAPs, all target localization runs with at least 10 trials, even those from the misalignment adaptation study, were pooled across subjects for analysis. S1's data, however, were not included in these directional analyses because of the scarcity and irregularity of the S1's visits. Thus, only data from S2 and S3 were pooled together.

Just as OCAPs had trends within and across days, intra-run trends also appeared. The directions of intra-run trends were therefore also compared to changes between runs and overall intraday trends. The same 3 horizontal/vertical plane vectors were analyzed as for interday-intraday OCAP changes, but applied to runs instead of days: the intra-run trend, the baseline change between runs, and the reset between the end of one run and start of the next. A fourth vector, the run's corresponding intraday vector, was also used to compare intra-run trends to each corresponding day's overall trend. As with interday-intraday comparisons, misalignment adaptation data were included in this analysis of intra-run trends. Unlike the comparisons

between days, data were pooled from S1, S2, and S3. S1 could be included in this analysis because visit irregularity should not affect relationships between runs in a single session.

6.2 Results

6.2.1 Separated horizontal and vertical OCAPs

OCAPs changed substantially over time, albeit very slowly. In the below plots, points each represent the initial OCAP estimate for one day. Black points were derived from trial runs for which subjects could use either hand interchangeably. Blue points were derived from dominant-hand responses, and red points from nondominant-hand responses. Not every dominant-hand data point had a corresponding nondominant-hand point because of time restrictions on test sessions. OLS linear or NLS sine models are shown in black as suggested visual aids. Gaps spanning 2013–2014 correspond to the misalignment adaptation study; data from the misalignment study were excluded from these OCAP analyses. S1 attended sessions scarcely and irregularly, so larger gaps are present in that subject's data.

Figure 20 – S1 horizontal OCAPs over time

Each point represents the initial ordinate value of one session's intraday model relating horizontal OCAP to time. Positive horizontal OCAP values correspond to rightward deviations from center. The OLS regression line was determined using all data points and has a slope of 0.006°/day.

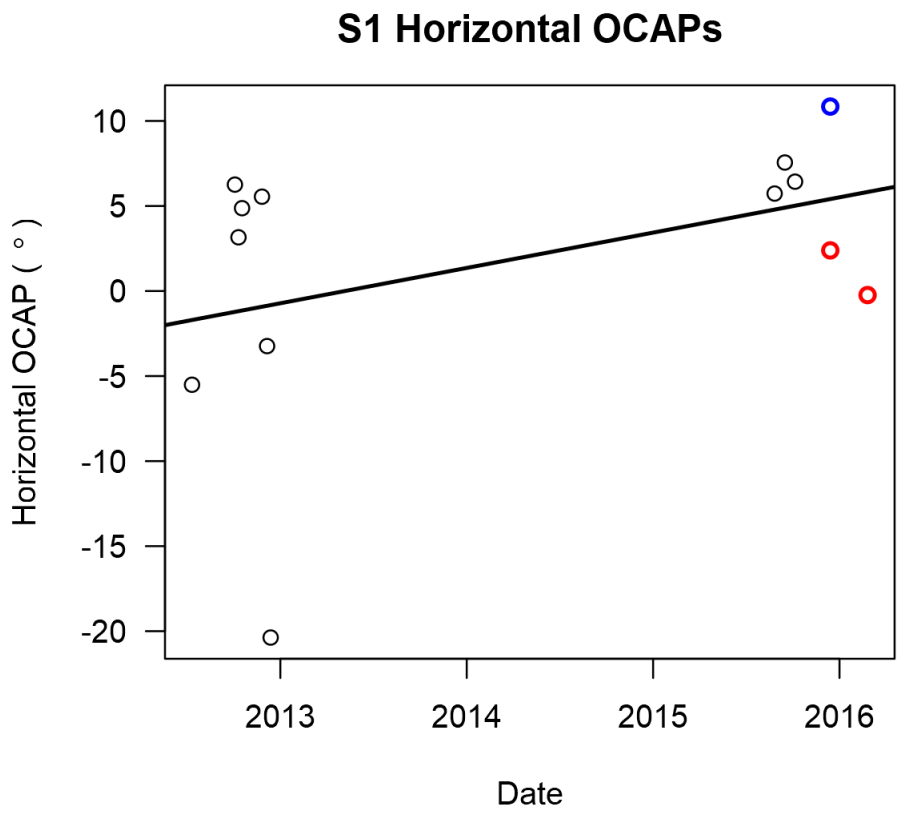

Figure 21 – S1 vertical OCAPs over time

Each point represents the initial ordinate value of one session's intraday model relating vertical OCAP to time. No clear relationship was apparent between vertical OCAPs and time, so no OCAP-time model was generated. Note that most vertical OCAPs shown are below the lower CAP limit imposed by the camera (−15.1°, considering rotation).

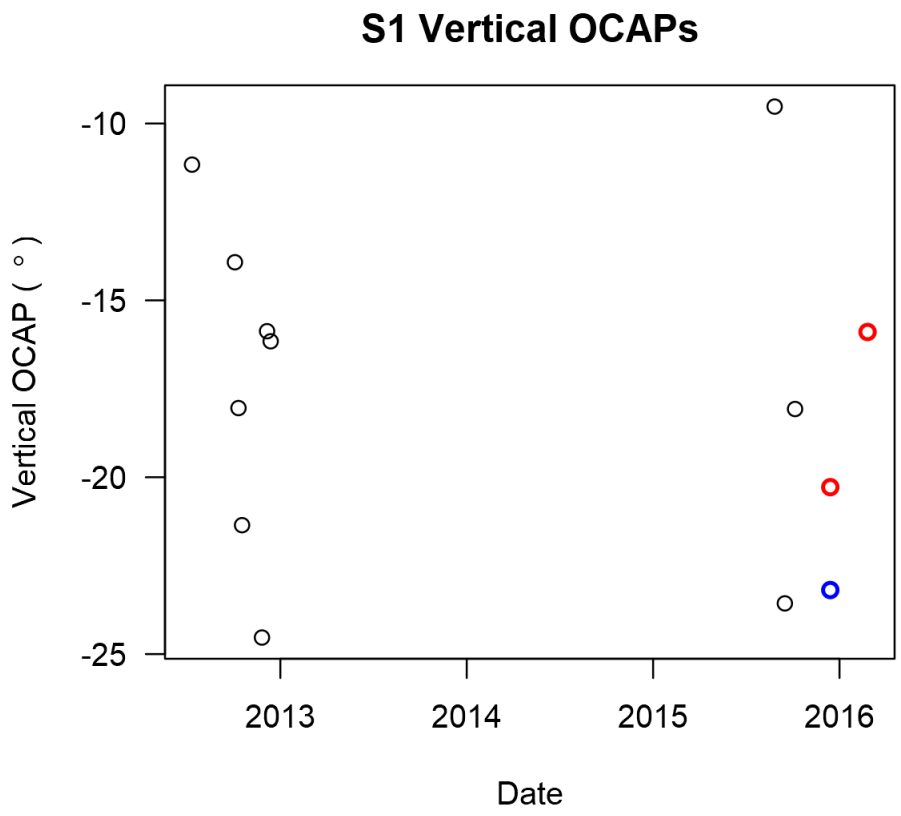

Figure 22 – S2 horizonal OCAPs over time

Each point represents the initial ordinate value of one session's intraday model relating horizontal OCAP to time. Positive horizontal OCAP values correspond to rightward deviations from center. OCAPs appeared to move rightward through 2012–2014, and then leftward from the beginning of 2015 onward. The nondominant hand (left hand) generally pointed more rightward than the dominant (right) hand within each session. The fitted sine model has an amplitude of 5°, 7-year period, and a zero-level of −15°. Only both-hands and dominant-hand responses were used for fitting the model.

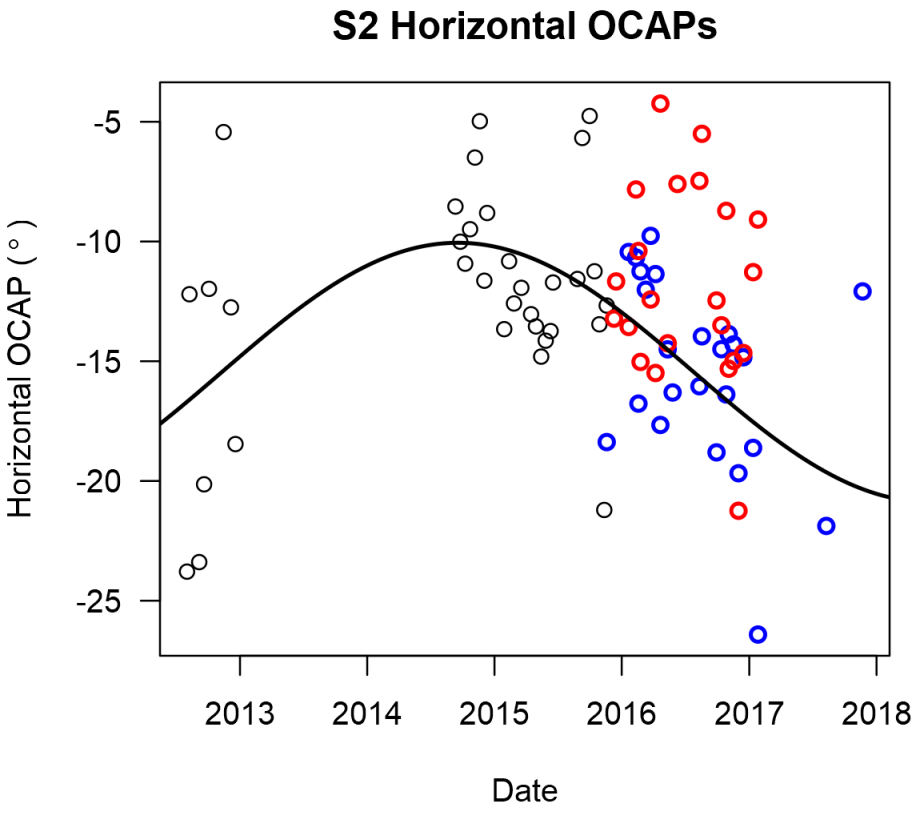

Figure 23 – S2 vertical OCAPs over time

Each point represents the initial ordinate value of one session's intraday model relating vertical OCAP to time. OCAPs appeared to move upward through 2012–2014, and then downward from the beginning of 2015 onward. The nondominant hand (left hand) generally pointed lower than the dominant (right) hand within each session. The fitted sine model has an amplitude of 8°, 6-year period, and a zero-level of 19°. Only both-hands and dominant-hand responses were used for fitting the model. Note that nearly all OCAPs were above the upward CAP limit of 11.90° for the A2E14 camera.

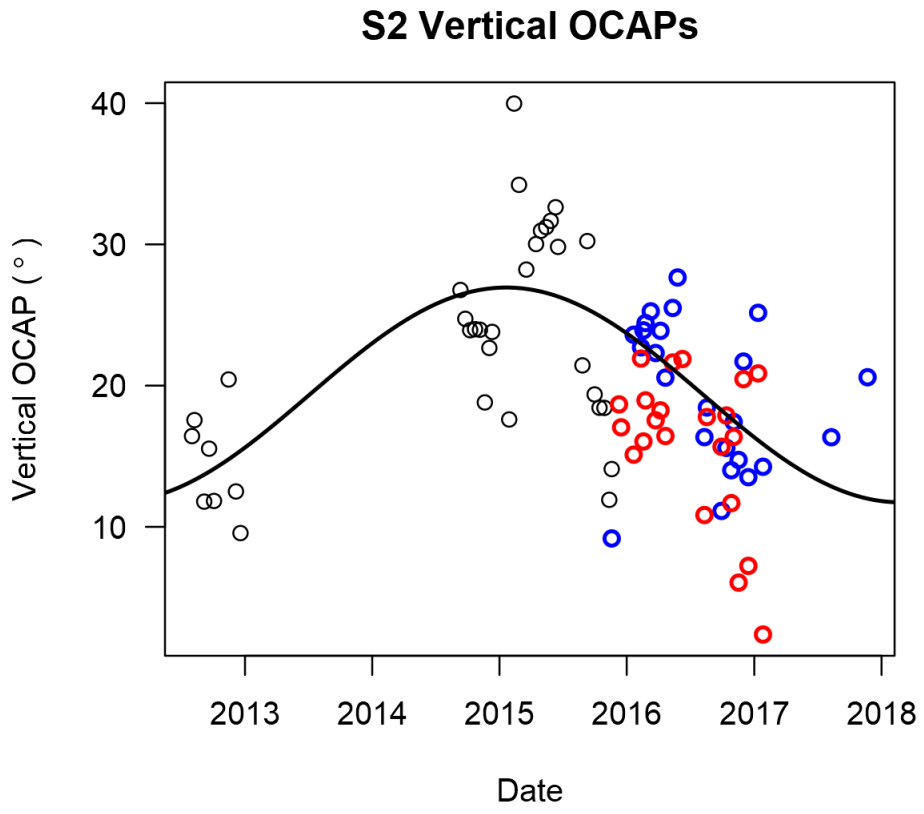

Figure 24 – S3 horizonal OCAPs over time

Each point represents the initial ordinate value of one session's intraday model relating horizontal OCAP to time. Positive horizontal OCAP values correspond to rightward deviations from center. OCAPs moved rightward over time, particularly through 2012–2016. The nondominant hand (right hand) consistently pointed more leftward than the dominant (left) hand. The OLS regression line was determined using only both-hands and dominant-hand responses and has a slope of 0.01°/day.

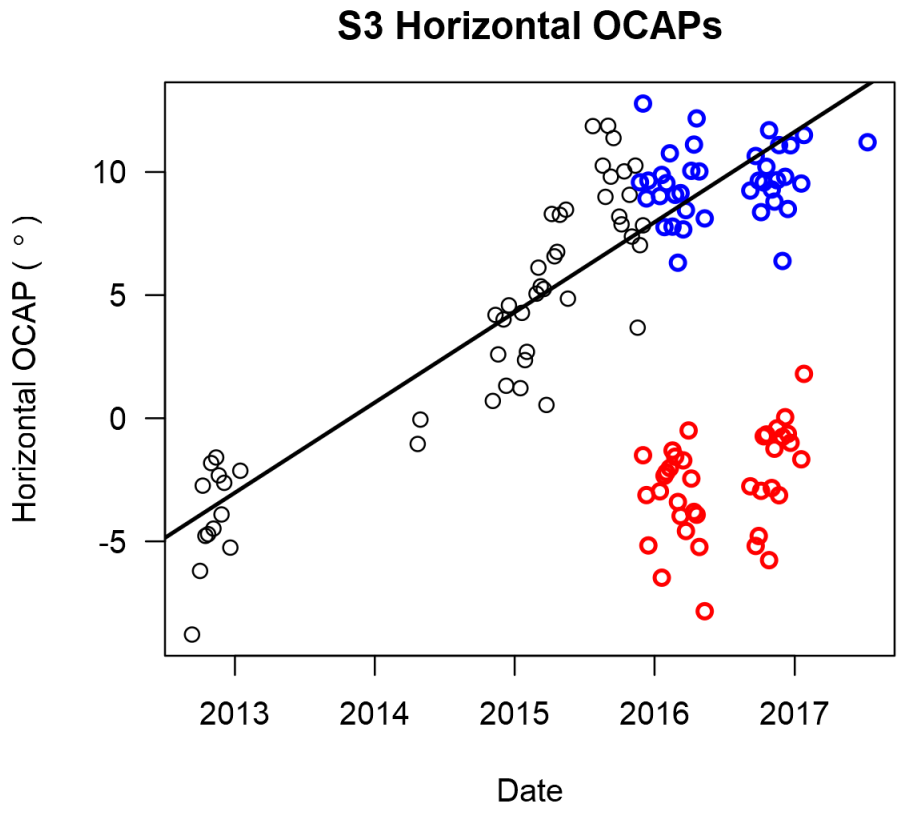

Figure 25 – S3 vertical OCAPs over time

Each point represents the initial ordinate value of one session's intraday model relating vertical OCAP to time. OCAPs appeared to gradually move downward throughout the period of data collection. The nondominant hand (right hand) generally pointed lower than the dominant (left) hand. The OLS regression line was determined using only both-hands and dominant-hand responses and has a slope of −0.003°/day.

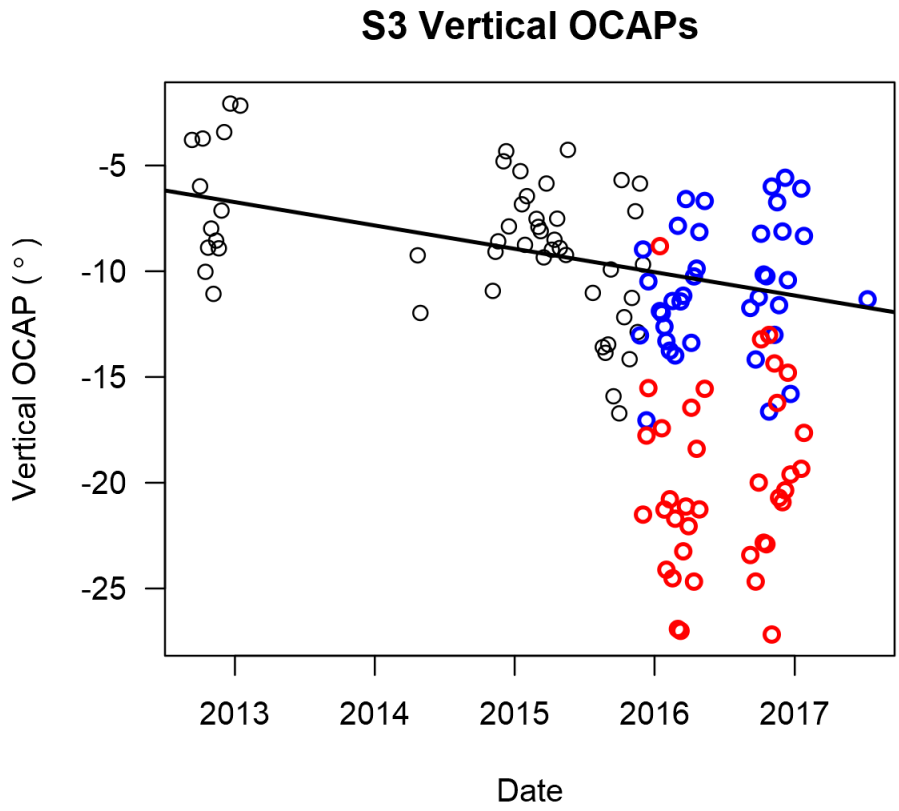

6.2.2 OCAP changes in the horizontal-vertical plane

Plots below show OCAPs relative to the camera FOV, such that the origin is the optical center of the camera. Similar to the treatment of misaligned-camera data in Section 5.2.1, aligned-camera data were first plotted with small time bins equal in size to the typical time between a subject's sessions. The small time bins demonstrated the gross variability of the data, and averaging with time bins 7–9- and 25- times larger permitted visualization of long-term trends. OCAP positions for timeframes with no data were interpolated and are shown as small black triangles. Dotted arrows either connect two periods that have no data or reflect artifacts from changes in equipment, and in both cases lack meaningful information. All data points reflect intraday models of both-hands or dominant-hand responses, but no purely nondominant-hand responses.

6.2.2.1 Small time bins

Just as with Figure 6–Figure 8, variability was too high to readily see any OCAP trends. Each point represents the average of OCAP model initial points from 0–2 testing sessions.

Figure 26 – OCAP changes, aligned camera, S1, 14-day time bins

Points represent average initial OCAPs within 14-day time bins. Each time bin contained 0–2 testing sessions. Variability was high with these small time bins. The long line of missing data points (triangles) corresponds to the particularly long period without proper camera alignment between 2012 and mid-2015.

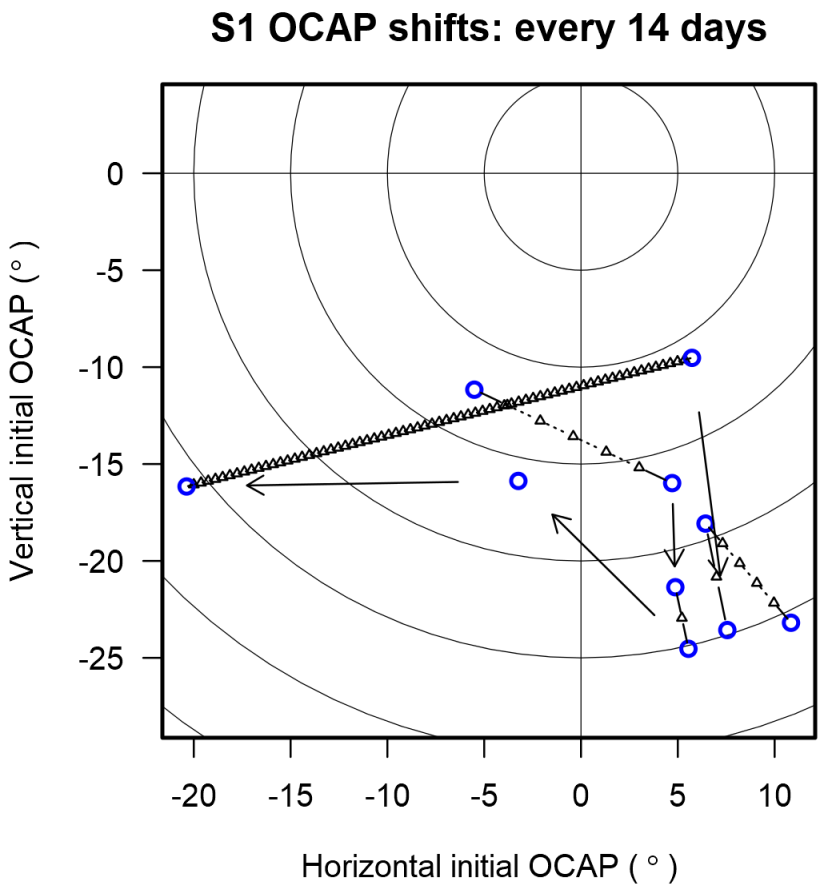

Figure 27 – OCAP changes, aligned camera, S2, 14-day time bins

Points represent average initial OCAPs within 14-day time bins. Each time bin contained 0–3 testing sessions. Variability was very high with these small time bins.

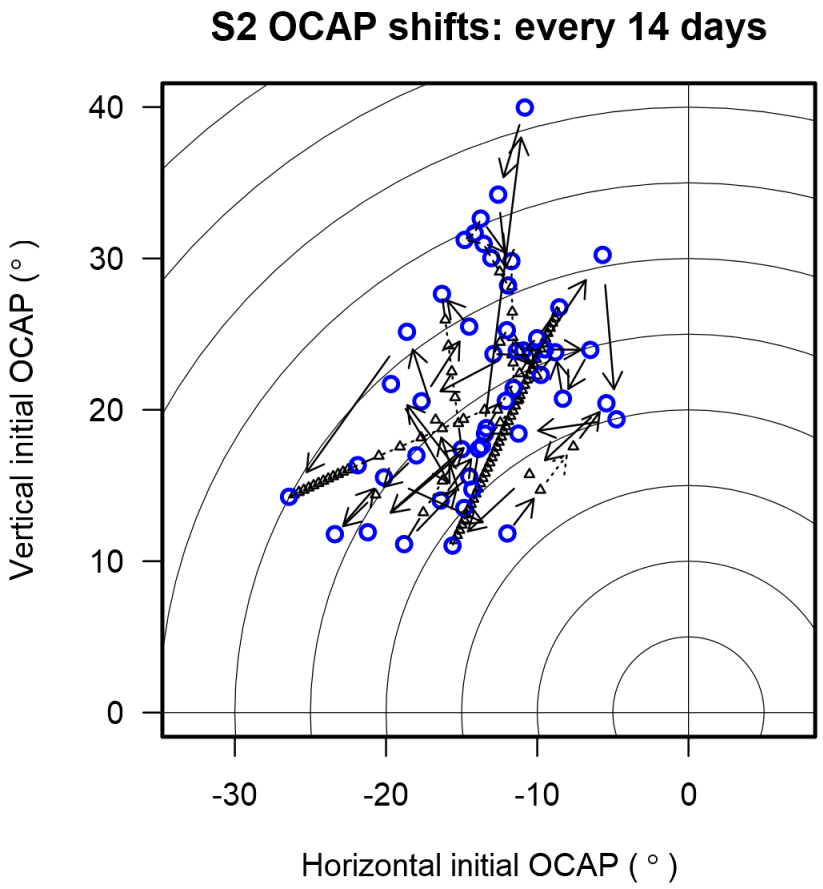

Figure 28 – OCAP changes, aligned camera, S3, 7-day time bins

Points represent average initial OCAPs within 7-day time bins. Each time bin contained 0–2 testing sessions. Variability was very high with these small time bins.

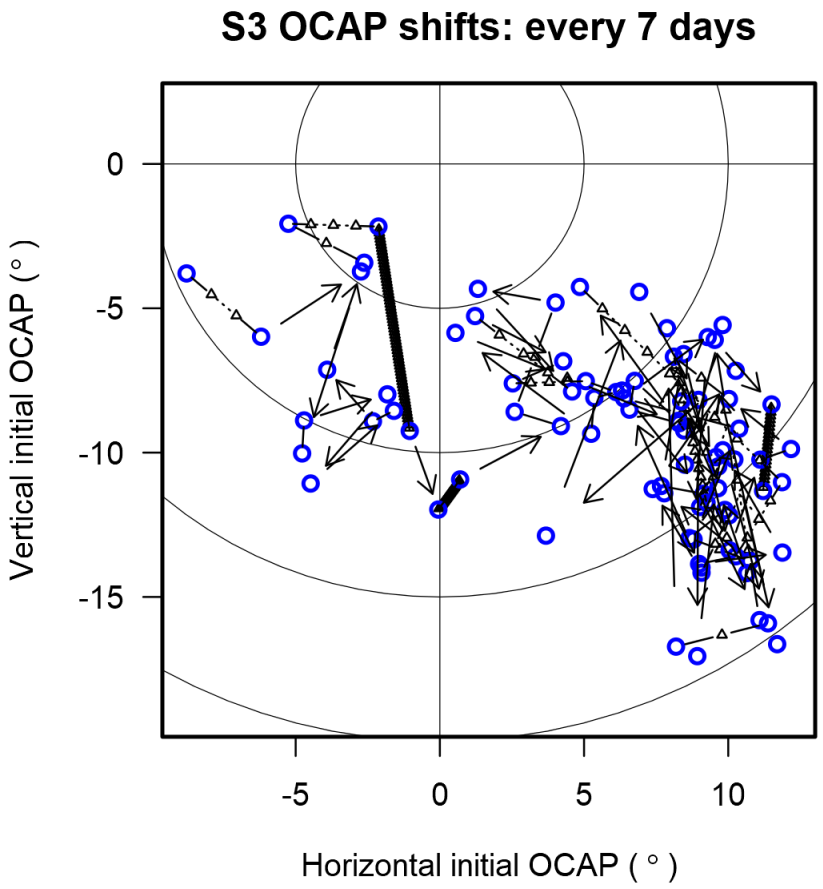

6.2.2.2 *Intermediate time bins*

Averaging within larger time bins reduced noise. Noise appeared reduced sufficiently to see some patterns after increasing time-bin size by a factor of 9. Because S1 did not participate in as many sessions or for as long a time, that subject's time bin size was only increased by a factor of 7. Dashed lines indicate the 95% confidence region about each time-bin average.

Figure 29 – OCAP changes, aligned camera, S1, 98-day time bins:

Points represent average initial OCAPs within 98-day time bins. Each time bin contained 0–4 testing sessions. S1's OCAPs grossly moved down and right during the observed period.

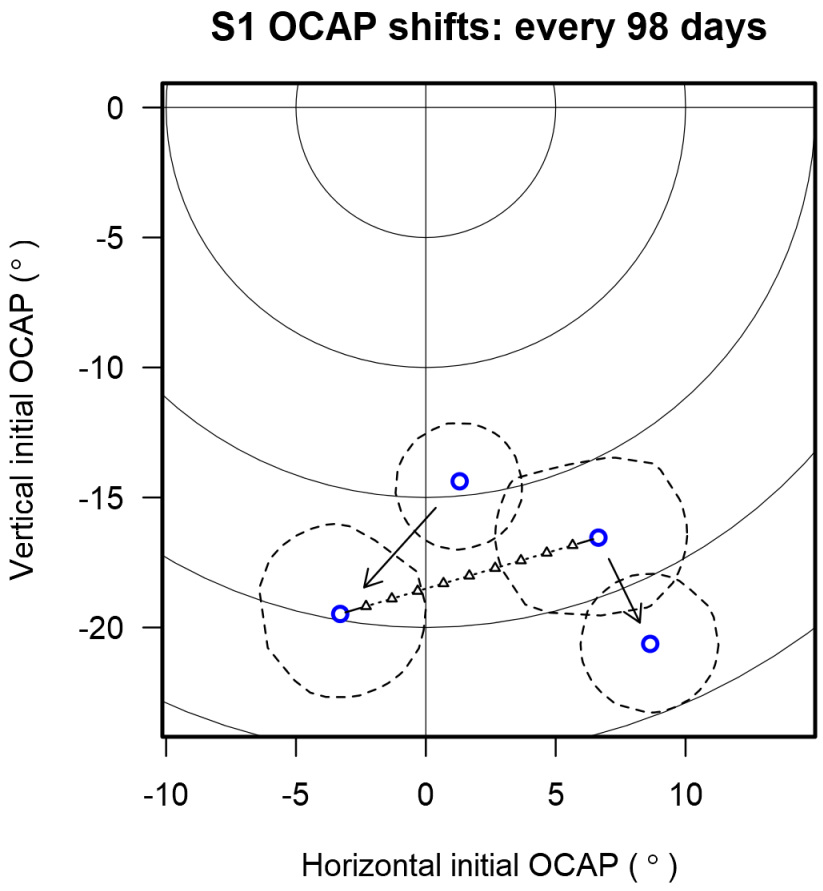

Figure 30 – OCAP changes, aligned camera, S2, 126-day time bins

Points represent average initial OCAPs within 126-day time bins. Each time bin contained 0–9 testing sessions. S2's OCAPs moved through a 4.5-year circuit, consistent with Figure 22 and Figure 23, before moving farther to the left. The dotted arrow projecting down from 35° eccentricity, 115° azimuth reflects the change to A2E14 equipment and new OCAP estimation limits imposed by the camera. The antepenultimate (30° eccentricity, 150° azimuth) and last (24° eccentricity, 120° azimuth) data points each only represent 20 data points, and therefore have large confidence regions.

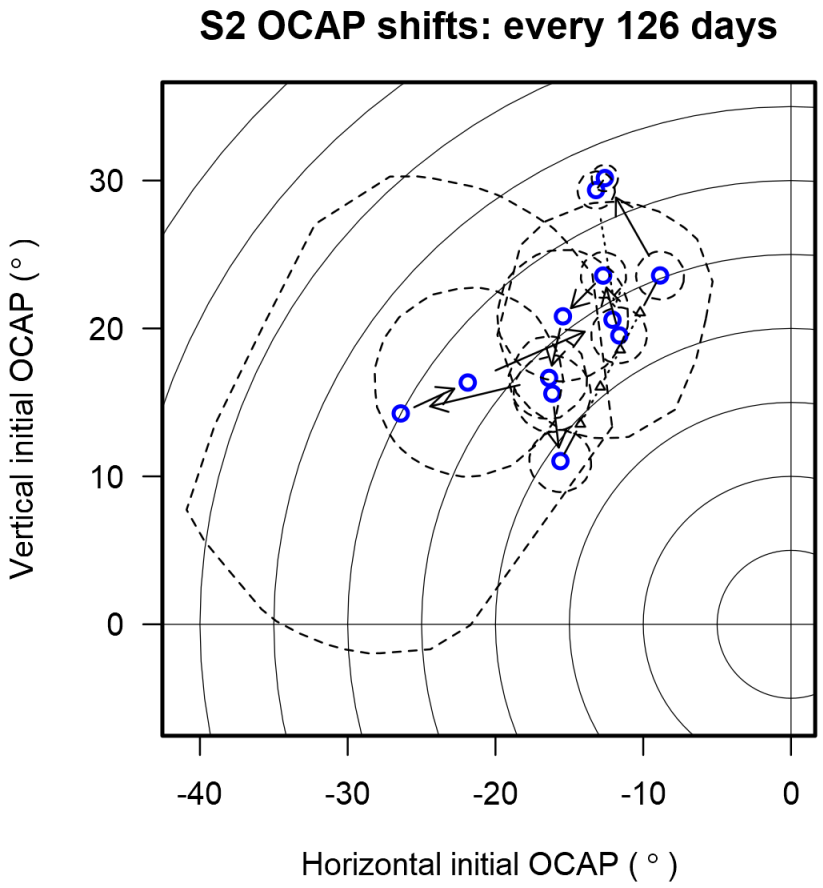

Figure 31 – OCAP changes, aligned camera, S3, 63-day time bins

Points represent average initial OCAPs within 63-day time bins. Each time bin contained 0–8 testing sessions. S3's OCAPs moved steadily rightward through May 2015 (days 0–981, points 1–16). The dotted arrow at about −45° azimuth corresponds the time when equipment was upgraded in July 2015. Although OCAPs were not beyond the limits of the new camera before the equipment change, the new OCAPs were slightly outside the range of the camera after the upgrade. This difference could be related to some combination of the existing rightward trend and the new glasses. For example, the camera's position relative to the eye may have differed slightly between A2E8 and A2E14 glasses.

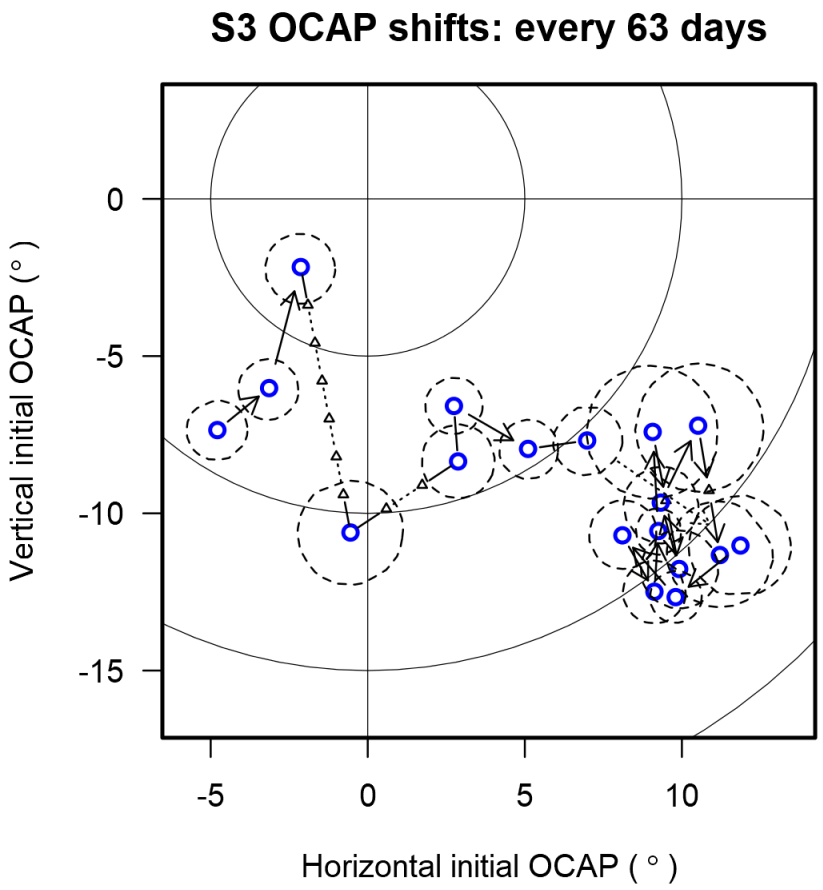

6.2.2.3 Large time bins

OCAPs were more aggressively averaged within time bins 25-times larger than their originals for further noise reduction. This reduced the plot data for S1 to only 2 meaningful points.

Figure 32 – OCAP changes, aligned camera, S1, 350-day time bins

Points represent average initial OCAPs within 350-day time bins. Time bins contained 7, 0, 0, and 4 testing sessions. S1 OCAPs in the final time bin were significantly more rightward than those in the first time bin.

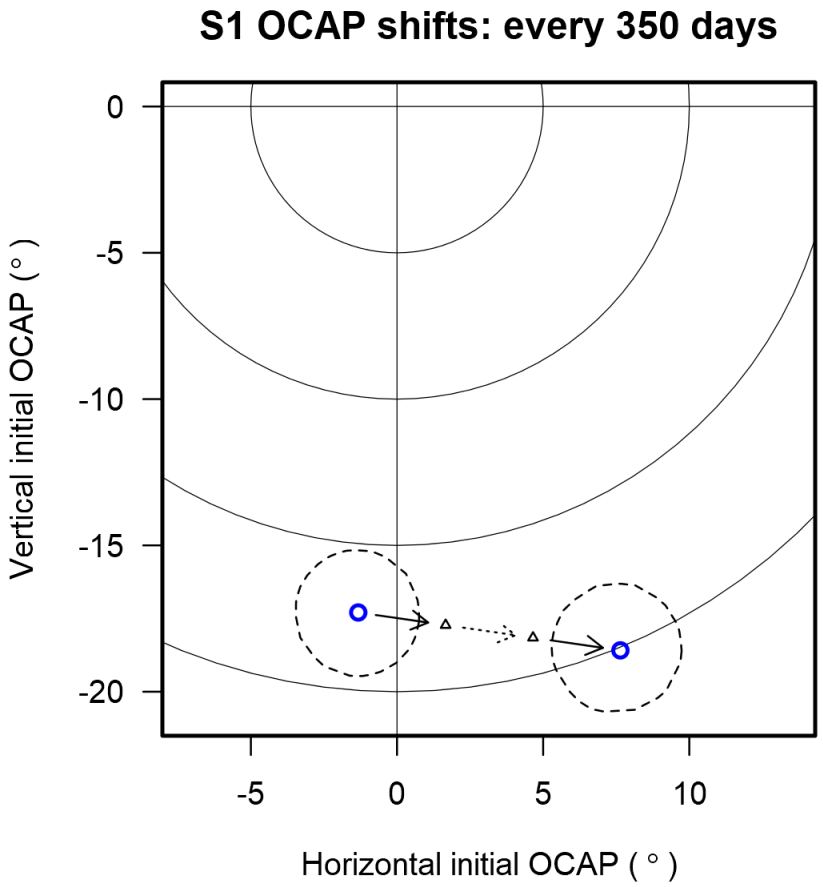

Figure 33 – OCAP changes, aligned camera, S2, 350-day time bins

Points represent average initial OCAPs within 350-day time bins. Each time bin contained 0–18 testing sessions. S2's OCAPs moved from approximately 22° eccentricity, 138° azimuth to 30° eccentricity, 112° azimuth during the first 3 years (points 1–3), and returned to where it started over the subsequent 2 years (points 3–5). The time when equipment was upgraded to A2E14 is marked by the dotted arrowing projecting down from 30° eccentricity, 112° azimuth.

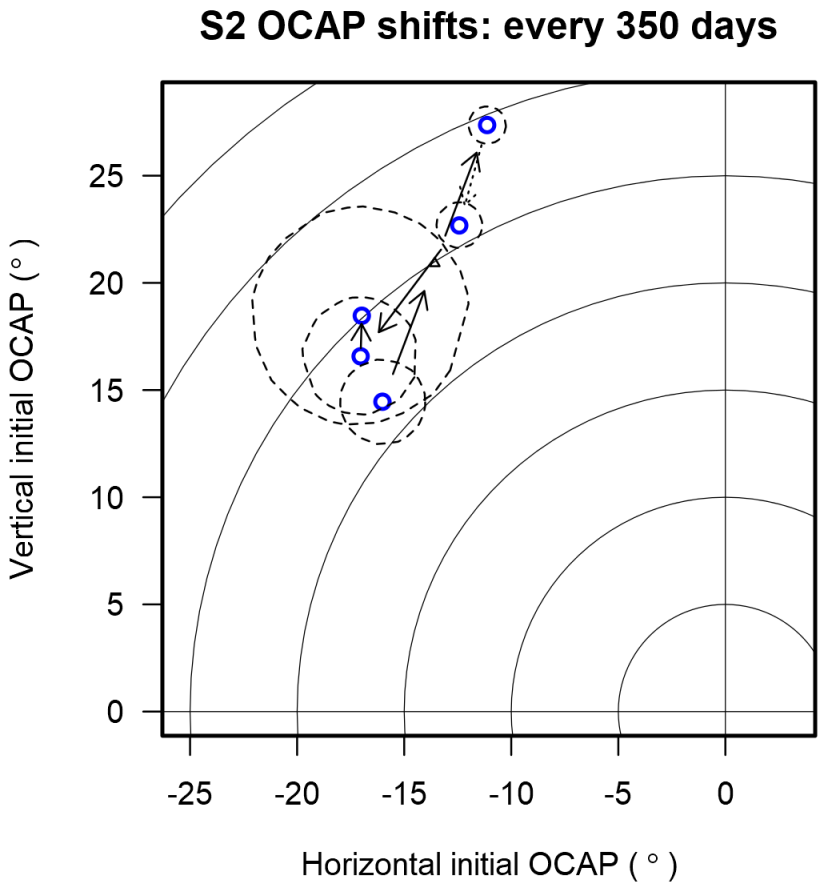

Figure 34 – OCAP changes, aligned camera, S3, 175-day time bins

Points represent average initial OCAPs within 175-day time bins. Each time bin contained 0–17 testing sessions. S3's OCAPs consistently moved rightward, albeit more slowly (points 7–11) after the first 3.5 years. The dotted line projecting from 10° eccentricity, −51° azimuth corresponds to when equipment was upgraded to A2E14.

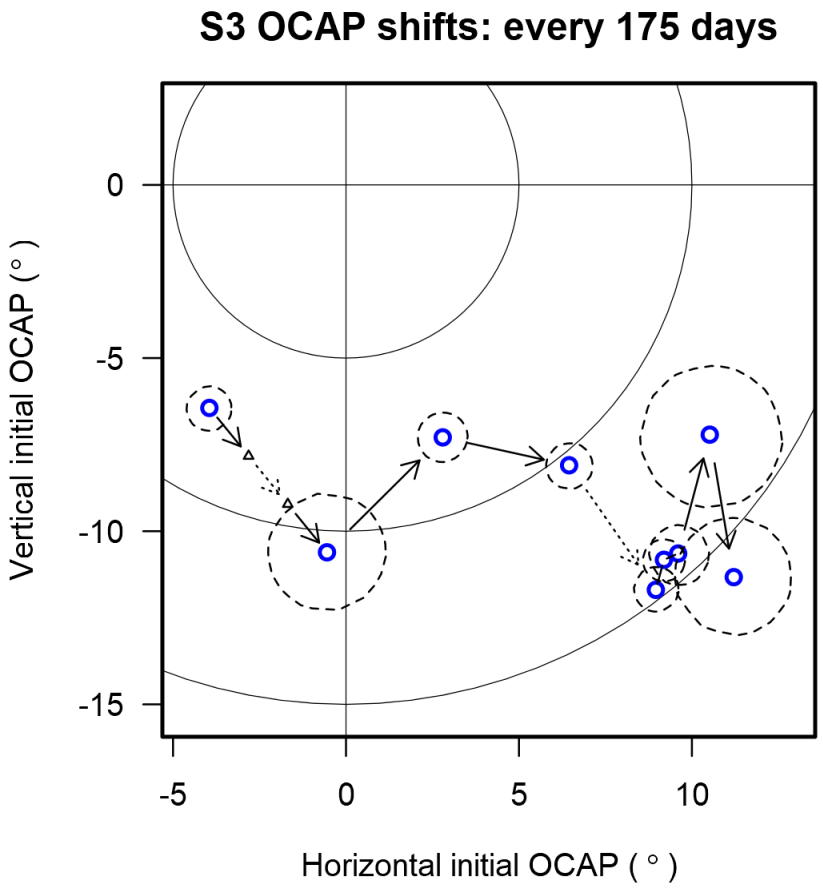

6.2.3 Constant-CAP applicability

Based on 30-day time-bin averages and confidence regions, the probability of one randomly chosen time-bin average being within another randomly chosen time-bin's confidence region is shown in Table 6 for each subject. Variability was too high for meaningful visualizations, similarly as in Figure 26 through Figure 28. Most confidence regions, however, were not trivially large. Only 15% or fewer random combinations of confidence region and extrinsic OCAP average for a subject involved an OCAP average that was within the confidence region.

Table 5 – 30-day time bins and confidence regions

All empty time bins were excluded from analysis. S1 had the fewest nonempty time bins. S2 and S3 had approximately the same number of time bins, but S3 tended to attend sessions more frequently than S2. S3's larger sample sizes within time bins, and tendency to be more precise, yielded much smaller OCAP average confidence regions.

Subject	Time bin count	Sessions per time bin			Confidence region width (°)		
		Median	Minimum	Maximum	Mean	Minimum	Maximum
S1	8	1	1	2	8.9	4.9	12
S2	32	2	1	3	8.6	2.4	21
S3	31	3	1	5	3.1	1.3	6.5

Table 6 – Extrinsic 30-day OCAP averages within confidence regions

Only a small portion of OCAP average and confidence region combinations involved an average that was encapsulated within the confidence region. OCAP averages to which a confidence region belonged were excluded from counts. All other OCAP averages were considered extrinsic to the confidence region. Total encapsulated counts are the sums of encapsulated extrinsic OCAP average counts across all confidences regions for a subject. Any particular OCAP average could be counted multiple times if it fell within multiple confidence regions.

Subject	Extrinsic 30-day OCAP averages within each confidence region			Total encapsulated	Total possible combinations	Percent encapsulated
	Median	Minimum	Maximum			
S1	0.5	0	2	6	56	11%
S2	4	0	18	158	992	16%
S3	1	0	14	77	930	8.3%

6.2.4 Intraday-interday OCAP directional relationships

Combining aligned- with misaligned-camera data for S2 and S3 maintained biases seen in Figure 18 and Figure 19. Intraday trend directions tended to have small differences from interday baseline directions, while differences from interday reset vectors averaged around 180°. Direction differences were stacked within 1° bins around the circle.

Figure 35 – OCAP inter- and intraday vector directions, S2 and S3

Distribution of direction differences between interday and intraday vectors. Direction differences had a mean confined to the 0° quadrant and were significantly nonuniform. Note the prominent stack at 0° difference.

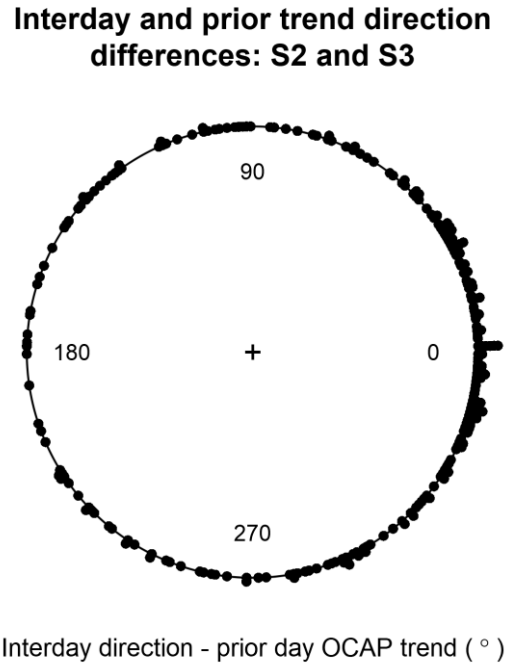

- Difference circular mean angle = 0.04°
  - Bootstrap 95% CI: −13°–14°
- Difference circular variance = 0.64
  - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Interday change, prior-day trend direction circular correlation coefficient = 0.20
  - No positive correlation probability < 0.002
- Fit probability:
  - uniform: < 0.01
  - von Mises: < 0.05

Figure 36 – OCAP interday reset, S2 and S3

Distribution of direction differences between interday reset and previous intraday trend vectors. Direction differences had a mean confined to the 180° quadrant and were significantly nonuniform. Note the prominent stack near 180° difference.

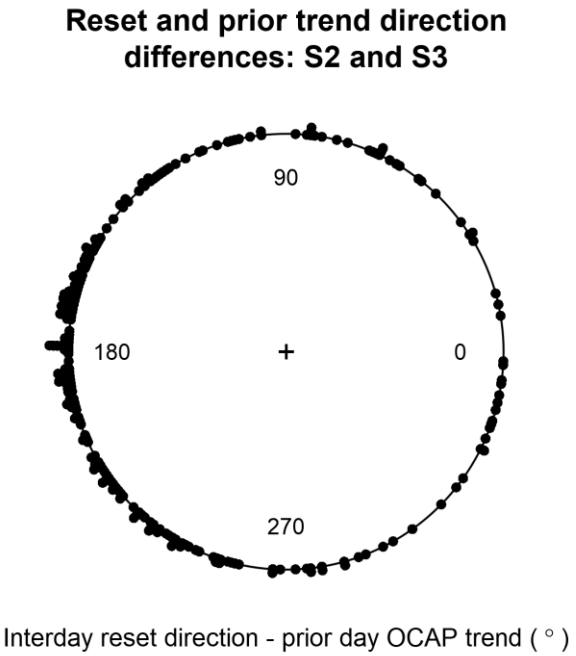

- Difference circular mean angle = 183°
  - Bootstrap 95% CI: 173°–193°
- Difference circular variance = 0.53
  - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Interday reset, prior-day trend direction circular correlation coefficient = 0.42
  - No positive correlation probability < 10^{-5}
- Fit probability:
  - uniform: < 0.01
  - von Mises: < 0.01

6.2.5 Intra-run, inter-run, intraday OCAP directional relationships

Intra- and inter-run OCAP behavior appeared to mimic that of intra- and interday OCAP trends. As with interday plots, direction differences were stacked within 1° bins around the circle. The vector between the start of one trial run to the next (inter-run vector) typically pointed in a direction similar to the first's intra-run trend, and intra-run trends tended to be reversed between trial runs. Unlike intraday/interday OCAP behavior, however, intra-run and inter-run vector directions were inversely correlated, despite direction differences remaining small. Inter-run vector directions correlated with those of reset vectors instead of intra-run vectors. Further, intra-run trend directions only had a modest tendency to align with intraday trend directions,

although the two still directly correlated. These differences reflect greater variability among intra-run vectors than intraday vectors.

Figure 37 – OCAP inter- and intra-run vector directions, S1, S2, and S3

Distribution of direction differences between inter-run and intra-run vectors. Direction differences had a mean confined to the 0° quadrant and were significantly nonuniform. Inter-run and intra-run vector directions were inversely correlated.

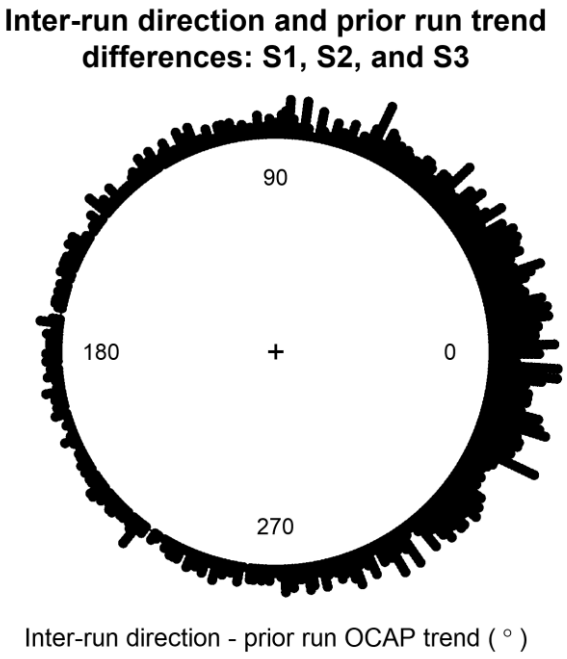

- Difference circular mean angle = 3.6°
  - Bootstrap 95% CI: −0.6°–7.9°
- Difference circular variance = 0.61
  - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Inter-run change, prior-run trend direction circular correlation coefficient = −0.24
  - No positive correlation probability > 0.99999
- Fit probability:
  - uniform: < 0.01
  - von Mises: < 0.01

Figure 38 – OCAP inter-run reset, S1, S2, and S3

Distribution of direction differences between run reset and intra-run vectors. Direction differences had a mean confined to the 180° quadrant and were significantly nonuniform. 180° was, however, just outside the mean 95% confidence interval.

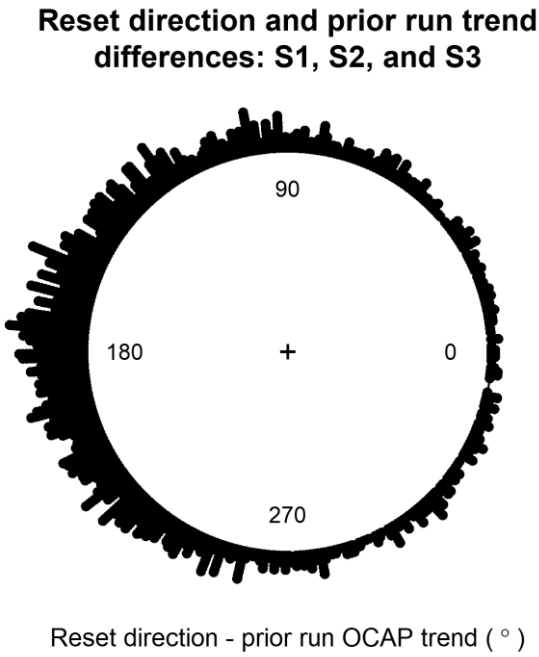

- Difference circular mean angle = 176°
    - Bootstrap 95% CI: 172°–179°
- Difference circular variance = 0.54
    - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Inter-run reset, prior-run trend direction circular correlation coefficient = 0.45
    - No positive correlation probability < 10^{-5}
- Fit probability:
    - uniform: < 0.01
    - von Mises: < 0.01

Figure 39 – OCAP intraday and intra-run trends, S1, S2, and S3

Distribution of direction differences between intraday and intra-run vectors. Direction differences had a mean confined to the 0° quadrant and were significantly nonuniform. The distribution, however, was still generously distributed around the circle.

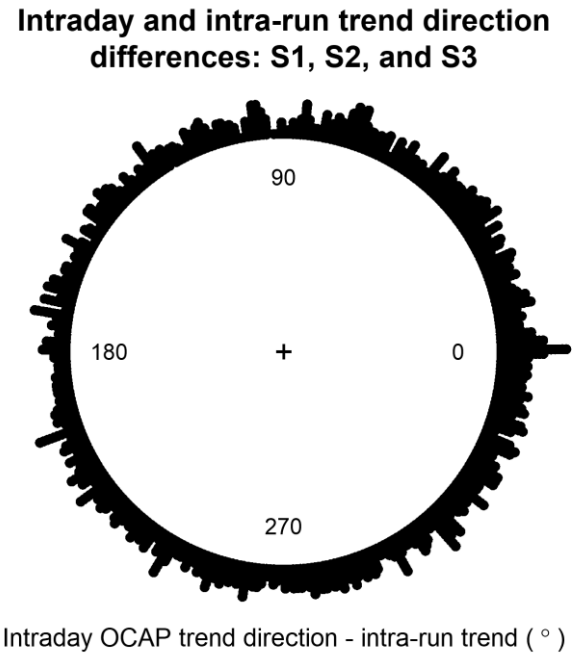

- Difference circular mean angle = −11°
  - Bootstrap 95% CI: −28°–4.9°
- Difference circular variance = 0.89
  - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Intraday, intra-run trend direction circular correlation coefficient = 0.05
  - No positive correlation probability < 0.02
- Fit probability:
  - uniform: < 0.01
  - von Mises: > 0.1

6.3 Discussion

OCAP estimates varied greatly within small time periods, but on average drifted slowly and consistently. The plots with time abscissae in Section 6.2.1, particularly for S2 and S3, show the consistency most clearly. S1's OCAPs varied too much and too few data points were collected to clearly view trends along a single axis. Some difference could be seen along the horizontal axis, in which OCAPs seemed to change on average by about 0.006°/day.

For S2, OCAP estimates within any small period were spread over 10°–15°, but the 10° swathes covered by estimates made gradual progressions over time. The masses of data points appeared to move upward and rightward before 2015, although much of that time was masked by the misalignment adaptation study. Through 2015 and afterward, OCAPs clearly moved downward

and leftward. Considering the parts of the sine models between maxima and minima, S2's OCAPs moved at average rates of 0.008–0.013°/day along a single dimension, or about 0.016°/day total. Although sine models were chosen to capture the possible changes in OCAP shift directions in 2015, there is no empirical reason to assume that periodicity exists. Further observation would be required to justify or revise these models.

Horizontal and vertical OCAP trends were most clear in data from S3. Despite the gap from the misalignment adaptation study, the rightward progression of OCAPs between 2012 and 2016 can be seen quite readily. Horizontal OCAPs seemed steady through 2016 and 2017, but this could be an artifact of the A2E14 camera's FOV: the OCAP was outside the camera's FOV, and this degrades the accuracy of OCAP estimation. Note that nondominant-hand responses, not limited by the range of the camera, continued a rightward trend during that period. Altogether, the horizontal OCAPs appeared to change by about 0.01°/day. Vertical OCAPs also displayed a consistent trend, but at a much slower rate of about 0.003°/day. Combined, the rate of OCAP change was dominated by the horizontal shift and remained close to 0.01°/day.

Figure 24 and Figure 25 for S3 show remarkable consistency in OCAPs before and after the data gap in 2013 and 2014. Trends appeared to continue across the gap undisturbed. This suggests that whatever determines how OCAPs move, at least for S3, was not substantially affected by the misalignment adaptation study. That would be consistent with the misalignment not inherently altering behavior, and observed adaptations being tied to auditory feedback. Subjects may have only been learning to alter pointing behavior for the target localization test, through auditory feedback, but did not generalize those changes to affect pointing overall. After auditory feedback was removed, they resumed their natural pointing behaviors. Ultimately, that involved OCAPs returning to where they would have been had the misalignment study never been

conducted. This reasoning, derived from S3's data, was used to justify the assumption that the differences before and after the misalignment study in S2's OCAPs were not caused by the misalignment study, and S2's OCAP actually did move upward and rightward before later moving downward and leftward. If the data before the misalignment study cannot be considered comparable to those after, one can only examine S2's largely linear trends present from 2014 and later.

Figure 22 through Figure 25 make clear the effect of using different hands for responding. Even in a single session, instructing the subject to respond with the nondominant hand yielded markedly different OCAP estimates than those obtained with the dominant hand. For both S2 and S3, each hand seemed to have a pointing bias directed across the midline, such that the right hand pointed farther left than the left hand. In both of these subjects, the dominant hand pointed higher than the nondominant hand. Biases differed by about 5° in each dimension for S2, and 10°–15° for S3. Dominant-hand responses lined up well with responses collected before hand use was controlled, while nondominant hand responses extended or existed entirely outside previously expected ranges of responses. Because of this consistency, and that dominant-hand use can be assumed as a default, both-hands and dominant-hand conditions were considered comparable and pooled together for analyzing trends.

Although steady progressions were not as clear, the horizontal-vertical plane plots in Section 6.2.2 give a more intuitive visualization of how OCAPs moved. Small time bins with only 1–2 sessions had highly variable OCAP averages, and gradual changes were thus made less intelligible than in previous plots by displaying that variance along both axes. Averaging within larger time bins made patterns more distinguishable. Horizontal-vertical plots for S1 may have displayed OCAP changes more clearly than the separate plots over time, showing distinct

rightward movement, and perhaps a drop downward. Consistent with the sinusoid models applied to S2's OCAP changes, that subject's OCAPs appeared to move through a circuit during the observed period. S3's OCAPs moved consistently rightward, as shown earlier, but the vertical component was less clear in these plots and mainly seemed to contribute noise. Notably, most time-bin averages had confidence intervals that showed clear separation from those of other time bins, indicating significant changes in OCAPs over time.

Comparing OCAP averages and confidence regions for 30-day time bins further demonstrated that OCAPs largely did not stay the same from month to month. Assessing the boundary of each confidence region and all other time-bin averages indicated that only 16% or fewer of the possible combinations yielded an extrinsic OCAP average that was within the selected confidence region. In application, one might wonder how well a CAP set based on a previous month's data would apply in another arbitrarily-picked month. Based on these results, performance differences would be less than statistically significant with a probability of only 8–16%. Confidence region sizes, of course, depend on the amount and quality of data collected, and these results do not speak to differences that would be of practical concern. Magnitude of OCAP differences that require realignment for practical purposes, rather than demonstrating statistical significance, would depend on the application and need for accuracy. Regardless, the range of 8–16% falls remarkably short of the 95% assumed if OCAPs are the same across months, and provides support for OCAPs changing with time.

As OCAPs changed across months, they also changed within days and even within runs. This was initially shown in Section 0. Adding data from this extended observation period to those of the misalignment study confirmed the original conclusions. Each testing session had its own trend in OCAP changes, and this trend pointed in the same direction as the change between sessions, on

average. Intraday trends, however, did not stack upon one another, and there was some amount of reset in OCAP progression between sessions. This reset can be seen in how the direction of the intraday trend and the change following the end of the session pointed in opposite directions, on average. As in the misalignment adaptation study, intraday and interday OCAP change directions were mostly not colinear, and typically had some angular difference. The 3 considered points (prior-session start, prior-session end, next-session start) had a triangular configuration, and vectors pointing among them had similar average magnitudes.

The directions of intra-run trends and inter-run changes showed similar behavior. Trends within runs pointed in the same direction as the change between runs, on average. The direction of change from the end of one run to the beginning of the next pointed on average in the opposite direction from that of the intra-run trend. Although variability was higher among run trends than intraday trends, there was at least some bias toward the intraday trend direction in intra-run trends. The connections among intra-run trends, inter-run changes, intraday trends, and changes between sessions adds confidence that OCAPs do change along some persistent shift, and that the observed changes are not just manifestations of measurement error.

The similar reset behavior between trial runs and between test sessions suggests that target localization testing may transiently accelerate the process of OCAP movement. The accelerated OCAP changes were observed within trial runs and test sessions, and they were reversed after each run and after each session. While intraday trends could be considered unrelated to testing, and perhaps were simply captured by testing, the same cannot be said of trial runs. Trial run trends were of similar magnitude to intraday trends, and resets between runs cannot be readily linked to anything other than the break in testing.

Localization testing is likely not the direct cause of OCAP shifts, however, given the persistence of shifting across unplanned breaks in testing. After the completion of regular testing in January 2017, S2 was not tested again until August 2017. S2's horizontal OCAP continued its leftward trend during that 7-month break to a point beyond most previous measurements. In mid-2014, S3 stopped testing for 6-months, and that subject's horizontal OCAP was markedly farther right than before the break.

Finally, although not enough data were collected with S1 to clearly tell whether a trend persisted during that that subject's break in testing, the rate of the apparent OCAP trend is consistent with those of S2 and S3. S1 had no localization testing between December 2013 and November 2014, and after one 40-trial run, no further testing until August 2015. The available data imply, over the entire period, S1's horizontal OCAP changed at a rate of 0.006°/day, and the vertical OCAP may not have changed. That rate of change agrees reasonably well with the overall rates of change for S2 and S3: both 0.01°/day. If S1's OCAP did in fact change as observed, then target localization testing could not have been the driving force of that change.

7 Eye tracking

Results of Chapter 6 show that subject OCAPs are not constant and slowly drift in the camera's FOV. One candidate source of OCAP changes is eye orientation. As discussed in Section 2.5, eye orientation within the head is taken into account for natural light localization. A visual prosthesis elicits percepts as if the eye were capturing light, and the brain still assigns egocentric locations to those percepts based on head and eye orientations. Blind individuals lack the visual stimulation necessary to fixate straight ahead, and that may carry over to lacking a means to calibrate what eye orientation corresponds to straight ahead. If a subject's average eye orientation drifts with time, that would require OCAPs to change accordingly, and might generate the OCAP shifts shown in Chapter 6.

Experiments described in this chapter measured eye orientations using two different systems for comparison with target localization behavior. The first was a tower-mounted eye-tracking setup that was incompatible with simultaneous target localization, but could be used in alternation with target localization over the course of months. The second tracked eyes during target localization and enabled automatic CAP adjustment, but required eye-tracking glasses that were only available for this testing for 2 days. Analyses of data from the tower-mounted eye tracker tested the hypotheses that the correlation between eye orientations and OCAPs was not 0, and that eye-orientation trend directions within runs or days were not independent of changes between runs or days. The latter hypothesis did not directly link eye orientations and OCAPs, but was necessary to conclude that eye orientations and OCAPs displayed the same types of trend behaviors. To connect trends between measurements, data from both eye-tracking systems were used to test the hypothesis that trend directions of eye orientations and OCAPs or localization errors were also not independent of each other. Prosthesis-integrated eye

tracking analyses also tested the hypothesis that adjusting CAPs with eye orientations would slow the rate of change of localization errors. The results suggest that OCAP shifts are related to changes in eye orientation, and CAP adjustment based on eye tracking can slow down the divergence between the CAP and OCAP.

7.1 Asynchronous OCAP measurement and eye tracking

The Argus II system does not utilize integrated eye tracking. In order to preliminarily investigate the relationship between eye orientation and OCAPs, a second setup that incorporated eye tracking was utilized during testing sessions. Eye tracking was performed at different times and under different conditions from OCAP measurements. Measured eye orientations and OCAPs within each session were then analyzed to check for similarities in behavior.

7.1.1 Methods

7.1.1.1 Subjects and equipment

Only S2 and S3 participated in this study, as introduced in Section 5.1.1. These eye-tracking experiments were performed during the same sessions as target localization tests discussed in Chapter 6, but only in 2016 and 2017. Subjects attended 20 and 31 sessions between February 2016 and January 2017, visiting the lab approximately every 1–2 weeks. S1 could not participate because of scheduling difficulties, and insufficient balance and strength to use the eye-tracking setup.

Eye tracking was performed using the same setup as for touch-precision measurements, described in Section 5.1.3 and shown in Figure 40. Each subject sat with the head pressed against a chinrest and forehead rest. The chin and forehead rests were part of an eye-tracking

tower that also possessed cameras, lenses, mirrors, and infrared illuminators for imaging each eye. To accommodate a display placed in front of the eyes, cameras were aimed downward from the top of the eye tracking tower. For each eye, a tilted hot mirror placed behind a lens reflected the eye's infrared image to a camera. The chinrest was adjusted to bring the subject's implanted eye to the same vertical level as the lens and mirror. To horizontally align the captured image with the eye, the entire lens, mirror, and camera assembly for the eye was synchronously adjusted left or right. Once the eye was in view of the camera, the subject's head was secured to the tower with a Velcro band.

Figure 40 – Tower-mounted eye-tracking setup

A subject is shown on the eye-tracking tower, indicating the apparent location of a percept. Two cameras on the top of tower captured images of the eyes via hot mirrors and lenses positioned in front of the eyes. The head secured by a chinrest, forehead rest, and headband. A transceiver coil was tied next to the implanted right eye using a eye patch. The transparent touchscreen was connected to the eye-tracking computer and collected subject responses. A small display in front of the subject's eyes was not used for blind subject, but was used for general equipment calibration with sighted subjects.

The captured image of the eye was processed by a computer connected to the eye-tracking tower. ViewPoint EyeTracker software from Arrington Research was used to identify and track features of the eye. Custom EyeFingerTrack software collected and saved relevant data from ViewPoint EyeTracker during experiments. As seen in Figure 41, ViewPoint EyeTracker fit a yellow ellipse to the detected pupil, and a red ellipse to one corneal reflection. Three infrared illuminators were aimed at the implanted eye, and the corresponding reflections can be seen in

Figure 41. Although tracking all three reflections, or glints, would have provided more information, the installed software was only capable of tracking one glint at a time. As such, the program was only configured to recognize a glint temporal or superotemporal to the pupil. Note that the image captured by the camera was vertically inverted, as a consequence of equipment geometry, so the superotemporal quadrant appears at the bottom left of Figure 41.

Figure 41 – Pupil and glint tracking

A screenshot is shown of the processed infrared image of the implanted eye. ViewPoint EyeTracker software searched for the pupil within the large red bounding box. The yellow ellipse was fitted to the borders of the detected pupil. The light-blue box's location was anchored to the position of the pupil, and defined the search area for identifying the corneal glint. The 3 glints in the image were generated by 3 separate infrared illuminators. Because of restrictions imposed on the equipment configuration, the captured image was vertically inverted. Accordingly, the top of the subject's eye is at the bottom of the figure.

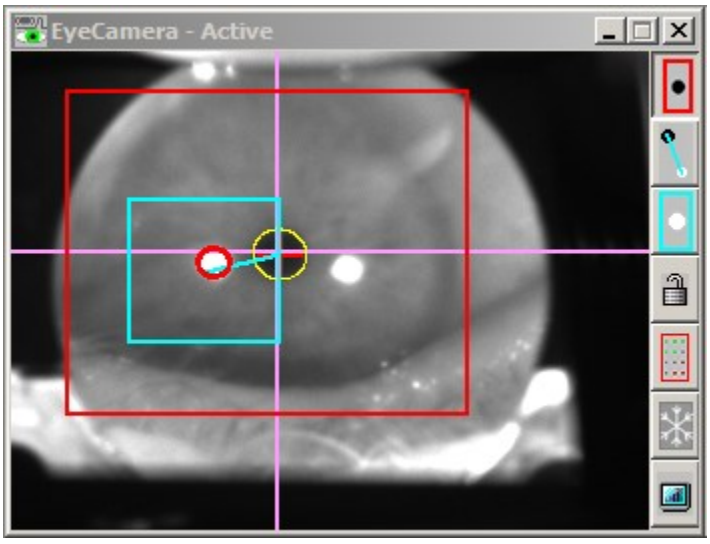

When subjects needed to respond to visual stimuli, a transceiver coil was added to the setup previously described. This coil was the same as that normally attached to the Argus II glasses, but was instead only connected to a long cable. Having the subject's head in contact with the eye-tracking tower precluded use of the prosthesis glasses, and thus prevented the normal method of providing stimulation from the prosthesis. Connecting via the separate transceiver coil bypassed this problem and permitted use of the Argus II stimulating array. Because the coil

was connected directly to the VPU and no video input was provided, stimulation could only be elicited by explicit commands from the attached programming computer. The programming computer allowed individual electrodes to be selected, and caused a common specified current waveform to be sent through each of them.

Subject touch responses were collected by a transparent touchscreen beyond the eye-tracking tower. The distance between the touchscreen and the opening of the lens casing, at the subject's eye, was fixed to 32.6 cm for all trials. The horizontal position of the touchscreen was adjusted so that subject responses landed approximately in the middle of the screen. The positions of the screen and the camera/lens assembly were measured during initial setup and after any changes to accurately measure responses relative to the position of the eye.

7.1.1.2 Study procedures

Eye-tracking experiments were conducted either with or without visual stimulation. In both cases, subjects were asked to provide responses on the touchscreen. Each time the subject tapped the screen, the EyeFingerTrack program extracted eye-tracking data from ViewPoint EyeTracker and recorded those data in conjunction with the touch position. Eye-tracking data were not recorded otherwise. An audio cue played every time the subject touched the screen to confirm data acquisition.

When no stimulation was provided, subjects were instructed to alternately touch close to the nose and touch the screen directly in front of the nose. Subjects could not actually touch their noses because of how equipment was placed on the tower. Subjects touched the screen 20 times with the specified hand for each trial run of this nose localization (NL) task. Each session typically involved 2 NL trial runs with each hand.

When stimulation was provided, the central 4 electrodes were activated simultaneously with just enough current that the subject would reliably perceive a phosphene. S2 required 220 μA of stimulation, whereas S3 only required 120 μA. Each square waveform lasted 250 ms, with 0.9 ms pulses, 0.45 ms per phase with no interphase gap, repeating at 20 Hz. Before each trial run, the subject responded to a sample stimulus with the specified hand to ensure proper placement of the touchscreen. Subjects were then given 20 stimuli to locate by tapping the touchscreen. Each stimulus was accompanied by an auditory cue to alert the subject. Trials were repeated when the subject did not see a phosphene or eye tracking failed for a trial. Trial runs were aborted and restarted if the touchscreen needed to be moved to capture responses. Subjects were asked to slightly adjust head positions if the pupil moved outside the view of the camera. Each session typically involved 2 runs per hand for this phosphene localization (PL) task.

These pointing tasks with eye tracking were done during the same sessions as target localization tests. The order of eye tracking and normal target localization alternated between sessions. The orders of specified hands and stimulation vs no-stimulation tasks were rotated also across sessions to reduce the average influence of timing and order on performance.

Pupil and glint locations measured for eye tracking were record with respect to the window shown in Figure 41. Horizontal and vertical positions could take values in the range [0–1], with (0,0) corresponding to the top left corner of the window. Translation of these window-based coordinates to degrees of eye rotation requires calibration, for which the subject looks at visual targets in specified locations.

Because such calibration was impossible for blind subjects, 4 sighted subjects were asked to look at stimuli presented in 17 locations on the tower-mounted display. A target first appeared at the center of one eye's half of the display. The subject foveated on the target and tapped the

touchscreen according to where the target appeared to be. Another target then appeared in an eccentric location. The subject fixated on the eccentric target and tapped its projected location on the touchscreen. This process, central target followed by eccentric target, was performed for each of 16 points located at ±4° and ±12° along each axis.

7.1.1.3 Analysis of eye orientations

Pupil and glint positions were recorded at the time of subject response for both NL and PL. These measurements, however, are sensitive to head movements that change the position of the eye relative to the camera. While head movements do change where the eye is pointed, translations of the head change the eye's focus far less than head or eye rotations. Pupil and glint positions are equally affected by translations, but because of the curvature of the cornea, glint movement is far less evident during rotation. The pupil position still changes dramatically in response to rotation. Head translations can therefore be separated from eye rotations by considering the position of the pupil relative to the glint. The difference between the pupil and glint positions is referred to as the glint-pupil vector (GPV). While GPV is quite insensitive to head translations in the plane parallel to the lens, it is sensitive to movement of the head closer to or away from the lens. As movement toward or away from the lens was limited by the forehead rest and headband, but no bite bar was present to prevent horizontal and vertical translations, eye orientation analyses focused purely on GPVs.

First attempts at translating raw GPV changes into changes in eye orientation utilized subject touch responses. The position of each touch was translated into degrees relative to the eye using known geometry. Deming regression between GPVs and touch responses, however, failed to yield reliable results. Subject touch precision, quantified in Section 5.2.2, may have been insufficient for resolving the raw GPV vs. degree turn relationship. Instead, sighted-subject GPVs

were compared to visual stimulus locations using OLS regression. Slopes and intercepts of those regression lines were then used to interpret changes in blind subjects' GPVs in terms of degrees. Regression line confidence bands were determined by bootstrap resampling of residuals along the linear model 10^5 times.

Intraday models were computed for GPVs in the same fashion as for OCAPs, described in Section 5.1.3. NL and PL GPVs were considered separately. No substantial difference was found in GPVs related to hand use, so data for right- and left-hand responses were pooled together. Gross GPV-OCAP relationships were plotted using GPV and OCAP model start positions. Correlations were further investigated between GPV model start or end positions and OCAP estimates for each target localization trial on the same day. The significance of the calculated Pearson's correlation coefficient was determined using up to 10^6 bootstrap replicates. Bootstrap replicates modeled 0 correlation by randomly pairing GPV model points and OCAP estimates across days. Individual OCAP estimates were chosen for comparison over OCAP model positions to increase test power. GPV model end positions were only used if start positions displayed substantially more variability, as determined by R^2 for OLS regressions of model positions vs. time across days.

To investigate the potential effect of compensating for eye orientation, average OCAPs and average differences between OCAPs and GPVs were calculated for the duration of the eye-tracking study. Only intraday model start positions were considered, and horizontal and vertical models were analyzed separately. Thus, OCAP model start positions were averaged, as were the differences between OCAP and GPV model start positions. The difference between model start positions was interpreted as a residual OCAP, or what additional CAP adjustment would need to be applied if the camera automatically adjusted its aiming based on GPV. Similarly, the rates of

change of model start positions across days were determined by OLS regression for both OCAPs and OCAP-GPV differences. Residual OCAPs and residual OCAP rates of change were presented as percentages of original OCAPs and OCAP rates of change.

GPV trends within and between sessions, and within and between runs, were investigated using the same vectors and directional analyses described in Sections 5.1.4 and 6.1.3. As mentioned above, GPV intraday models were calculated separately for NL and PL, but no distinction was made based on which hand was used. Intra-run models, of course, necessarily had no mixtures of conditions. Beyond investigating the reset behavior already observed with OCAPs, GPV intra-run and intraday models were compared to OCAP models to determine any relationship in trend directions. Intra-run directions of each measurement were compared to the intraday direction of the other experiment. NL and PL data remained separate for analyses, resulting in two data pools: (1) NL intra-run, OCAP intraday and OCAP intra-run, NL intraday direction differences; (2) PL intra-run, OCAP intraday and OCAP intra-run, PL intraday direction differences.

7.1.2 Results

7.1.2.1 Raw GPV reading calibration

Raw GPV readings from sighted subjects correlated very well with presented stimulus locations, as shown in Figure 42 and Figure 43. Horizontal GPVs and stimulus locations had a Pearson's correlation coefficient of 0.98 ($p < 10^{-6}$), and the correlation along the vertical axis was 0.95 ($p < 10^{-6}$). OLS regression lines suggested that rotation of one degree corresponds to about 0.37% of the window's width and 0.47% of the window's height. These relationships were used for analyzing all GPVs measured from S2 and S3.

Figure 42 – GPV-eye rotation horizontal relationship

Points each indicate the mean horizontal GPV for each horizontal stimulus location for one sighted subject. Each subject is shown with its own color and symbol. Dashed lines indicate the OLS regression line 95% confidence band. The solid black regression line has a slope of 0.0037.

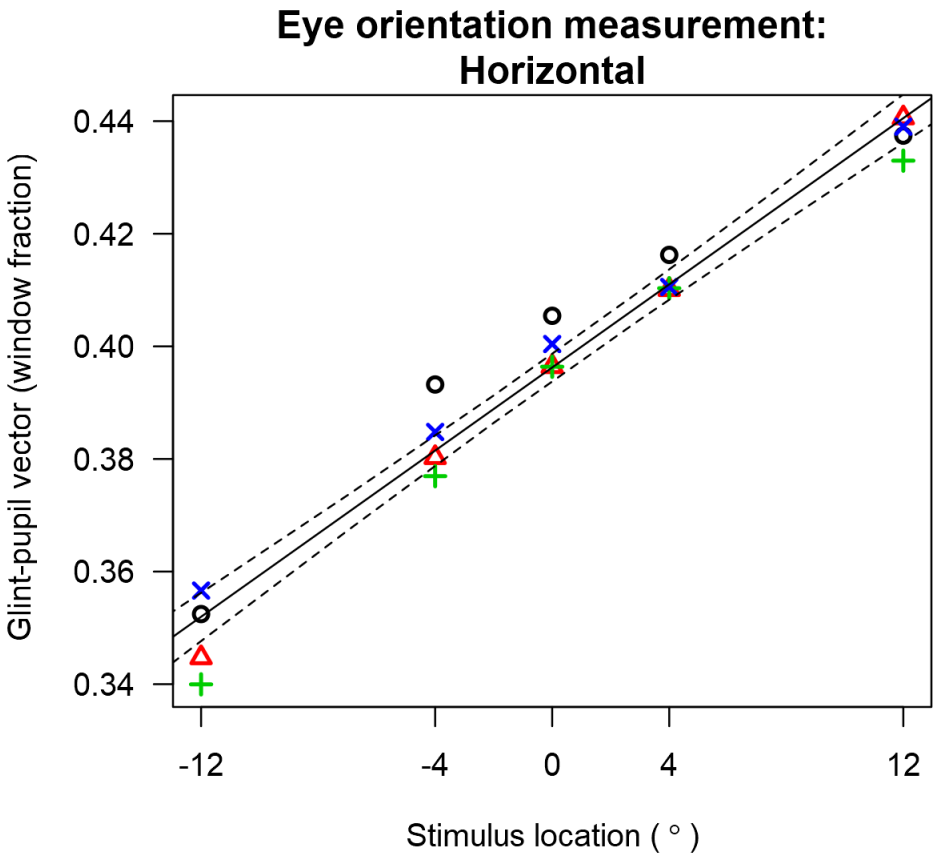

Figure 43 – GPV-eye rotation vertical relationship

Points each indicate the mean vertical GPV for each vertical stimulus location for one sighted subject. Each subject is shown with its own color and symbol. Dashed lines indicate the OLS regression line 95% confidence band. The solid black regression line has a slope of 0.0047. Vertical offsets among subject datasets reflect different vertical positions of the eye with respect to lens and display. Disregarding offsets, changes in GPV readings per degree were still very similar across subjects.

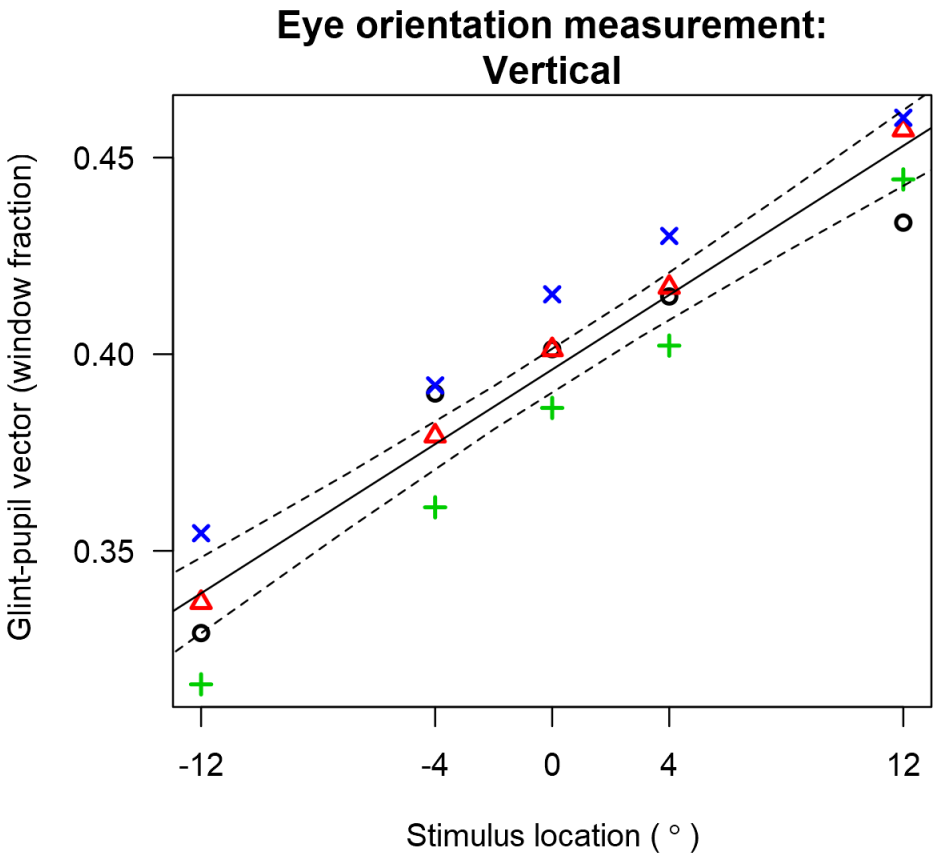

7.1.2.2 GPV-OCAP correlations

Although the relationships between eye orientations and OCAPs were not clean, some distinct correlations did emerge. Because subjects tended to look in different directions for the NL and PL tasks, and NL and PL GPVs were substantially different, GPVs from these tasks were considered separately. Looking between subjects, PL GPVs grossly preserved OCAP differences. As seen in Figure 44 and Figure 45, differences between subject OCAPs were echoed in PL GPV

differences. S2's PL GPVs were not as far leftward as OCAPs were, but PL GPVs corresponded more appropriately along the vertical dimension.

Figure 44 – Gross horizontal PL GPV, OCAP relationships, S1 and S2

Each point represents the initial positions of 1 day's horizontal OCAP and horizontal PL GPV intraday linear model with respect to time. Data for S2 are shown in blue, and data for S3 in red. There is some overlap in horizontal PL GPVs that does not exist in OCAPs. The average pattern S2's OCAPs being more left than S3's is preserved in PL GPVs, however.

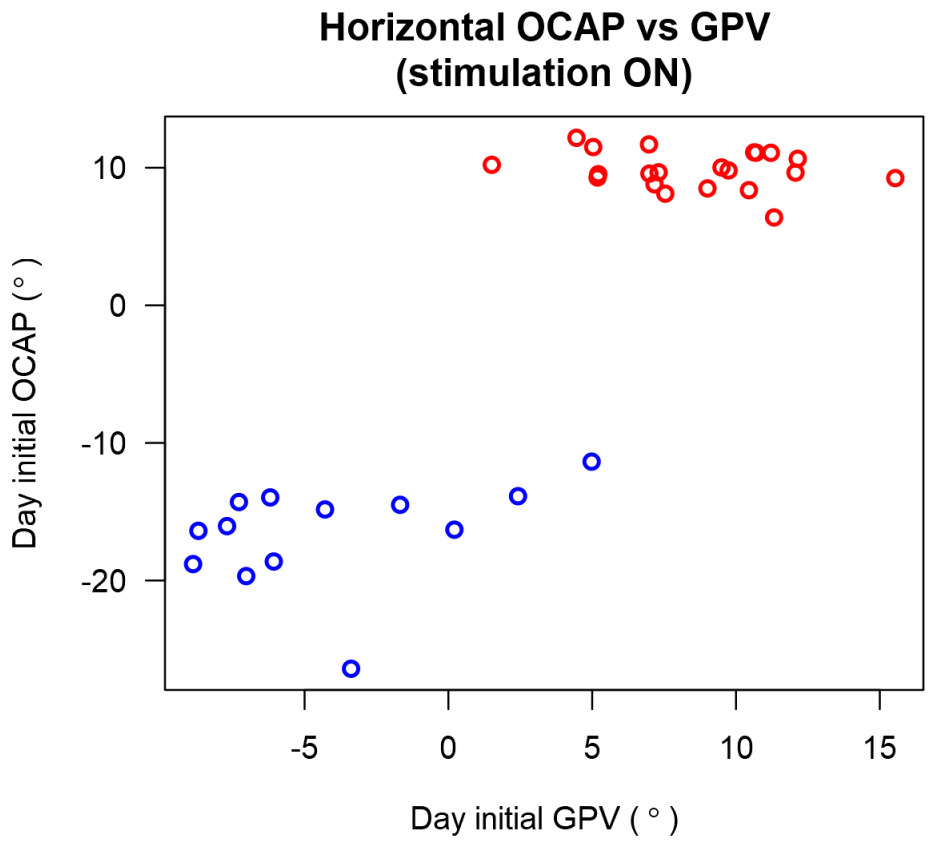

Figure 45 – Gross vertical PL GPV, OCAP relationships, S1 and S2

Each point represents the initial positions of 1 day's vertical OCAP and vertical PL GPV intraday linear model with respect to time. Data for S2 are shown in blue, and data for S3 in red. Vertical PL GPVs correspond to OCAPs more accurately than horizontal PL GPVs.

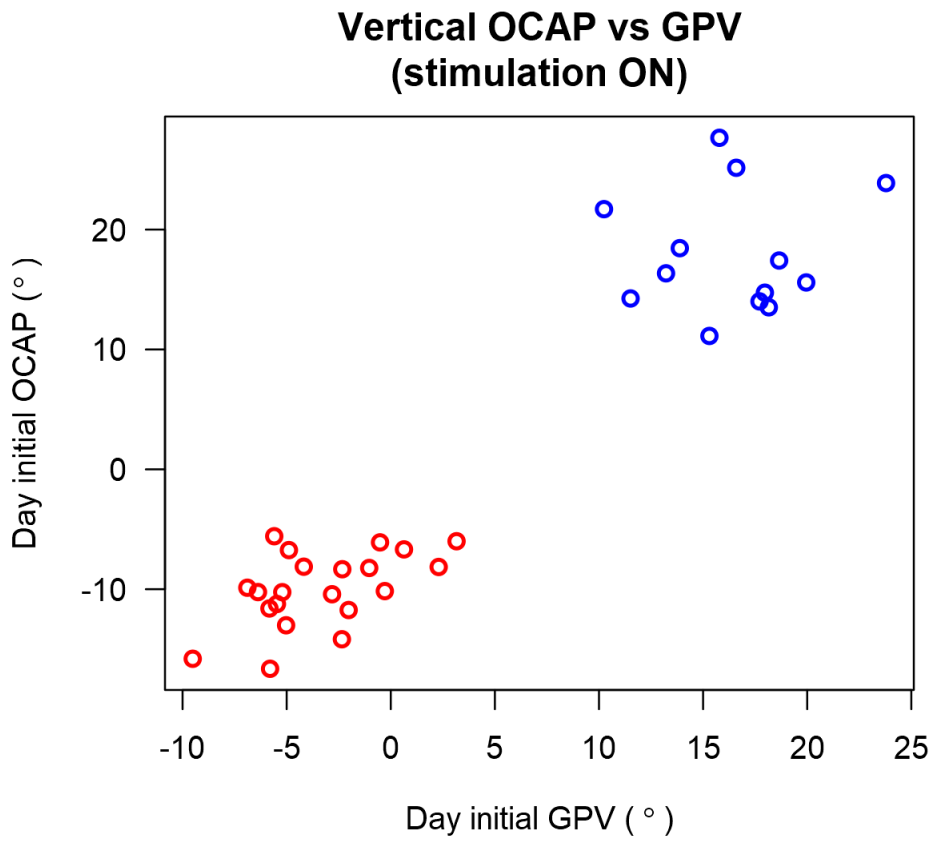

NL GPVs, shown in Figure 46 and Figure 47 did not correspond to OCAPs as well as PL GPVs. Along the horizontal dimension, NL GPVs varied little between S2 and S3. Vertical NL GPVs were closer to 0 than their PL GPV counterparts, making the divide between S2 and S3 less pronounced and reducing the correlation between NL GPVs and OCAPs. Given this apparent difference between NL and PL GPVs, quantitative GPV-OCAP analyses focused only PL GPVs.

Figure 46 – Gross horizontal NL GPV, OCAP relationships, S1 and S2

Each point represents the initial positions of 1 day's horizontal OCAP and horizontal NL GPV intraday linear model with respect to time. Data for S2 are shown in blue, and data for S3 in red. Unlike PL GPVs, NL GPVs showed little difference between subjects, despite large OCAP differences.

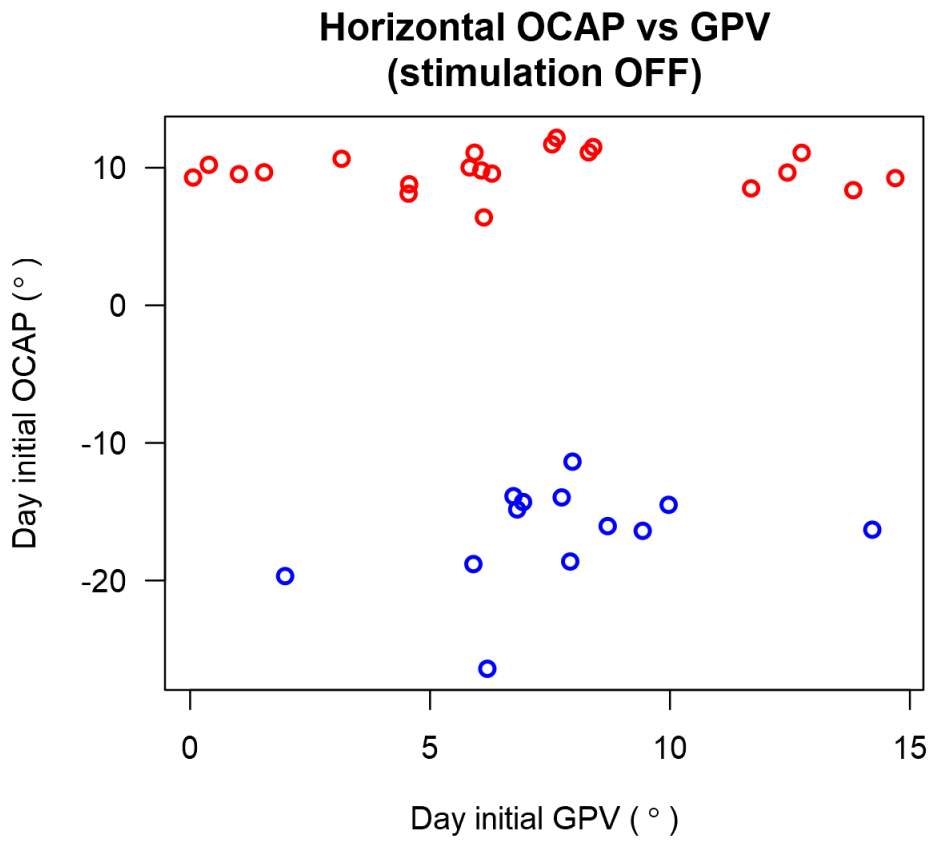

Figure 47 – Gross vertical NL GPV, OCAP relationships, S1 and S2

Each point represents the initial positions of 1 day's vertical OCAP and vertical NL GPV intraday linear model with respect to time. Data for S2 are shown in blue, and data for S3 in red. NL GPVs were on average higher for S2 than S3, as were OCAPs, but the correlation between GPVs and OCAPs was less pronounced than with PL GPVs,

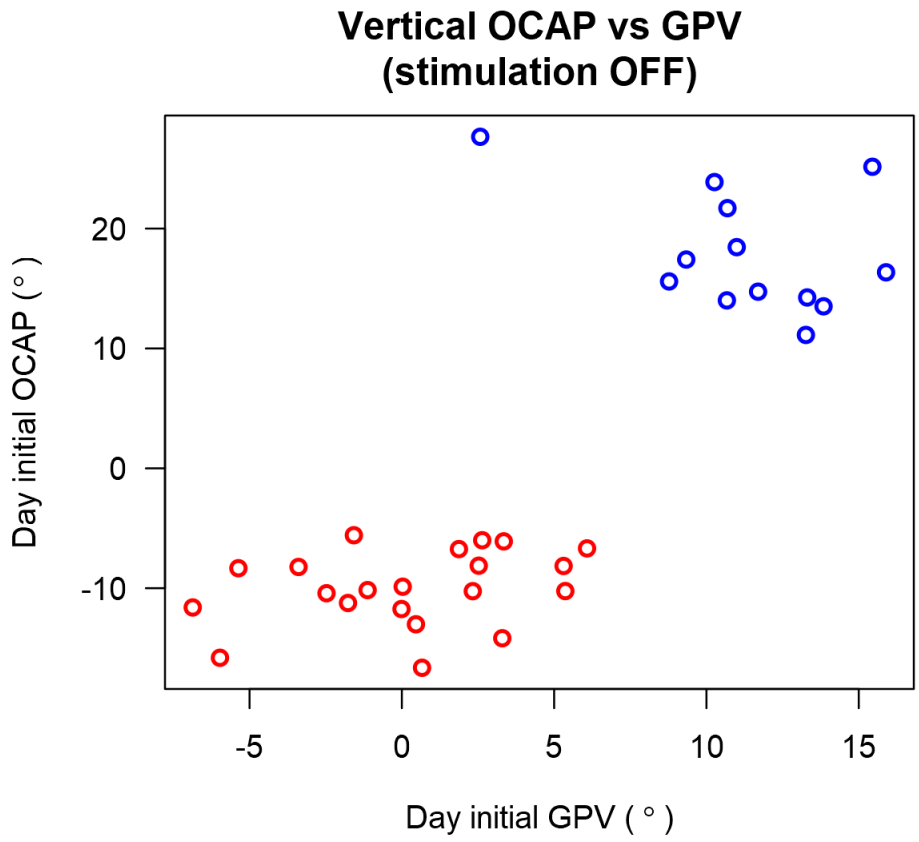

Although GPV model initial positions were favored for analysis, some consideration of model variability was required. In concept, the initial position of the intraday model should be less variable than central or final positions because it would not be as sensitive to how many trials were performed. Some GPV-OCAP correlations were made worse, however, when simple mean values were replaced by model initial positions in calculations. To investigate why this might happen, OLS regressions of model initial and final positions with respect to time across days were evaluated. Table 7 shows R^2 values for these regressions. As expected, almost all models had smaller residual variability with initial positions than final positions. This was reversed, however, for S3 in PL GPV models. S3's PL GPV model initial positions along the horizontal axis had an R^2 value almost an order of magnitude smaller than final positions.

Table 7 – Variability of intraday model points across days

R^2 values for OLS regressions of GPV intraday model initial and final positions with respect to time across days. R^2 values were larger for nearly all regressions using model initial positions, compared to those of regressions using final positions. The only exception was S3 along the horizontal axis with PL GPVs.

Subject	Horizontal GPV, date R^2				Vertical GPV, date R^2			
	NL model		PL model		NL model		PL model	
	Initial	Final	Initial	Final	Initial	Final	Initial	Final
S2	0.33	0.17	0.30	0.03	0.15	0.04	0.17	0.03
S3	0.004	0.0005	0.009	0.08	0.23	0.006	0.03	0.03

Considering the R^2 values in Table 7, calculation of correlation coefficients used all OCAP estimates for a given day matched with that day's PL GPV model initial position. The only exception to this was the calculation for S3 along the horizontal axis, for which PL GPV model final positions were used. Correlation coefficients and zero-correlation probabilities are shown in Table 8. Correlation coefficients were low, particularly given the variance individual-trial OCAP estimates, but were high enough to suggest a relationship between PL GPVs and OCAPs.

Table 8 – PL GPV, OCAP correlation coefficients

Pearson's correlation coefficient is given for each subject and dimension, comparing PL GPVs to OCAPs. OCAP estimates for each target localization trial were matched with the PL GPV model initial or final position for same day. Probabilities were determined by bootstrap resampling OCAP estimates assuming independence from PL GPVs. Although correlation coefficients were low, PL GPVs were related to OCAPs.

Subject	Dimension	GPV model point	Pearson's r	Probability
S2	Horizontal	Initial	0.07	< 0.008
	Vertical	Initial	0.14	$< 10^{-6}$
S3	Horizontal	Final	0.08	$< 8 \times 10^{-6}$
	Vertical	Initial	0.09	$< 10^{-6}$

PL GPV initial model positions were subtracted from OCAP model initial positions to determine how much of a correction PL GPV data could provide. Residual OCAPs and amounts of correction are shown in Table 9. On average, adjusting the camera based on the PL GPV would have accounted for 59% of the original OCAP.

Table 9 – Residual OCAPs after PL GPV correction

OCAP intraday model initial positions are shown alongside residual OCAPs after subtracting PL GPV intraday model initial positions. On average, adjusting the camera based on the PL GPV would have accounted for 59% of the original OCAP.

Subject	Dimension	Average OCAP initial position (°)		Correction
		Original	Residual	
S2	Horizontal	−17	−12	25%
	Vertical	18	1.6	91%
S3	Horizontal	9.8	1.3	87%
	Vertical	−10	−6.6	34%

Rates of OCAP change were similarly recalculated after subtracting corresponding PL GPV initial model positions. OLS regression lines for OCAP intraday model initial positions over time, with or without adjustment based on PL GPV, were calculated across days. Only trends that were

significantly different from 0 were considered for PL GPV adjustment. As shown in Table 10, PL GPVs accounted for an average of 94% of OCAP changes over time.

Table 10 – Residual OCAP rates of change after PL GPV correction

OLS regression was used to calculated average changes of OCAP model initial positions over time, with and without subtracting PL GPVs. S3's vertical rate of change was excluded because initial OCAPs did not show any significant trend with time. Subtracting PL GPVs counteracted an average of 94% of the remaining rates of change.

Subject	Dimension	OCAP rate of change (°/day)		Correction
		Original	Residual	
S2	Horizontal	−0.026	−0.0009	96%
	Vertical	−0.026	−0.009	66%
S3	Horizontal	−0.002	0.0004	120%

7.1.2.3 Intraday-interday GPV directional relationships

Relationships among intraday GPV trends showed the same patterns as OCAPs. NL and PL GPV trend relationships between days were calculated separately and then pooled for analysis. Intraday trend directions tended to have small differences from interday baseline directions, while differences from interday reset vectors averaged around 180°. Direction differences were stacked within 1° bins around the circle.

Figure 48 – GPV inter- and intraday vector directions, S2 and S3

Distribution of direction differences between GPV interday and intraday vectors. Direction differences had a mean confined to the 0° quadrant and were significantly nonuniform. Direction variance was significantly low. Direction correlation was positive, but not significant.

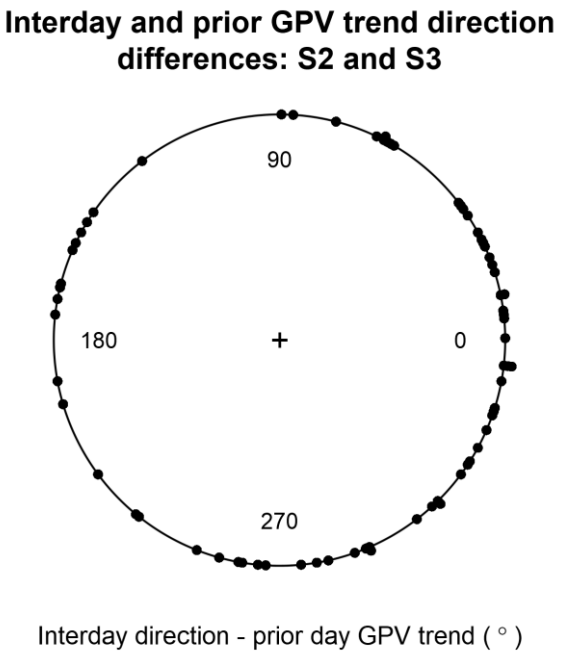

- Difference circular mean angle = −11°
  - Bootstrap 95% CI: −44°–19°
- Difference circular variance = 0.70
  - Probability of observing variance as low or lower assuming no correlation < 0.002
- Inter-run change, prior-run trend direction circular correlation coefficient = 0.17
  - No positive correlation probability > 0.07
- Fit probability:
  - uniform: < 0.01
  - von Mises: > 0.1

Figure 49 – GPV interday reset, S2 and S3

Distribution of direction differences between GPV interday reset and previous intraday trend vectors. Direction differences had a mean confined to the 180° quadrant and were significantly nonuniform. Direction variance was significantly low and the positive correlation was significantly high.

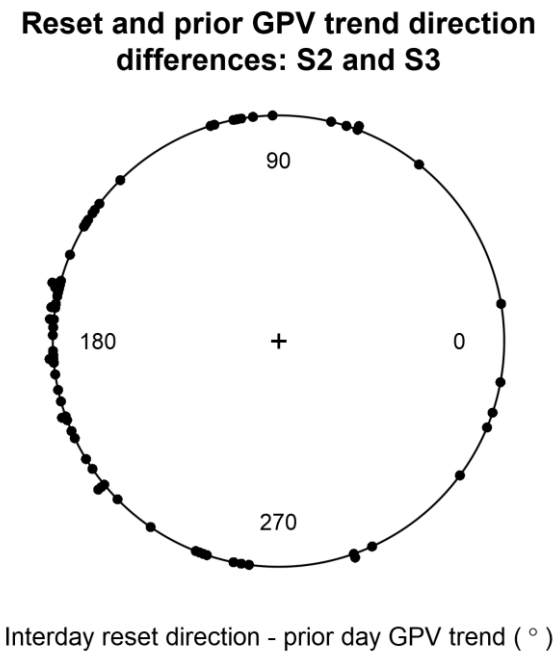

- Difference circular mean angle = 180°
  - Bootstrap 95% CI: 162°–197°
- Difference circular variance = 0.50
  - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Inter-run reset, prior-run trend direction circular correlation coefficient = 0.51
  - No positive correlation probability < 10^{-5}
- Fit probability:
  - uniform: < 0.01
  - von Mises: > 0.05

7.1.2.4 Intra-run, inter-run, intraday GPV directional relationships

Intra- and inter-run GPV behavior was the same as that of OCAPs. The vector between the start of one trial run to the next (inter-run vector) typically pointed in a direction similar to the first's intra-run trend, and intra-run trends tended to be reversed between trial runs. Just as with OCAPs, intra-run and inter-run vector directions were still inversely correlated. Inter-run vector directions correlated with those of reset vectors instead of intra-run vectors. Also like OCAPs, intra-run trend directions only had a modest tendency to align with intraday trend directions. Direction differences plotted below were stacked within 1° bins around the circle.

Figure 50 – GPV inter- and intra-run vector directions, S2 and S3

Distribution of direction differences between inter-run and intra-run GPV vectors. Direction differences had a mean confined to the 0° quadrant and were significantly nonuniform. Direction variance was significantly low, and directions were inversely correlated.

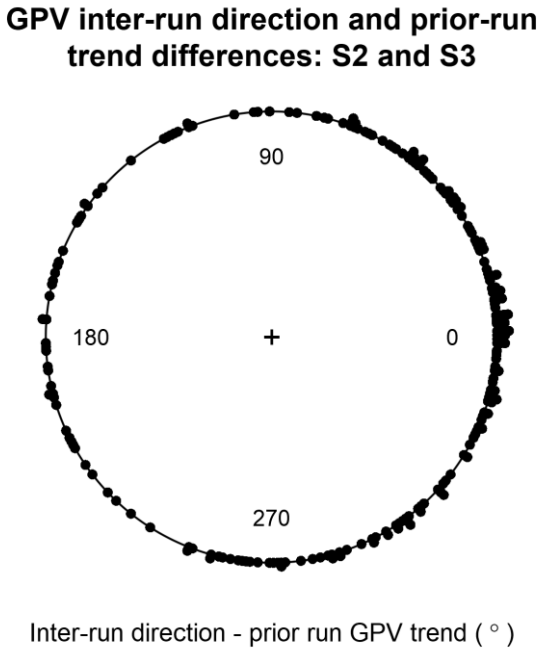

- Difference circular mean angle = −0.5°
  - Bootstrap 95% CI: −16°–14°
- Difference circular variance = 0.67
  - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Inter-run change, prior-run trend direction circular correlation coefficient = −0.22
  - No positive correlation probability > 0.999
- Fit probability:
  - uniform: < 0.01
  - von Mises: > 0.05

Figure 51 – GPV inter-run reset, S2 and S3

Distribution of direction differences between GPV run reset and intra-run vectors. Direction differences had a mean confined to the 180° quadrant and were significantly nonuniform. Direction variance was significantly low and the positive correlation was significantly high.

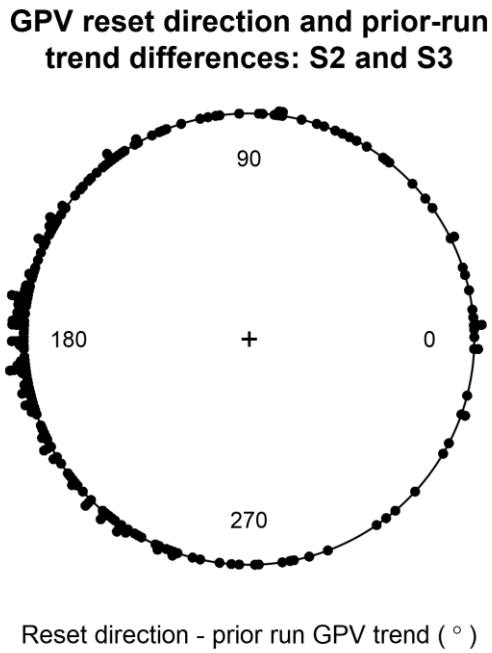

- Difference circular mean angle = 178°
    - Bootstrap 95% CI: 167°–188°
- Difference circular variance = 0.54
    - Probability of observing variance as low or lower assuming no correlation < 10^{-5}
- Inter-run reset, prior-run trend direction circular correlation coefficient = 0.50
    - No positive correlation probability < 10^{-5}
- Fit probability:
    - uniform: < 0.01
    - von Mises: < 0.01

Figure 52 – GPV intraday and intra-run trends, S2 and S3

Distribution of direction differences between intraday and intra-run vectors. Direction differences had a mean confined to the 0° semicircle and were significantly nonuniform. The negative correlation resulted from pooling NL and PL data.

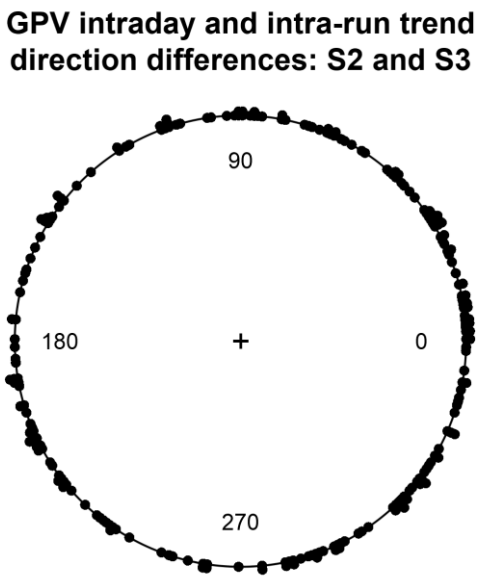

- Difference circular mean angle = 1.4°
  - Bootstrap 95% CI: −48°–47°
- Difference circular variance = 0.88
  - Probability of observing variance as low or lower assuming no correlation > 0.06
- Intraday, intra-run trend direction circular correlation coefficient = −0.06
  - No positive correlation probability > 0.8
- Fit probability:
  - uniform: < 0.03
  - von Mises: > 0.1

7.1.2.5 Cross-experiment GPV-OCAP directional relationships

Comparing intraday and intra-run trends across tests (OCAP/GPV measurements) did not yield very clean results. Intraday and intra-run trend directions did not match well within tests, and the lack of synchronization between OCAP and GPV measurements added to that variability.

Each OCAP and NL or PL GPV intra-run trend was compared to that day's NL or PL GPV and OCAP intraday trend, respectively. Comparisons using NL GPV intraday and intra-run trends (Figure 53) showed a significant correlation between the GPV and OCAP measurements, and direction differences were significantly nonuniform. These direction differences, while having a mean confined to the 0° semicircle, did not have a mean confidence interval that included 0°.

For comparison, Figure 54 shows comparisons using system-on PL GPV trends. Unlike for system-off NL GPV trends, differences between PL GPV and OCAP trends were quite uniform.

Figure 53 – Cross-measurement intraday and intra-run trends, NL GPVs and OCAPs, S2 and S3

Distribution of direction differences between intraday and intra-run trend vectors compared across NL GPV and OCAP measurements. Each point represents the difference between one trial run's trend direction from one measurement and the intraday trend direction of the other measurement. Direction differences had a mean confined to the 0° semicircle and were significantly nonuniform. 0° was, however, not inside the mean 95% confidence interval. Direction variance was significantly low and the positive correlation was significantly high.

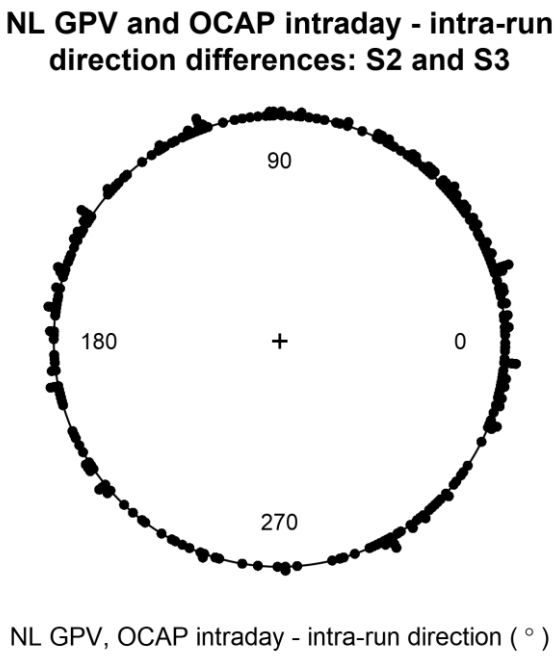

- Difference circular mean angle = 50°
  - Bootstrap 95% CI: 23°–81°
- Difference circular variance = 0.83
  - Probability of observing variance as low or lower assuming no correlation < 0.0006
- Intraday, intra-run trend direction circular correlation coefficient = 0.18
  - No positive correlation probability < 10^{-3}
- Fit probability:
  - uniform: < 0.01
  - von Mises: < 0.05

Figure 54 – Cross-measurement intraday and intra-run trends, PL GPVs and OCAPs, S2 and S3

Distribution of direction differences between intraday and intra-run trend vectors compared across PL GPV and OCAP measurements. Each point represents the difference between one trial run's trend direction from one measurement and the intraday trend direction of the other measurement. Direction differences were uniformly distributed around the circle.

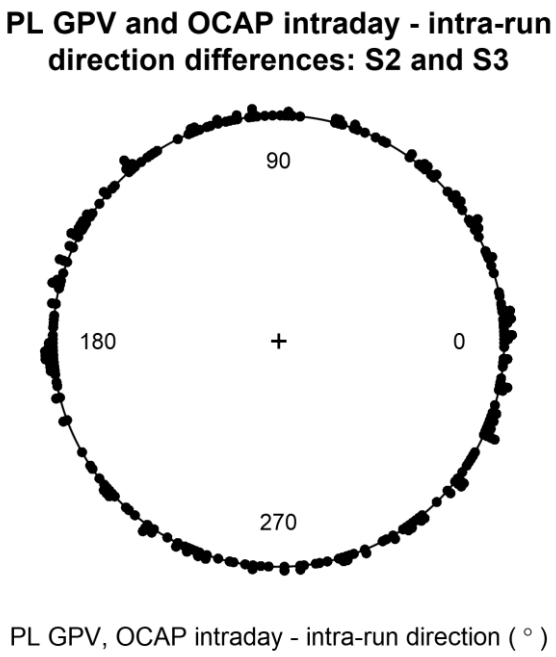

- Difference circular mean angle = 16°
  - Bootstrap 95% CI: −148°–182°
- Difference circular variance = 0.98
  - Probability of observing variance as low or lower assuming no correlation > 0.9
- Intraday, intra-run trend direction circular correlation coefficient = −0.15
  - No positive correlation probability > 0.99
- Fit probability:
  - uniform: > 0.1
  - von Mises: < 0.05

7.2 Simultaneous OCAP measurement and eye tracking

Separate eye-tracking and target localization tasks described in Section 7.1 demonstrated some connection between eye orientations and OCAPs. GPV and OCAP measurements taken at different times on different setups, however, fell short of permitting precise inspection of the eye orientation-OCAP relationship. Building upon the investigation started using the tower-mounted eye tracker, eye-tracking glasses were configured to permit simultaneous eye tracking and target localization. Not only did this allow monitoring eye orientations during the target localization task, but it also permitted automatic adjustment of the subject's CAP in response to eye movements. The acute study described below showed that eye orientation trends measured during trial runs matched trends in pointing biases, as well as demonstrated the utility of linking CAPs to eye orientations.

7.2.1 Methods

7.2.1.1 Subjects and equipment

Both S2 and S3 participated in this experiment, along with an additional subject, S4. S1 was unavailable for scheduling during these tests. S4 was implanted with an Argus II retinal prosthesis in July 2014 under an FDA HDE. Like S1–3, S4's prosthesis was implanted in the right eye. Each subject attended only 1 session on September 13th (S4) or 14th (S2 and S3) in 2016.

Standard Argus II glasses were replaced with SMI Eye Tracking Glasses 2.0. These lensless glasses had miniature cameras aimed at the eyes, and the frame around each eye was lined with 6 infrared illuminators. The glasses sent images of each eye to a connected laptop computer for processing in eye-tracking software. The software constructed a model of each eye in a process of self-calibration that did not require responses to stimuli. Instead, subjects moved their eyes

over large ranges for about 1–2 minutes. The software used the locations of the pupil relative to the six corneal reflections to develop its model of each eye, and then subsequently to estimate directions of gaze. Different nosepieces were used to ensure eyes were well-positioned with respect to the eye-tracking cameras.

An additional scene camera was attached to the bridge of the glasses, similarly to the configuration on the Argus II glasses. The lens attached to the camera permitted a FOV ranging 73° × 55°. The camera sent a 640 × 480 pixel image to a laptop computer, and an 18° × 11° area was isolated for use with the subject's Argus II. This image subsection was sent to the VPU in place of the image normally obtained by the Argus II camera. The VPU's CAP setting was always centered at (0°,0°) to ensure that subsection provided by the computer was processed in its an entirety, and nothing else. The transceiver coil described in Section 7.1.1.1 for use with PL was used for this setup to transmit camera-driven stimulation commands from the VPU to the subject's implant.

The computer processing the scene camera's input assumed the functional role of implementing any CAP setting. A simple, constant CAP setting mimicked the paradigm of the Argus II: one subsection of the image was always processed, and head movement was the only method of scanning a scene. Alternatively, the computer could use eye-tracking to adjust the CAP in response to eye movements. Eye-tracking measurements were made at 60 Hz, and CAP adjustments were made every 0.01–0.02 seconds. Modifying the CAP based on eye movements added the capacity for scanning a scene with both head and eye movements. These two modes, constant and dynamics CAPs, were referred to as head-only and eye-head scanning conditions. Eye tracking data were collected regardless of condition. Eye-tracking-based CAP adjustments in eye-head scanning were added to a baseline constant CAP that was set to improve accuracy.

The eye-tracking were also equipped with an inertial measurement unit (IMU). The IMU provided linear acceleration data along all 3 spatial axes, as well as angular velocities along yaw, pitch, and roll dimensions. Eye-tracking, IMU, and target localization data were all synchronized to have matching timestamps.

7.2.1.2 Study procedures

Subjects were asked to perform target localization using the eye-tracking glasses. Subjects sat 36–40 cm from the touchscreen and only responded with their dominant hands. Each subject performed 5–6 trial runs of 20 trials each. At least 2 trial runs were performed with each scanning condition, with at least 2 eye-head scanning trial runs directly followed by 2 head-only scanning runs. The experimenter informed the subject of scanning condition before trial run. When eye-head scanning was enabled, subjects were encouraged to use eye movements for scanning the screen.

7.2.1.3 Analysis of synchronized eye-tracking, localization, and head-motion data

Ordinary least squares linear models were calculated for localization errors and subjects' points of regard (PORs) with respect to time within each trial run. Separate models were calculated for horizontal and vertical dimensions, and then combined as a 3D vector. The directions of the vectors in the horizontal-vertical plane were computed, and vectors for localization errors and PORs were compared within each trial run, using the same direction difference calculations described in Section 5.1.4. Analyses focused on the pooled data of all 3 subjects.

Eye-tracking readings were recorded continuously throughout trial runs, including time before and after trials. The time when a target was in view was not recorded. Trends in eye orientations

within trial runs were calculated in two ways: either using all readings from a trial run, or only using readings from specified intervals relative to trial starts or ends.

Head motion, and corresponding eye movement via the vestibulo-ocular reflex (VOR), was suspected to interfere with measuring subtle changes in average eye orientation, particularly in the head-only scanning condition. To analyze the contribution of VOR to observed eye movements, eye movement velocities were calculated from eye orientation readings and timestamps. Eye movements related to VOR can be expected to inversely correlate with head rotations. Normalized eye movement and head angular velocities from one example trial with head-only scanning were plotted to demonstrate such a suggestion of VOR influence. Head angular velocities were inverted in the plot to show how the two velocities corresponded more clearly.

IMU-measured angular velocities (yaw and pitch) were correlated with computed eye movement velocities (horizontal and vertical) to quantify contributions of VOR to eye movements within subsections of trials. If some time period typically had reduced VOR influence, such as at the beginning or end trials, focusing analyses on trial-to-trial differences within that time period may be more effective for monitoring trends in average eye orientation. Within each subject and scanning mode, trials were broken into 0.5s periods and pooled. Within each period, Pearson's correlation coefficient was calculated for head angular and POR velocities.

Mean correlation coefficients were calculated for periods covering time up to the median trial completion time of 7.0s. Each mean was based on 12 observations, one coefficient for each of: 3 subjects × 2 scanning modes (eye-head/head-only) × 2 dimensions (horizontal/vertical). Significance was tested by creating 10^5 bootstrap period samples using stratified resampling.

Each bootstrap sample contained 12 coefficients, one from each data stratum, each taken from a random period through the median trial completion time. Correlation coefficients closer to 0 in the initial period may indicate less influence of VOR on POR during that time. Trend direction differences between errors and eye orientations were reanalyzed considering only the readings within the period with the small absolute correlation.

Rates intra-run error trends were also compared between scanning conditions. More trial runs were performed with eye-head scanning, so testing the significance of the difference between conditions focused on resampling eye-head scanning data. A bootstrap distribution of eye-head scanning error rates of change was developed using stratified resampling within subject-dimension (H/V) strata, and the observed error rate of change was for head-only scanning was compared to that distribution.

7.2.2 Results

7.2.2.1 Trend direction differences considering all eye-orientation readings

Taken altogether, localization error trends tended to point in the same direction as trends in eye orientation. When considering all eye-tracking readings throughout each trial run, error and eye trends only matched when eye-head scanning was enabled. Eye readings from head-only scanning runs did not have prevalent trends pointed any consistent direction relative to that of the trend in errors. The bias seen when considering direction differences in all trial runs was largely carried by eye-head scanning runs. Direction differences show below were plotted within 1° bins around the circle.

Figure 55 – Trend direction differences, all eye readings, both scanning conditions, S2, S3, and S4

Distribution of direction differences between eye-orientation and localization-error intra-run trends considering all eye-tracking readings, from both head-only and eye-head scanning runs. Direction differences had a mean confined to the 0° semicircle and were significantly nonuniform.

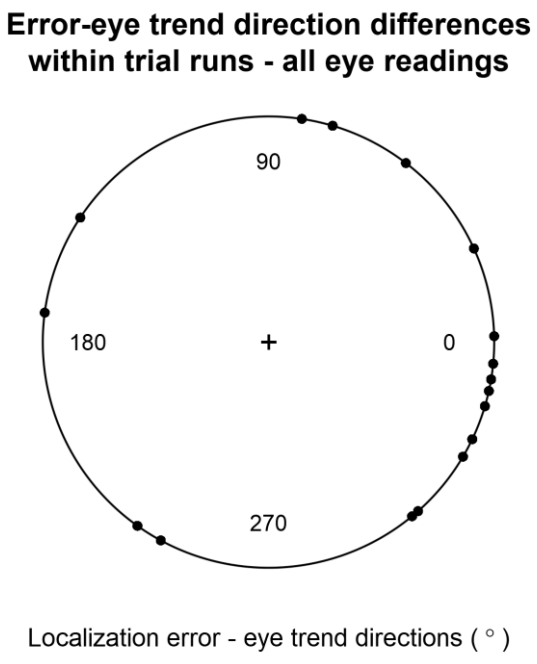

- Difference circular mean angle = −8.3°
  - Bootstrap 95% CI: −53°–45°
- Difference circular variance = 0.58
  - Probability of observing variance as low or lower assuming no correlation > 0.06
- Error-eye intra-run trend directions circular correlation coefficient = 0.12
  - No positive correlation probability > 0.3
- Fit probability:
  - uniform: < 0.05
  - von Mises: > 0.1

Figure 56 – Trend direction differences, all eye readings, head-only scanning, S2, S3, and S4

Distribution of direction differences between eye-orientation and localization-error intra-run trends for head-only scanning runs, considering all eye-tracking readings. There was no clear relationship between eye-orientation and localization-error trends.

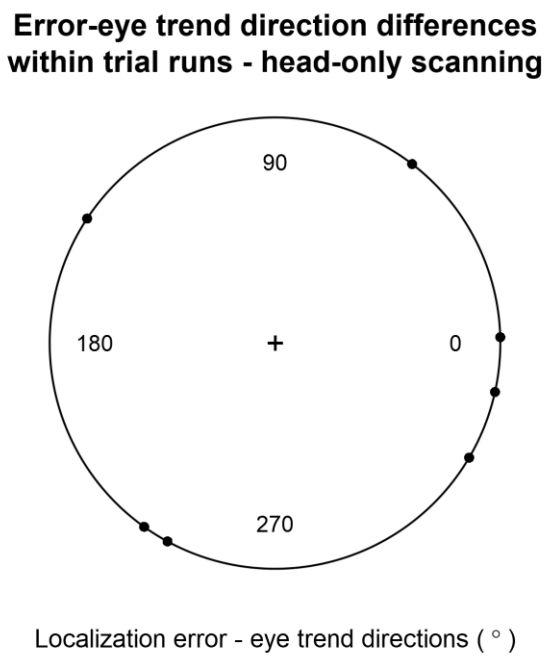

- Difference circular mean angle = −34°
    - Bootstrap 95% CI: −184°–103°
- Difference circular variance = 0.73
    - Probability of observing variance as low or lower assuming no correlation > 0.7
- Error-eye intra-run trend directions circular correlation coefficient = −0.56
    - No positive correlation probability > 0.8
- Fit probability:
    - uniform: > 0.1
    - von Mises: > 0.1

Figure 57 – Trend direction differences, all eye readings, eye-head scanning, S2, S3, and S4

Distribution of direction differences between eye-orientation and localization-error intra-run trends for eye-head scanning runs, considering all eye-tracking readings. Direction differences had a mean confined to the 0° semicircle and were significantly nonuniform.

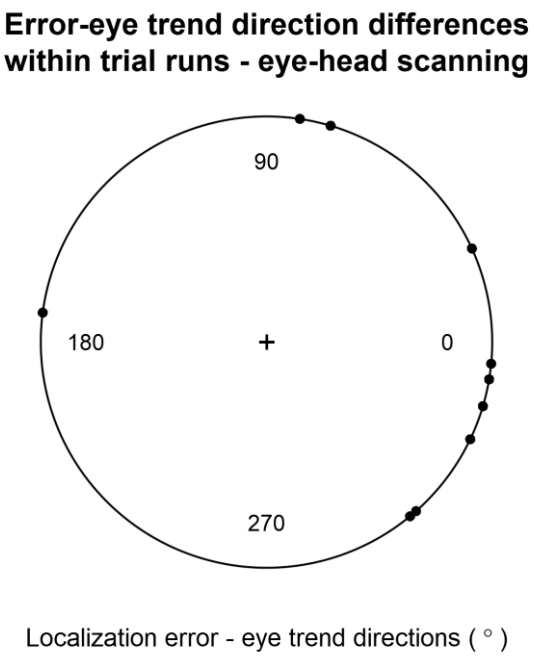

- Difference circular mean angle = 0.3°
  - Bootstrap 95% CI: −30°–61°
- Difference circular variance = 0.45
  - Probability of observing variance as low or lower assuming no correlation > 0.05
- Error-eye intra-run trend directions circular correlation coefficient = 0.17
  - No positive correlation probability > 0.2
- Fit probability:
  - uniform: < 0.05
  - von Mises: > 0.1

7.2.2.2 Head motion and VOR

Eye-movement velocities had prominent inverse correlations with head angular velocities in head-only trial runs, as shown with inverted head velocities in Figure 58 and Figure 59. Eye movements dominated by VOR should move in the opposite direction of any head movement, thus VOR-based eye movements should match normalized head movements multiplied by −1. The latter relationship is shown for simplicity. Eye and head movements did not have as clear a relationship in eye-head scanning trial runs.

Figure 58 – Head and eye horizontal velocity measurements

Eye-movement velocities for one example head-only scanning trial run are shown as black circles. Inverse IMU-measured angular velocities are shown as a blue line. The blue line was chosen to represent head velocities instead of separate blue points for clarity. Both velocities were normalized (centered and rescaled) and are plotted against time. Eye movement velocities and inverse head velocities generally matched well in head-only scanning trials.

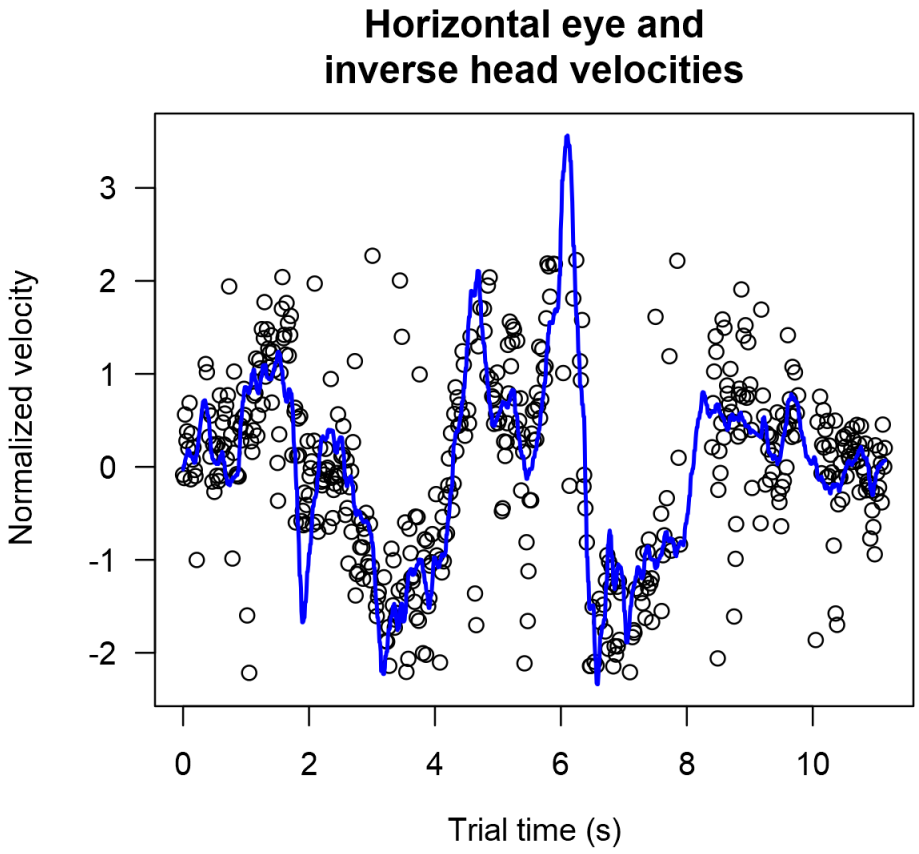

Figure 59 – Head and eye vertical velocity measurements

Eye-movement velocities for the same head-only scanning trial run as for Figure 58 are shown as black circles. Inverse IMU-measured angular velocities are shown as a blue line. The blue line was chosen to represent head velocities instead of separate blue points for clarity. Both velocities were normalized (centered and rescaled) and are plotted against time. Eye movement velocities and inverse head velocities generally matched well in head-only scanning trials.

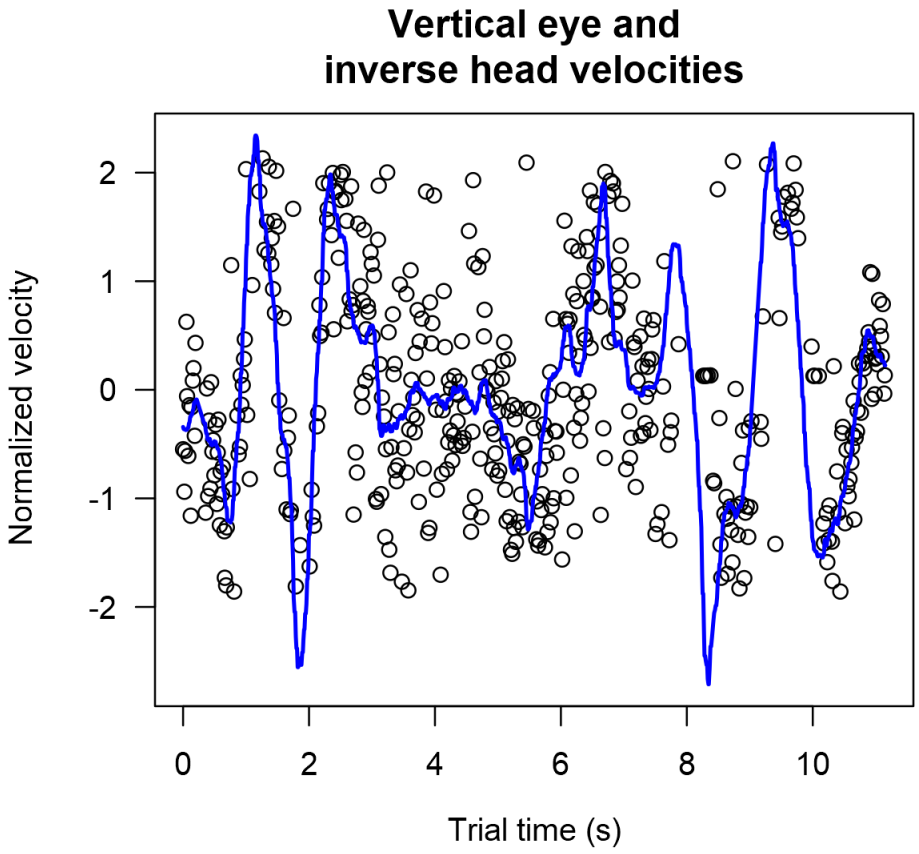

Head-only scanning trial runs were collectively analyzed to determine whether any 0.5s time periods through 7s had significantly less correlation between head and eye movements. As shown in Figure 60, the first 0.5s period after the beginning of the trial did not show as strong a correlation as later periods. The mean correlation coefficient for this 0–0.5s period was 20% less negative than the mean of the other period coefficients ($p < 0.0006$), and thus likely had less contribution of VOR.

Figure 60 – Head and eye velocity correlation coefficients

Mean correlation coefficients of head angular velocities and eye-movement velocities pooled within 0.5s periods of elapsed trial time. Head and eye-movement velocities had significantly less negative correlation coefficients in the 0–0.5s period than other periods.

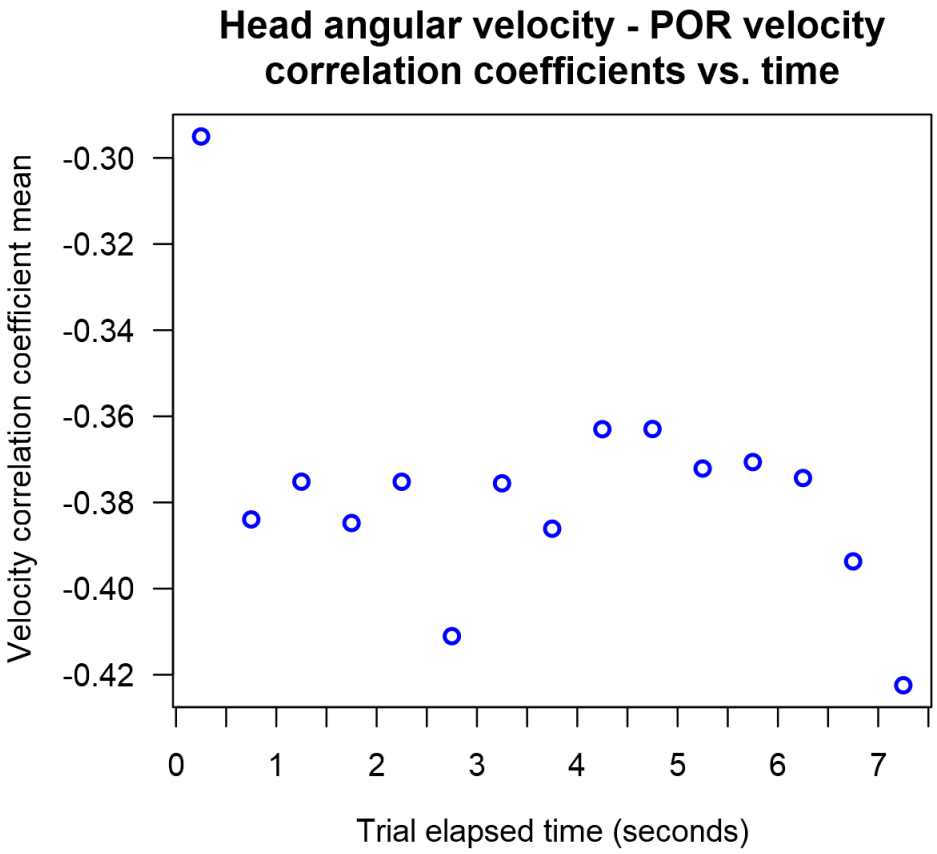

7.2.2.3 Trend direction differences with focus on trial times 0–0.5s

Using eye-orientation data only from the first 0.5s of each trial to development trial run linear models yielded trend directions that were significantly biased toward the directions of localization error trends. Changing the analyzed timed window did not affect the localization error trends, because 1 data point represented the entirety of each trial's error. When pooling direction differences from both scanning conditions, the distribution of differences was significantly nonuniform and biased toward 0°. Variance was significantly low, and the positive

correlation was significantly high. Similar patterns can be seen when head-only scanning is isolated, but results are not as significant because there fewer data points. The distribution for eye-head scanning still appears biased, but no results were significant with this time filter.

Figure 61 – Trend direction differences, trial time 0–0.5s, both scanning conditions, S2, S3, and S4

Distribution of direction differences between eye-orientation and localization-error intra-run trends considering eye-tracking readings from only the first 0–0.5s of each trial, from both head-only and eye-head scanning runs. Direction differences had a mean confined to the 0° semicircle and were significantly nonuniform. Direction variance was significantly low and the positive correlation was significantly high.

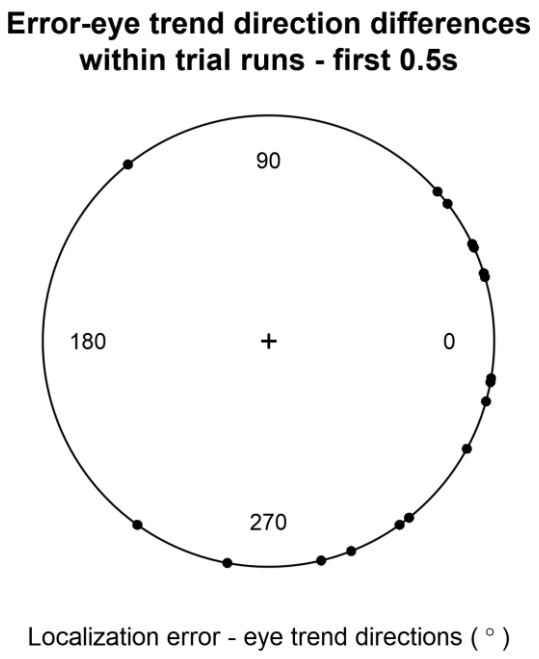

- Difference circular mean angle = −17°
  - Bootstrap 95% CI: −51°–11°
- Difference circular variance = 0.41
  - Probability of observing variance as low or lower assuming no correlation < 0.005
- Error-eye intra-run trend directions circular correlation coefficient = 0.64
  - No positive correlation probability < 0.003
- Fit probability:
  - uniform: < 0.01
  - von Mises: > 0.1

Figure 62 – Trend direction differences, trial time 0–0.5s, head-only scanning, S2, S3, and S4

Distribution of direction differences between eye-orientation and localization-error intra-run trends from head-only scanning runs, considering eye-tracking readings from only the first 0–0.5s of each trial. Direction differences had a mean confined to the 0° semicircle and were significantly nonuniform. Correlation between trend directions was significantly high.

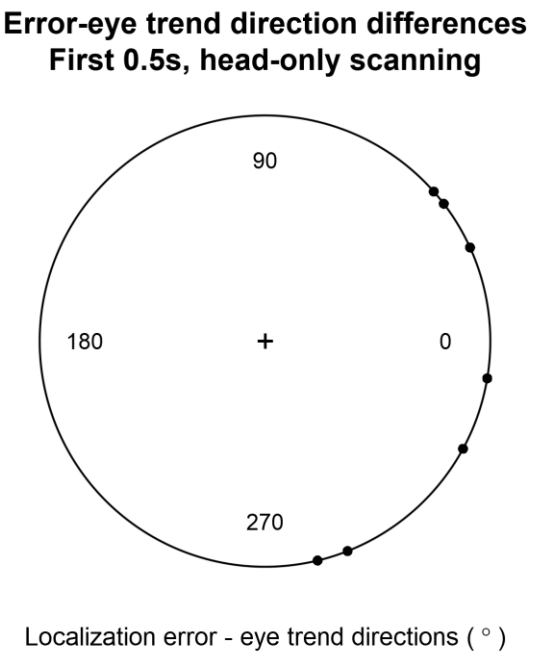

- Difference circular mean angle = −10°
  - Bootstrap 95% CI: −50°–22°
- Difference circular variance = 0.29
  - Probability of observing variance as low or lower assuming no correlation > 0.08
- Error-eye intra-run trend directions circular correlation coefficient = 0.76
  - No positive correlation probability < 0.03
- Fit probability:
  - uniform: < 0.05
  - von Mises: > 0.1

Figure 63 – Trend direction differences, trial time 0–0.5s, eye-head scanning, S2, S3, and S4

Distribution of direction differences between eye-orientation and localization-error intra-run trends from eye-head scanning runs, considering eye-tracking readings from only the first 0–0.5s of each trial. Direction differences had a mean confined to the 0° semicircle and were significantly nonuniform.

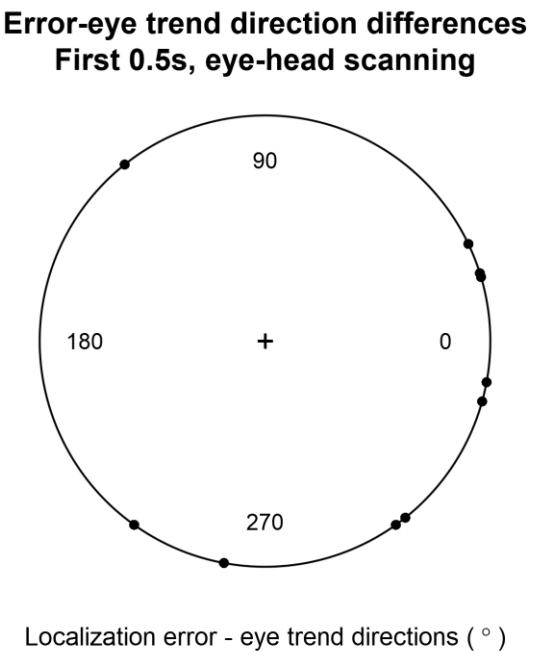

- Difference circular mean angle = −24°
  - Bootstrap 95% CI: −85°–22°
- Difference circular variance = 0.50
  - Probability of observing variance as low or lower assuming no correlation > 0.08
- Error-eye intra-run trend directions circular correlation coefficient = −37
  - No positive correlation probability > 0.8
- Fit probability:
  - uniform: > 0.05
  - von Mises: > 0.1

7.2.2.4 Rates of localization error change

Localization errors had an average intra-run rate of change of 0.03°/day, with a maximum of 0.15°/day. Considering scanning modes separately, errors changed an average rate of 0.024°/day for eye-head scanning, and 0.040°/day with head-only scanning. Based on bootstrap resampling eye-head scanning data, observing a mean rate of change as high as 0.040°/day was less than 5×10^{-6}. Removing the most extreme head-only rate of change (0.15°/day) as an outlier reduced the head-only average rate of change to 0.032°/day, which was still significantly higher than the eye-head average of change ($p < 0.02$).

7.3 Discussion

Repeated measurements from the eye-tracking tower and the acute head-mounted eye-tracking experiment provided evidence linking OCAP behavior to changes in eye orientations. Comparing the separate GPV and OCAP measurements showed a relationship in how each changed over the course of weeks and months. Given differences in how each were measured, though, the dual-setup configuration only weakly suggested intraday similarities in trends. In contrast, although the head-mounted eye tracker could only be used for 2 days, simultaneous eye tracking and target localization showed a robust relationship between eye orientations and localization errors on the timescale of a single trial run.

Correlating eye orientations from the eye-tracking tower to OCAPs was very sensitive to the task used for measuring GPVs. Figure 44 and Figure 45 showed that GPVs measured in response to visual stimulation, when comparing between subjects, grossly matched the biases observed with OCAPs. Figure 46 and Figure 47, on the other hand, showed that GPVs during nose localization matched OCAPs far less reliably. Horizontal NL GPVs did not seem to adhere to any relationship at all with OCAPs. Vertical NL GPVs reflected OCAP biases, but were held closer to 0° than their PL counterparts. Even though no specific instruction for where to look was provided during NL, the task may have motivated subjects to orient their eyes closer to their noses. NL may have thereby anchored eye orientations away from those favored during stimulation, reducing or even completely eliminating biases expected from subjects based on OCAPs.

Because PL GPVs were generally better reflections of OCAPs than NL GPVs, correlations across sessions focused only on PL GPVs. Individual OCAP and GPV measurements could not be paired in any reasonable way, so OCAP estimates for each trial were paired with an endpoint of the

same day's intraday PL GPV model. The reverse, matching PL GPVs to OCAP model endpoints, was disfavored simply because more OCAP data points were taken each day, and the increased data count increased analysis power. When considering OCAPs and their corresponding initial GPV position, or final position in the one case for S3, OCAPs and GPVs significantly correlated along both dimensions.

Significant correlations indicated some link between the datasets, but did not indicate the relevance of any such link. In an attempt to address the importance of the relationship between OCAPs and eye orientations, OCAPs were reanalyzed assuming that any change in GPV could be negated by an automatic adjustment of the camera. These hypothetical adjustments would have all reduced the CAP modification necessary to optimize accuracy, as GPVs were never oriented opposite to OCAPs. Residual calculated OCAPs were on average 59% smaller than the original OCAPs, reductions ranging 25–91%. Calculated residual rates of OCAP change were similarly reduced by an average of 94%. If overshoots were to be penalized, converting S3's overcompensation of 120% to an 80% reduction, then the mean residual rate of OCAP change would be 81% smaller than the original.

Simulated camera adjustments based on given GPVs were not always helpful, however, and did present a problem with OCAPs that did not significantly change. Adjustment of S3's camera along the vertical axis would have counteracted the nonsignificant observed rate of change by over 400%. One might hope that such disparity in eye-orientation and OCAP measurements might not arise when eye tracking and compensation are present and active during the pointing task. Nevertheless, the remainder of the simulated adjustments suggest that eye orientations can still account for a large portion of OCAP eccentricities and changes. A system could

conceivably compensate for 59% of OCAP eccentricity and 81% of OCAP changes by using only eye-tracking data.

Further supporting the role of eye orientations in how OCAPs shift over time, GPVs displayed the same behavior in intraday and intra-run trends as OCAPs. Just like OCAPs, GPVs shifted within days, and the direction of an intraday trend typically had a small angular difference from the change between sessions. From the end of one session to the beginning the next, GPVs tended to reverse the direction of change seen with the prior day's GPV trend. OCAPs not only displayed this reset pattern between days, but also between localization trial runs. Interestingly, GPVs also demonstrated resets between NL and PL trial runs. When comparing intra-run trends to intraday trends, both OCAPs and GPVs displayed a weak, yet significant, bias for trends to point in the same direction. GPV intra-run and intraday trends were negatively correlated, while OCAP trends were positively correlated, but this negatively correlation disappeared when NL and PL trends were analyzed separately. Similar behavior on trial-run and test-session scales would be required if eye orientations drive changes in OCAPs, and these data demonstrate such commonalities exist.

Directly linking eye orientation trends with OCAP trends was not as simple as relating trends within experiments. Intraday trends between experiments were not similar enough for any significant results. Individual trial-run trends needed to be considered to increase test power. Intra-run trend directions were all therefore compared to the intraday trend direction of the other experiment. Significant results only manifested when NL and PL GPV trends were separated. Unlike with correlations across days, however, only NL GPV trend directions significantly matched OCAP directions; no such correlation was evident with PL GPVs. NL GPV and OCAP direction differences were widely distributed, but the difference mean and variance

and direction correlation were all significantly suggestive of a bias for trends to align. PL GPV and OCAP direction differences were uniformly distributed. Although anchoring eye orientation to the nose may have counteracted any trends across days, the reduced variance in eye orientation may have made subtle intraday GPV trends more detectable.

Separate eye-tracking and target-localization setups were clearly insufficient for demonstrating a direct link between eye orientation and OCAP changes, at least on short timescales. Simultaneous eye tracking and target localization with the head-mounted eye tracker, however, was well-suited for exploring a direct relationship. Although limited accessibility to the head-mounted eye tracker made regular assessments impractical, and thus precluded examining long-term eye and localization behavior, enough data were collected to show similarities within trial runs.

Without any filtering, localization errors and eye orientations had intra-run trends that tended to point in the same direction. Some tendency for trend alignment may have been present in head-only scanning, but if so, only enough to bolster the alignment seen with eye-head scanning. No significant bias in trend direction differences could be seen within head-only scanning. Eye orientations during head-only scanning were likely dominated by VOR, but overall trends in head motion did not appear to translate into trends in errors. VOR-induced eye movements at the time of viewing the target likely would have produced corresponding localization errors, but all other VOR eye movements would have been largely unrelated to response error.

Eye movements with eye-head scanning were less subject to VOR, and eye displacements caused by VOR would not necessarily create localization errors in this condition. If eye orientations tended to move in a particular direction through a trial run, localization errors

would not be expected to follow in the same way as for head-only scanning. For example, if a subject's average eye orientation tends to move upward through a trial run, the average height of the CAP and the scanned location on the screen would also go up. Although the eye is pointed higher, the target will only be seen if it too is higher. Rather than localization errors being caused by eye-camera misalignments, in which case you would expect errors to grow in the same direction as eye-orientation shifts, errors may be related more to pointing imprecision and only viewing a side of the target. If the average gaze angle is higher, it becomes more likely that a subject will first strike upon the top of the target. As targets were about 7° in diameter, a viewing bias can start a subject with up to a 3.5° pointing bias, not considering any other sources of pointing bias. The alignment in trends seen in eye-head scanning were thus more likely eye-orientation shifts affecting viewing bias rather than eye-camera misalignment.

VOR appeared to have less influence on head-only scanning trials within the first 0.5s of each trial. Analyzing trial-to-trial trends using only that timeframe yielded trend directions that matched error trends very well. Root-mean-square head angular velocities were well above subject minimums during this period, so this period was not ideal because subjects were most still. Subjects may have been moving their eyes in advance of head motions to look for the target. Head movements would be initiated with some small delay after eye movements, so the two would not correlate, and the eye movements could still provide meaningful information. If there is a change in resting eye orientation, that change might be expected to persist through the trial run and generate errors through eye-camera misalignments. The direction and magnitude of initial eye movements, relating to search patterns, could also provide information on viewing bias during the trial. Trends observed with trial-start periods could thus provide information that affects localization errors. Trend directions were significantly aligned after

applying this time filter, and correlations may be even stronger if times of stimulation could be isolated.

Comparing rates of error trends between head-only scanning and eye-head scanning indicated that errors progressed at faster rate with head-only scanning. Eye-head scanning reduced the localization error rate of change by 24%–40%. While this reduction is more modest than that predicted by subtracting tower-mount GPVs from OCAPs, it is consistent in suggesting that eye-tracking can significantly slow down error rates of change, but not completely prevent shifts in localization errors. Greater reduction may have been observed with eye-head scanning if smaller targets were used and errors within the radius of the target, but not at the center of the target, were less prominent in analyses.

8 General discussion

The studies described in Chapters 5–7 characterized several time-dependent attributes of CAPs and localization accuracy. First, the results in Chapter 5 demonstrated that Argus II users could not passively adapt to camera misalignments. CAPs need to be specifically configured for each user in order to optimize accuracy. Further, as shown in Chapter 6, any constant CAP applied for a user might only be applicable in a small fraction of other randomly chosen timepoints. OCAPs changed with time in slow, but steady progressions. Eye-tracking experiments detailed in Chapter 7 suggest that at least part of the observed OCAP changes were related to how average eye orientations were drifting. Eye orientations correlated with OCAPs across days, and the two measurements followed the same reset behavior between sessions and between trial runs. When eyes were tracked during target localization runs, eye orientations and localization errors tended to move in the same directions. Altogether, this research suggests that optimal accuracy with remote-camera prostheses is complicated by eye-orientation behavior, and that prosthesis-integrated eye tracking can at least partially compensate for the gradual degradation of accuracy.

8.1 Lack of passive adaptation

When one considers perceptual misalignment, it is readily imagined as being observable by the subject. Knowing that there is some misalignment, the imagined visual scene is shifted to no longer correspond with reality. For example, one can imagine that a subject reaches for a plate, but reaches much too far to the right. That imagined error is easy to comprehend, and the solution would seem obvious. The subject must try reaching farther left until the reaching motion is properly aimed toward the plate. Melding that idea with camera misalignment would

lead one to expect that a prosthesis user should detect a misalignment and learn corrective behavior, at least through trial and error.

Such a mental image of misalignment and correction is consistent with experiments involving altered visual fields. As discussed in Section 2.5, normally-sighted subjects can successfully adapt to inverted or translated visual fields within hours or days, depending on the degree of alteration. Although the vision provided by a prosthesis is extremely diminished in quality, it is consistent enough for misalignments to have readily observable effects on accuracy. One might therefore be surprised when prosthesis users fail to notice or adapt to camera misalignments, and accuracy remains poor without any indication of improvement.

Held and Hein[136] pointed out that reafference was required for adaptation to prism glasses. This immediately provides one reason why prosthesis users should not be able to adapt to camera misalignments. If camera misalignments are considered equivalent to the visual shifts caused by prism glasses, users would only be able to properly adjust their motor commands after viewing their actions through the prosthesis. Current prosthetic vision, however, is not informative enough for users to recognize their body parts or their actions, and Luo's[234] experiment with an LED on the user's finger suggested that generating reafference is not a simple task.

Beyond not being able to recognize their actions, prosthesis users have difficulty recognizing the source of any particular stimulation. Elicited phosphenes often lack fine details, such that individual or small groups of electrodes produce images of blobs or arcs rather than clean dots[170]. Localization is also not precise, as seen through the SDs of target localization responses shown in Table 2. Between not being able to see a clear picture, with no orienting cues to provide structure, and not localizing percepts consistently, it is little wonder that users do not readily detect camera misalignments. With localization error SDs of 5°–9° after substantial

practice, users probably have some learned expectation for not localizing objects consistently. Even when something is seen unusually far from its true location, users cannot rule out the possibility that they are actually looking at something other than the known target. Despite being faced with consistent failure in localizing known objects, users may not recognize the nature of the misalignment, and simply assume that they are not seeing anything meaningful.

While current prosthesis technology does not permit users to intuitively compensate for camera misalignments, some modifications might prove helpful toward that end. Not only was integrated eye tracking useful for correlating eye movements and error changes, localization with eye-head scanning also drastically improved response precision[235]. Response error distributions were concentrated into widths closer to those of the 7° targets. If users can localize stimuli more precisely, they might gain more confidence regarding the locations of percepts or even orientation cues from the environment. Increased confidence could allow them to notice camera misalignments more readily. Although eye-head scanning would not directly improve reafference, the ability to detect misalignment would place more cognitive tools for adaptation at the user's disposal.

Attaching an LED to a finger or wearing a lit glove likely will not provide users with reafference, partly because signals from the marker can be confused with other light sources. Although still subject to the messy spatial and temporal dynamics of retinal stimulation, driving prosthesis activity based on infrared signals instead of visible light could reduce confusing environment noise. Infrared-based stimulation has been shown to significantly help in the localization of people and hot objects, when compared to standard visible-light use[236]. Users would still not likely see any well-formed image when infrared signals are processed, but the chances of them being able to relate stimulation to their own hand movement should be greater than with

normal imagery. Stimulation might not correlate sufficiently with movements to qualify as reafference useful for cerebellar learning, but it could be another resource for helping users to perceive and understand camera misalignment.

Special rehabilitation for identifying and responding to camera misalignment was not explored in any of these experiments. Subjects' responses to in-lab auditory feedback showed that they could improve accuracy when errors where highlighted, although that learning may not have generalized beyond specific test conditions. It is possible that users could instead be explicitly trained to detect discrepancies between elicited visual percepts and tactile signals. Certain exercises might help to mentally calibrate the mapping of percepts to egocentric world locations. Eye-tracking camera control and infrared imagery could be used to facilitate training and exercises to improve the chances of effective adaptation. Such active adaptation to misalignments, however, could require considerable time and effort from users. If eye tracking, or any other method, can automatically adjust CAPs for misalignments, users would probably be more satisfied with an automated solution.

The lack of passive adaptation to camera misalignment implies that many retinal prostheses and any intracranial visual prostheses will require either regular camera alignment or some rehabilitation regimen to optimize and maintain accuracy. Future devices, however, might bestow greater qualities of vision upon recipients. Prostheses might eventually be able to create meaningful, structured visual fields with adequate representation of the body for reafference. Camera misalignments would then be closer to vision displaced by prism glasses, and users might be able to quickly and intuitively adapt to misalignments. Until such progress is achieved though, proper camera alignment will remain important for prosthetic visual function.

8.2 OCAP fluctuations

After tracking OCAPs over several years, it became apparent that pointing biases did not have long-term constant values. Run OCAP estimates varied within subjects over ranges as wide as 37°, and 30-day averages still occupied ranges 12°–31° wide. Most 30-day 95% confidence regions, however, were smaller than 10° wide. OCAP changes were not entirely random, given the progressions plotted in Figure 22–Figure 25. Something was causing OCAPs to gradually move across parts of the camera's FOV. Although pointing biases changed depending on what hand was used, the same slow patterns appeared in separate right- and left-hand responses.

Changes in eye orientations was a likely candidate for causing OCAP variations. When the implanted eye is deviated from the camera's orientation, localization errors do correlate with the direction of eye deviation. If the average eye orientation changed over months or years, it would make sense that pointing biases and OCAPs would change accordingly. Eye tracking conducted in alternation and simultaneously with target localization confirmed that eye orientation behaviors matched those of OCAPs, and adjusting CAPs with eye orientations slowed changes in localization errors.

Both OCAPs and eye orientations demonstrated interesting reset patterns between sessions and between runs. During trial runs, OCAPs would progress several degrees along a trend. That trend would be partially reversed before the start of the next trial run. The same pattern was found among OCAPs within and between sessions. Taken by itself, such trends and resets in OCAPs could be interpreted to come from fluctuating motor biases or perhaps trends in retinal adaptation to stimulation. The independent presence of reset behavior in eye orientations,

particularly when no prosthetic stimulation was provided, argues strongly for eye movements generating the observed inter-run and interday OCAP behavior.

If eye movements might drive OCAP changes at the level of intra-run and intraday trends, it is not unreasonable to imagine that they also drive OCAP changes through months and years. PL eye orientations did correlate with OCAPs during the 11 months when eye tracking was performed in alternation with target localization. The differences between setups and methods of data collection, however, made relationships weak and indirect. Simultaneous eye tracking and target localization, as described in Section 7.2, would be required through many repeated sessions to properly demonstrate the link between eye orientations and long-term OCAP changes. If the PL-measured eye orientations had been used to adjust CAPs, residual OCAPs eccentricities would have been about 59% smaller than original OCAPs, and rates of OCAP change would have been reduced by about 94%.

The best support for eye movements being tied to localization errors, at least on the timescale of minutes, comes from the acute testing done with prosthesis-integrated eye tracking. Trends within trial runs, even with only 17 data points, showed significant similarities in direction. Errors continued to shift with eye-head scanning, but rates of change were 24%–40% smaller than with head-only scanning. Future testing might benefit from using smaller targets for localization testing, as well as recording the times in which the prosthesis was providing stimulation. With smaller targets, greater reduction in error shifts may observed with eye-head scanning. Knowing the orientation of the eye at the time of stimulation, along with the position of the target, would help investigate the effect eye orientation on localization errors with greater detail.

Ultimately, these studies demonstrated that localization errors and eye orientations are related in some way. In order to properly determine whether eye movements cause shifts in localization errors and OCAPs, however, OCAPs would need to be evaluated over extended periods when no eye movements are performed or intended. As no easy method exists for removing the effects of corollary discharge on percept localization over long periods, detailed correlation analyses may be the only practical way to probe the contribution of eye orientation to localization errors. Even if one cannot demonstrate causality, it may be sufficient to know that eye movements contain information on how percepts will be located. Such information can be used to rapidly adjust camera settings, regardless of whether eye movements are actually the proximal cause of OCAP and error shifts.

8.3 Prosthesis-integrated eye tracking

Eye tracking merged with current Argus II equipment proved very helpful for demonstrating a link between eye movements and localization error changes. Although the prototype setup was stationary and restricted to lab use, the proof of concept can be used to develop a mobile configuration to replace standard Argus II glasses and processors. If users could always be equipped with eye-tracking glasses, the benefits of eye-head scanning could become standard for home use. Localization errors could be tracked to determine whether errors continue to shift after compensation for eye movements, and eye-tracking data would provide much more information on eye movement behavior outside of lab settings.

Based on the acute experiment described in Section 7.2, eye-head scanning both improves pointing precision and slows the rate of error change within trial runs. In order for a system to accommodate all the pointing biases observed in Chapter 6, a larger FOV would be needed from

the scene camera. The camera used with the prototype eye-tracking setup had a FOV of 73° × 55°. This was larger than the Argus II camera FOV, 49° × 38°, but still not large enough to handle the most extreme calculated OCAPs. The eye tracker's scene camera could accommodate CAPs ranging ±28° × ±22°. That horizontal range was barely enough to satisfy estimated horizontal OCAPs, but vertical OCAPs for S1–S3 required eccentricities as high as 25°–40°. Baseline CAPs would need to be set at these eccentric positions, and the camera FOV would need to span even greater ranges to allow for transient eye movements. Given that using a wider FOV lens would increase image distortion, future designs may need multiple scene cameras or allow for physical rotation of the camera along multiple axes.

Although eye-head scanning and subtracting GPVs from OCAPs reduced necessary CAP adjustments and slowed pointing-bias changes, these reductions were not consistently complete. The lack of complete error compensation when comparing tower-mount GPVs to OCAPs could be explained by the differences between setups. The persistence of error changes with eye-head scanning, however, implies other factors need consideration. The effects of target size and viewing bias mentioned in Section **Error! Reference source not found.** could be investigated by repeating simultaneous eye tracking and target localization with different target sizes. Insufficient CAP adjustments related to boundaries of the scene camera may have also contributed to localization errors. Removing the effects of the camera boundaries would be less simple, requiring hardware changes that expand the scene camera's FOV.

Beyond visual contributions to pointing bias, motor commands may also be out of proportion with respect to eye orientations. Without reafference to aid cerebellar calibration of motor commands, pointing motions might extend to greater eccentricities than the user intends. Equivalently, eyes may not move as far as intended or implied in efference copies. Careful

examination of eye orientations at times of stimulation, target locations, and response positions would be necessary to determine if there is any consistent tendency of a subject to point beyond where one would expect based on eye orientations. If subjects almost always make errors in the direction of eye deviations, but with eccentricities greater than or equal to those of the eye, some personalized scaling factor may be required for CAP adjustments. Further analysis can then determine whether localization errors continue to shift or eye-orientation CAP adjustment adequately stabilizes pointing biases.

Future improved visual prostheses will hopefully provide much more detailed and precise visual information to their users. With structured visual fields and reafference, users could likely adapt to constant components of camera misalignment, and that capacity for adaptation may anchor pointing biases and eliminate the time-dependent trends recently observed. With that improved perception, however, adjusting CAPs to account transient eye movements will still become even more important. Currently, misalignments generated by individual eye movements create localization errors, but users are largely unaware of the problem. Once visual quality improves, users will become far more sensitive to misalignment. If the visual scene does not change with eye movements, users will experience disorienting illusions of the entire scene moving with the eyes. Eye tracking integrated into visual prostheses will thus be important for current technology as well as any developed with improved neural tissue interfaces.

References

1. Drasdo N, Fowler CW. Non-linear projection of the retinal image in a wide-angle schematic eye. *Br J Ophthalmol* 1974;58:709-714.

2. Hildebrand C, Remahl S, Waxman SG. Axo-glial relations in the retina-optic nerve junction of the adult rat: electron-microscopic observations. *J Neurocytol* 1985;14:597-617.

3. Straatsma BR, Foos RY, Heckenlively JR, Taylor GN. Myelinated Retinal Nerve Fibers. *Am J Ophthalmol* 1981;91:25-38.

4. Dacey DM. Primate retina: cell types, circuits and color opponency. *Prog Retin Eye Res* 1999;18:737-763.

5. Dacey D. Origins of Perception: Retinal Ganglion Cell Diversity and the Creation of Parallel Visual Pathways. In: Gazzaniga MS (ed), *The Cognitive Neurosciences Iii*: MIT Press; 2004:281-301.

6. Bird AC, Weale RA. On the retinal vasculature of the human fovea. *Exp Eye Res* 1974;19:409-417.

7. Provis JM, Hendrickson AE. The foveal avascular region of developing human retina. *Arch Ophthalmol* 2008;126:507-511.

8. Kolb H, Dekorver L. Midget ganglion cells of the parafovea of the human retina: A Study by electron microscopy and serial section reconstructions. *The Journal of Comparative Neurology* 1991;303:617-636.

9. Provis JM, Dubis AM, Maddess T, Carroll J. Adaptation of the Central Retina for High Acuity Vision: Cones, the Fovea and the Avascular Zone. *Prog Retin Eye Res* 2013;35:63-81.

10. Yamada E. Some structural features of the fovea centralis in the human retina. *Arch Ophthalmol* 1969;82:151-159.

11. Strauss O. The Retinal Pigment Epithelium in Visual Function. *Physiol Rev* 2005;85:845-881.

12. Nickla DL, Wallman J. The Multifunctional Choroid. *Prog Retin Eye Res* 2010;29:144-168.

13. Weiter JJ, Delori FC, Wing GL, Fitch KA. Retinal pigment epithelial lipofuscin and melanin and choroidal melanin in human eyes. *Invest Ophthalmol Vis Sci* 1986;27:145-152.

14. Hu D-N, Simon JD, Sarna T. Role of Ocular Melanin in Ophthalmic Physiology and Pathology. *Photochem Photobiol* 2008;84:639-644.

15. Watson PG, Young RD. Scleral structure, organisation and disease. A review. *Exp Eye Res* 2004;78:609-623.

16. Veleri S, Lazar CH, Chang B, Sieving PA, Banin E, Swaroop A. Biology and therapy of inherited retinal degenerative disease: insights from mouse models. *Disease Models & Mechanisms* 2015;8:109-129.

17. Koyama Y, Kubo K, Komori M, Yasuda H, Mukai Y. Effect of protonation on the isomerization properties of n-butylamine Schiff base of isomeric retinal as revealed by direct HPLC analyses: Selection of isomerization pathways by retinal proteins. *Photochem Photobiol* 1991;54:433-443.

18. Nakanishi K, Balogh-Nair V, Arnaboldi M, Tsujimoto K, Honig B. An external point-charge model for bacteriorhodopsin to account for its purple color. *J Am Chem Soc* 1980;102:7945-7947.

19. Wald G, Brown PK. Human rhodopsin. *Science* 1958;127:222-249.

20. Merbs SL, Nathans J. Absorption spectra of human cone pigments. *Nature* 1992;356:433-435.

21. Bailes HJ, Lucas RJ. Human melanopsin forms a pigment maximally sensitive to blue light (λ_{max} ≈ 479 nm) supporting activation of G_q_{/11} and G_{i/o} signalling cascades. *Proceedings of the Royal Society B: Biological Sciences* 2013;280.

22. Munk C, Isberg V, Mordalski S, et al. GPCRdb: the G protein-coupled receptor database – an introduction. *Br J Pharmacol* 2016;173:2195-2207.

23. Birnbaumer L. G Proteins in Signal Transduction. *Annu Rev Pharmacol Toxicol* 1990;30:675-705.

24. Jastrzebska B, Orban T, Golczak M, Engel A, Palczewski K. Asymmetry of the rhodopsin dimer in complex with transducin. *The FASEB Journal* 2013;27:1572-1584.

25. Gunkel M, Schöneberg J, Alkhaldi W, et al. Higher-Order Architecture of Rhodopsin in Intact Photoreceptors and Its Implication for Phototransduction Kinetics. *Structure* 2015;23:628-638.

26. Terakita A. The opsins. *Genome Biol* 2005;6:213.

27. Shichida Y, Morizumi T. Mechanism of G-protein Activation by Rhodopsin. *Photochem Photobiol* 2007;83:70-75.

28. Struts AV, Salgado GFJ, Martínez-Mayorga K, Brown MF. Retinal Dynamics Underlie Inverse-Agonist to Agonist Switch in Rhodopsin Activation. *Nat Struct Mol Biol* 2011;18:392-394.

29. Schöneberg J, Heck M, Hofmann Klaus P, Noé F. Explicit Spatiotemporal Simulation of Receptor-G Protein Coupling in Rod Cell Disk Membranes. *Biophys J* 2014;107:1042-1053.

30. Gross OP, Burns ME. Control of rhodopsin's active lifetime by arrestin-1 expression in mammalian rods. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 2010;30:3450-3457.

31. Lee KA, Nawrot M, Garwin GG, Saari JC, Hurley JB. Relationships between visual cycle retinoids, rhodopsin phosphorylation and phototransduction in mouse eyes during light and dark-adaptation. *Biochemistry* 2010;49:2454-2463.

32. Kliewer A, Reinscheid RK, Schulz S. Emerging Paradigms of G Protein-Coupled Receptor Dephosphorylation. *Trends Pharmacol Sci* 2017;38:621-636.

33. Arshavsky VY, Burns ME. Current understanding of signal amplification in phototransduction. *Cellular Logistics* 2014;4:e29390.

34. Kosloff M, Alexov E, Arshavsky VY, Honig B. Electrostatic and Lipid Anchor Contributions to the Interaction of Transducin with Membranes: MECHANISTIC IMPLICATIONS FOR ACTIVATION AND TRANSLOCATION. *J Biol Chem* 2008;283:31197-31207.

35. Catty P, Pfister C, Bruckert F, Deterre P. The cGMP phosphodiesterase-transducin complex of retinal rods. Membrane binding and subunits interactions. *J Biol Chem* 1992;267:19489-19493.

36. Skiba NP, Hopp JA, Arshavsky VY. The Effector Enzyme Regulates the Duration of G Protein Signaling in Vertebrate Photoreceptors by Increasing the Affinity between Transducin and RGS Protein. *J Biol Chem* 2000;275:32716-32720.

37. Chen C-K, Wieland T, Simon MI. RGS-r, a retinal specific RGS protein, binds an intermediate conformation of transducin and enhances recycling. *Proc Natl Acad Sci U S A* 1996;93:12885-12889.

38. Clack JW, Springmeyer ML, Clark CR, Witzmann FA. Transducin subunit stoichiometry and cellular distribution in rod outer segments. *Cell Biol Int* 2006;30:829-835.

39. Deterre P, Bigay J, Forquet F, Robert M, Chabre M. cGMP phosphodiesterase of retinal rods is regulated by two inhibitory subunits. *Proc Natl Acad Sci U S A* 1988;85:2424-2428.

40. Bender AT, Beavo JA. Cyclic Nucleotide Phosphodiesterases: Molecular Regulation to Clinical Use. *Pharmacol Rev* 2006;58:488-520.

41. Kolandaivelu S, Chang B, Ramamurthy V. Rod Phosphodiesterase-6 (PDE6) Catalytic Subunits Restore Cone Function in a Mouse Model Lacking Cone PDE6 Catalytic Subunit. *The Journal of Biological Chemistry* 2011;286:33252-33259.

42. Norton AW, D'Amours MR, Grazio HJ, Hebert TL, Cote RH. Mechanism of Transducin Activation of Frog Rod Photoreceptor Phosphodiesterase: Allosteric interactions between the inhibitory γ subunit and the noncatalytic cGMP-binding sites. *J Biol Chem* 2000;275:38611-38619.

43. Krispel CM, Chen D, Melling N, et al. RGS Expression Rate-Limits Recovery of Rod Photoresponses. *Neuron* 2006;51:409-416.

44. Cukkemane A, Seifert R, Kaupp UB. Cooperative and uncooperative cyclic-nucleotide-gated ion channels. *Trends Biochem Sci* 2011;36:55-64.

45. Shuart NG, Haitin Y, Camp SS, Black KD, Zagotta WN. Molecular mechanism for 3:1 subunit stoichiometry of rod cyclic nucleotide-gated ion channels. *Nature Communications* 2011;2:457.

46. Gross OP, Pugh Jr EN, Burns ME. cGMP in mouse rods: the spatiotemporal dynamics underlying single photon responses. *Front Mol Neurosci* 2015;8:6.

47. Sterling P, Laughlin S. *Principles of Neural Design*: MIT Press; 2015.

48. Nakatani K, Yau KW. Guanosine 3',5'-cyclic monophosphate-activated conductance studied in a truncated rod outer segment of the toad. *The Journal of Physiology* 1988;395:731-753.

49. Baylor DA, Lamb TD, Yau KW. Responses of retinal rods to single photons. *The Journal of Physiology* 1979;288:613-634.

50. Field GD, Rieke F. Mechanisms Regulating Variability of the Single Photon Responses of Mammalian Rod Photoreceptors. *Neuron* 2002;35:733-747.

51. Ohyama T, Picones A, Korenbrot JI. Voltage-dependence of Ion Permeation in Cyclic GMP–gated Ion Channels Is Optimized for Cell Function in Rod and Cone Photoreceptors. *The Journal of General Physiology* 2002;119:341-354.

52. Ohyama T, Hackos DH, Frings S, Hagen V, Kaupp UB, Korenbrot JI. Fraction of the Dark Current Carried by Ca(2+) through Cgmp-Gated Ion Channels of Intact Rod and Cone Photoreceptors. *The Journal of General Physiology* 2000;116:735-754.

53. Derebe MG, Zeng W, Li Y, Alam A, Jiang Y. Structural studies of ion permeation and Ca2+ blockage of a bacterial channel mimicking the cyclic nucleotide-gated channel pore. *Proceedings of the National Academy of Sciences* 2011;108:592-597.

54. Schneeweis DM, Schnapf JL. Photovoltage of rods and cones in the macaque retina. *Science-AAAS-Weekly Paper Edition* 1995;268:1053-1055.

55. Koschak A, Pinggera A, Schicker K, Striessnig J. Role of L-Type Ca2+ Channels in Sensory Cells. In: Weiss N, Koschak A (eds), *Pathologies of Calcium Channels.* Berlin, Heidelberg: Springer Berlin Heidelberg; 2014:47-75.

56. Krizaj D, Copenhagen DR. Calcium regulation in photoreceptors. *Frontiers in bioscience : a journal and virtual library* 2002;7:d2023-d2044.

57. Schein S, Ahmad KM. A Clockwork Hypothesis: Synaptic Release by Rod Photoreceptors Must Be Regular. *Biophys J* 2005;89:3931-3949.

58. Chen M, Krizaj D, Thoreson W. Intracellular calcium stores drive slow non-ribbon vesicle release from rod photoreceptors. *Front Cell Neurosci* 2014;8.

59. Babai N, Morgans CW, Thoreson WB. Calcium-induced calcium release contributes to synaptic release from mouse rod photoreceptors. *Neuroscience* 2010;165:1447.

60. Mata NL, Radu RA, Clemmons RS, Travis GH. Isomerization and Oxidation of Vitamin A in Cone-Dominant Retinas: A Novel Pathway for Visual-Pigment Regeneration in Daylight. *Neuron* 2002;36:69-80.

61. Kiser PD, Golczak M, Palczewski K. Chemistry of the Retinoid (Visual) Cycle. *Chem Rev* 2014;114:194-232.

62. Malechka VV, Moiseyev G, Takahashi Y, Shin Y, Ma J-x. Impaired Rhodopsin Generation in the Rat Model of Diabetic Retinopathy. *The American Journal of Pathology* 2017;187:2222-2231.

63. Eroglu A, Harrison EH. Carotenoid metabolism in mammals, including man: formation, occurrence, and function of apocarotenoids: Thematic Review Series: Fat-Soluble Vitamins: Vitamin A. *J Lipid Res* 2013;54:1719-1730.

64. Hyung S-J, Deroo S, Robinson CV. Retinol and Retinol-Binding Protein Stabilize Transthyretin via Formation of Retinol Transport Complex. *ACS Chem Biol* 2010;5:1137-1146.

65. O'Byrne SM, Blaner WS. Retinol and retinyl esters: biochemistry and physiology: Thematic Review Series: Fat-Soluble Vitamins: Vitamin A. *J Lipid Res* 2013;54:1731-1743.

66. Bouillet P, Oulad-Abdelghani M, Vicaire S, et al. Efficient Cloning of cDNAs of Retinoic Acid-Responsive Genes in P19 Embryonal Carcinoma Cells and Characterization of a Novel Mouse Gene, Stra1 (Mouse LERK-2/Eplg2). *Dev Biol* 1995;170:420-433.

67. Kawaguchi R, Yu J, Honda J, et al. A Membrane Receptor for Retinol Binding Protein Mediates Cellular Uptake of Vitamin A. *Science* 2007;315:820-825.

68. Kiser PD, Palczewski K. Retinoids and Retinal Diseases. *Annual Review of Vision Science* 2016;2:197-234.

69. Dove AW. New structure found in plain sight. *The Journal of Cell Biology* 2004;164:332-332.

70. Chelstowska S, Widjaja-Adhi M, Silvaroli J, Golczak M. Molecular Basis for Vitamin A Uptake and Storage in Vertebrates. *Nutrients* 2016;8:676.

71. Palczewska G, Maeda T, Imanishi Y, et al. Noninvasive multi–photon fluorescence microscopy resolves retinol and retinal–condensation products in mouse eyes. *Nat Med* 2010;16:1444-1449.

72. Lhor M, Salesse C. Retinol dehydrogenases: Membrane-bound enzymes for the visual function. *Biochem Cell Biol* 2014;92:510-523.

73. Sahu B, Maeda A. Retinol Dehydrogenases Regulate Vitamin A Metabolism for Visual Function. *Nutrients* 2016;8.

74. Golczak M, Kiser PD, Lodowski DT, Maeda A, Palczewski K. Importance of Membrane Structural Integrity for RPE65 Retinoid Isomerization Activity. *The Journal of Biological Chemistry* 2010;285:9667-9682.

75. Saari JC, Nawrot M, Stenkamp RE, Teller DC, Garwin GG. Release of 11-cis-retinal from cellular retinaldehyde-binding protein by acidic lipids. *Mol Vis* 2009;15:844-854.

76. Gonzalez-Fernandez F, Betts-Obregon B, Yust B, et al. Interphotoreceptor Retinoid-Binding Protein Protects Retinoids from Photodegradation. *Photochem Photobiol* 2015;91:371-378.

77. Palczewski K. Retinoids for Treatment of Retinal Diseases. *Trends Pharmacol Sci* 2010;31:284-295.

78. Quazi F, Lenevich S, Molday RS. ABCA4 is an N-retinylidene-phosphatidylethanolamine and phosphatidylethanolamine importer. *Nature Communications* 2012;3:925.

79. Sparrow JR, Wu Y, Kim CY, Zhou J. Phospholipid meets all-trans-retinal: the making of RPE bisretinoids. *J Lipid Res* 2010;51:247-261.

80. Biswas-Fiss EE, Kurpad DS, Joshi K, Biswas SB. Interaction of Extracellular Domain 2 of the Human Retina-specific ATP-binding Cassette Transporter (ABCA4) with All-trans-retinal. *The Journal of Biological Chemistry* 2010;285:19372-19383.

81. Wu Q, Chen C, Koutalos Y. All-Trans Retinol in Rod Photoreceptor Outer Segments Moves Unrestrictedly by Passive Diffusion. *Biophys J* 2006;91:4678-4689.

82. Chen C, Adler L, Goletz P, Gonzalez-Fernandez F, Thompson DA, Koutalos Y. Interphotoreceptor retinoid binding protein removes all-trans retinol and retinal from rod outer segments preventing lipofuscin precursor formation. *J Biol Chem* 2017.

83. Haeseleer F, Huang J, Lebioda L, Saari JC, Palczewski K. Molecular Characterization of a Novel Short-chain Dehydrogenase/Reductase That Reduces All-trans-retinal. *J Biol Chem* 1998;273:21790-21799.

84. Kaylor JJ, Yuan Q, Cook J, et al. Identification of DES1 as a Vitamin A Isomerase in Müller Glial Cells of the Retina. *Nat Chem Biol* 2013;9:30-36.

85. Kaylor JJ, Cook JD, Makshanoff J, Bischoff N, Yong J, Travis GH. Identification of the 11-cis-specific retinyl-ester synthase in retinal Müller cells as multifunctional O-acyltransferase (MFAT). *Proc Natl Acad Sci U S A* 2014;111:7302-7307.

86. Xue Y, Shen SQ, Jui J, et al. CRALBP supports the mammalian retinal visual cycle and cone vision. *The Journal of Clinical Investigation* 2015;125:727-738.

87. Xue Y, Sato S, Razafsky D, et al. The role of retinol dehydrogenase 10 in the cone visual cycle. *Sci Rep* 2017;7:2390.

88. Wu SM. Synaptic Organization of the Vertebrate Retina: General Principles and Species-Specific Variations: The Friedenwald Lecture. *Invest Ophthalmol Vis Sci* 2010;51:1264-1274.

89. Sinha R, Hoon M, Baudin J, Okawa H, Wong ROL, Rieke F. Cellular and Circuit Mechanisms Shaping the Perceptual Properties of the Primate Fovea. *Cell* 2017;168:413-426.e412.

90. Asteriti S, Gargini C, Cangiano L. Mouse rods signal through gap junctions with cones. *eLife* 2014;3:e01386.

91. Jiang Z, Yang J, Purpura LA, Liu Y, Ripps H, Shen W. Glycinergic feedback enhances synaptic gain in the distal retina. *The Journal of Physiology* 2014;592:1479-1492.

92. Dacheux R, Raviola E. Physiology of HI horizontal cells in the primate retina. *Proceedings of the Royal Society of London B: Biological Sciences* 1990;239:213-230.

93. Kolb H, Fernandez E, Schouten J, Ahnelt P, Linberg KA, Fisher SK. Are there three types of horizontal cell in the human retina? *J Comp Neurol* 1994;343:370-386.

94. Boije H, Shirazi Fard S, Edqvist P-H, Hallböök F. Horizontal Cells, the Odd Ones Out in the Retina, Give Insights into Development and Disease. *Front Neuroanat* 2016;10:77.

95. Chapot CA, Euler T, Schubert T. How do horizontal cells 'talk' to cone photoreceptors? Different levels of complexity at the cone–horizontal cell synapse. *The Journal of Physiology* 2017;595:5495-5506.

96. Kamermans M, Fahrenfort I. Ephaptic interactions within a chemical synapse: hemichannel-mediated ephaptic inhibition in the retina. *Curr Opin Neurobiol* 2004;14:531-541.

97. DeVries SH. Exocytosed Protons Feedback to Suppress the Ca2+ Current in Mammalian Cone Photoreceptors. *Neuron* 2001;32:1107-1117.

98. Jouhou H, Yamamoto K, Homma A, Hara M, Kaneko A, Yamada M. Depolarization of isolated horizontal cells of fish acidifies their immediate surrounding by activating V-ATPase. *The Journal of Physiology* 2007;585:401-412.

99. Tachibana M, Kaneko A. gamma-Aminobutyric acid acts at axon terminals of turtle photoreceptors: difference in sensitivity among cell types. *Proc Natl Acad Sci U S A* 1984;81:7961-7964.

100. Verweij J, Hornstein EP, Schnapf JL. Surround Antagonism in Macaque Cone Photoreceptors. *The Journal of Neuroscience* 2003;23:10249-10257.

101. Liu X, Hirano AA, Sun X, Brecha NC, Barnes S. Calcium channels in rat horizontal cells regulate feedback inhibition of photoreceptors through an unconventional GABA- and pH-sensitive mechanism. *The Journal of Physiology* 2013;591:3309-3324.

102. Endeman D, Fahrenfort I, Sjoerdsma T, Steijaert M, ten Eikelder H, Kamermans M. Chloride currents in cones modify feedback from horizontal cells to cones in goldfish retina. *The Journal of Physiology* 2012;590:5581-5595.

103. Kemmler R, Schultz K, Dedek K, Euler T, Schubert T. Differential Regulation of Cone Calcium Signals by Different Horizontal Cell Feedback Mechanisms in the Mouse Retina. *The Journal of Neuroscience* 2014;34:11826-11843.

104. Kuffler SW. Neurons in the retina: organization, inhibition and excitation problems. *Cold Spring Harb Symp Quant Biol* 1952;17:281-292.

105. Haverkamp S, Haeseleer F, Hendrickson A. A comparison of immunocytochemical markers to identify bipolar cell types in human and monkey retina. *Vis Neurosci* 2004;20:589-600.

106. Field GD, Gauthier JL, Sher A, et al. Functional connectivity in the retina at the resolution of photoreceptors. *Nature* 2010;467:673.

107. Crook JD, Manookin MB, Packer OS, Dacey DM. Horizontal Cell Feedback without Cone Type-Selective Inhibition Mediates “Red–Green” Color Opponency in Midget Ganglion Cells of the Primate Retina. *The Journal of Neuroscience* 2011;31:1762-1772.

108. Euler T, Haverkamp S, Schubert T, Baden T. Retinal bipolar cells: elementary building blocks of vision. *Nat Rev Neurosci* 2014;15:507-519.

109. Dacey DM, Crook JD, Packer OS. Distinct synaptic mechanisms create parallel S-ON and S-OFF color opponent pathways in the primate retina. *Vis Neurosci* 2014;31:139-151.

110. Varela C, Blanco R, De la Villa P. Depolarizing effect of GABA in rod bipolar cells of the mouse retina. *Vision Res* 2005;45:2659-2667.

111. Puthussery T, Venkataramani S, Gayet-Primo J, Smith RG, Taylor WR. Na(V)1.1 Channels in Axon Initial Segments of Bipolar Cells Augment Input to Magnocellular Visual Pathways in the Primate Retina. *The Journal of Neuroscience* 2013;33:16045-16059.

112. Field GD, Sher A, Gauthier JL, et al. Spatial Properties and Functional Organization of Small Bistratified Ganglion Cells in Primate Retina. *The Journal of Neuroscience* 2007;27:13261-13272.

113. Graham DM, Wong KY. Melanopsin-expressing, intrinsically photosensitive retinal ganglion cells (ipRGCs). 2016.

114. Spitschan M, Jain S, Brainard DH, Aguirre GK. Opponent melanopsin and S-cone signals in the human pupillary light response. *Proceedings of the National Academy of Sciences* 2014;111:15568-15572.

115. Ghodrati M, Khaligh-Razavi S-M, Lehky SR. Towards building a more complex view of the lateral geniculate nucleus: Recent advances in understanding its role. *Prog Neurobiol* 2017;156:214-255.

116. Schneider KA, Richter MC, Kastner S. Retinotopic Organization and Functional Subdivisions of the Human Lateral Geniculate Nucleus: A High-Resolution Functional Magnetic Resonance Imaging Study. *The Journal of Neuroscience* 2004;24:8975-8985.

117. Benjamin CFA, Singh JM, Prabhu SP, Warfield SK. Optimization of tractography of the optic radiations. *Hum Brain Mapp* 2014;35:683-697.

118. Mahaney KB, Abdulrauf S. *Anatomic relationship of the optic radiations to the atrium of the lateral ventricle: Description of a novel entry point to the trigone*; 2008:195-202; discussion 202.

119. Callaway EM. Local circuits in primary visual cortex of the macaque monkey. *Annu Rev Neurosci* 1998;21:47-74.

120. Brown JF. The visual perception of velocity. *Psychol Forsch* 1931;14:199-232.

121. Proske U, Gandevia SC. The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force. *Physiol Rev* 2012;92:1651-1697.

122. Blumer R, Maurer-Gesek B, Gesslbauer B, et al. Palisade Endings Are a Constant Feature in the Extraocular Muscles of Frontal-Eyed, But Not Lateral-Eyed, Animals. *Invest Ophthalmol Vis Sci* 2016;57:320-331.

123. Crapse TB, Sommer MA. Corollary discharge across the animal kingdom. *Nature Reviews Neuroscience* 2008;9:587.

124. Sun LD, Goldberg ME. Corollary Discharge and Oculomotor Proprioception: Cortical Mechanisms for Spatially Accurate Vision. *Annual Review of Vision Science* 2016;2:61-84.

125. Sommer MA, Wurtz RH. Influence of the thalamus on spatial visual processing in frontal cortex. *Nature* 2006;444:374.

126. Raphan T, Cohen B. The vestibulo-ocular reflex in three dimensions. *Exp Brain Res* 2002;145:1-27.

127. Crowell JA, Banks MS, Shenoy KV, Andersen RA. Visual self-motion perception during head turns. *Nat Neurosci* 1998;1:732-737.

128. Bridgeman B, Stark L. Ocular proprioception and efference copy in registering visual direction. *Vision Res* 1991;31:1903-1913.

129. Morgan CL. Constancy of egocentric visual direction. *Percept Psychophys* 1978;23:61-68.

130. Lewald J, Ehrenstein WH. Visual and proprioceptive shifts in perceived egocentric direction induced by eye-position. *Vision Res* 2000;40:539-547.

131. Stratton GM. Some preliminary experiments on vision without inversion of the retinal image. *Psychol Rev* 1896;3:611-617.

132. Stratton GM. Vision without inversion of the retinal image. *Psychol Rev* 1897;4:341-360.

133. Von Helmholtz H. *Handbuch der physiologischen Optik*: Voss; 1867.

134. Wooster M. Certain factors in the development of a new spatial co-ordination. *Psychol Monogr* 1923;32:i-96.

135. Gibson JJ. Adaptation, after-effect and contrast in the perception of curved lines. *J Exp Psychol* 1933;16:1-31.

136. Held R, Hein AV. Adaptation of disarranged hand-eye coordination contingent upon re-afferent stimulation. *Percept Mot Skills* 1958;8:87-90.

137. Gilligan T, Cristino F, Rafal R, Bultitude J. Still seeing straight: No role for ocular proprioception in prism adaptation? *Journal of Vision* 2016;16:535-535.

138. McLaughlin SC, Webster RG. Changes in straight-ahead eye position during adaptation to wedge prisms. *Percept Psychophys* 1967;2:37-44.

139. Sisak S, Banin E, Blumenthal EZ. A two-compartment model of the human retina. *Med Hypotheses* 2004;62:808-816.

140. Feltgen N, Walter P. Rhegmatogenous Retinal Detachment—an Ophthalmologic Emergency. *Deutsches Ärzteblatt International* 2014;111:12-22.

141. Abouzeid H, Wolfensberger TJ. Macular recovery after retinal detachment. *Acta Ophthalmol Scand* 2006;84:597-605.

142. van Bussel EM, van der Valk R, Bijlsma WR, La Heij EC. IMPACT OF DURATION OF MACULA-OFF RETINAL DETACHMENT ON VISUAL OUTCOME: A Systematic Review and Meta-analysis of Literature. *Retina* 2014;34:1917-1925.

143. Chen K-H, Chen L-R. Bilateral retinal detachment with subsequent blindness in a pregnant woman with severe pre-eclampsia. *Taiwanese Journal of Obstetrics and Gynecology* 2013;52:142-144.

144. Jalali S. Retinal Detachment. *Community Eye Health* 2003;16:25-26.

145. Fritsche LG, Fariss RN, Stambolian D, Abecasis GR, Curcio CA, Swaroop A. Age-Related Macular Degeneration: Genetics and Biology Coming Together. *Annual Review of Genomics and Human Genetics* 2014;15:151-171.

146. Bui TV, Han Y, Radu RA, Travis GH, Mata NL. Characterization of Native Retinal Fluorophores Involved in Biosynthesis of A2E and Lipofuscin-associated Retinopathies. *J Biol Chem* 2006;281:18112-18119.

147. Rimayanti U, Kiuchi Y, Yamane K, et al. Inner retinal layer comparisons of eyes with exudative age-related macular degeneration and eyes with age-related macular degeneration and glaucoma. *Graefe's Archive for Clinical and Experimental Ophthalmology* 2014;252:563-570.

148. Fuerst NM, Serrano L, Han G, et al. Detailed functional and structural phenotype of Bietti crystalline dystrophy associated with mutations in CYP4V2 complicated by choroidal neovascularization. *Ophthalmic Genet* 2016;37:445-452.

149. Halford S, Liew G, Mackay DS, et al. Detailed Phenotypic and Genotypic Characterization of Bietti Crystalline Dystrophy. *Ophthalmology* 2014;121:1174-1184.

150. Shintani K, Shechtman DL, Gurwood AS. Review and update: Current treatment trends for patients with retinitis pigmentosa. *Optometry - Journal of the American Optometric Association* 2009;80:384-401.

151. Narayan DS, Wood JPM, Chidlow G, Casson RJ. A review of the mechanisms of cone degeneration in retinitis pigmentosa. *Acta Ophthalmol (Copenh)* 2016;94:748-754.

152. Hamel C. Retinitis pigmentosa. *Orphanet J Rare Dis* 2006;1:40.

153. Jones BW, Pfeiffer RL, Ferrell WD, Watt CB, Marmor M, Marc RE. Retinal remodeling in human retinitis pigmentosa. *Exp Eye Res* 2016;150:149-165.

154. Crawford Downs J, Roberts MD, Sigal IA. Glaucomatous cupping of the lamina cribrosa: A review of the evidence for active progressive remodeling as a mechanism. *Exp Eye Res* 2011;93:133-140.

155. Flammer J, Mozaffarieh M. What Is the Present Pathogenetic Concept of Glaucomatous Optic Neuropathy? *Surv Ophthalmol* 2007;52:S162-S173.

156. Fan BJ, Wiggs JL. Glaucoma: genes, phenotypes, and new directions for therapy. *The Journal of Clinical Investigation* 2010;120:3064-3072.

157. Verma R, Gupta M, Chaudhari TS. Neurogenic vision loss: Causes and outcome. An experience from a tertiary center in Northern India. *J Neurosci Rural Pract* 2014;5:340-348.

158. Zhang X, Kedar S, Lynn MJ, Newman NJ, Biousse V. Homonymous hemianopias. *Clinical–anatomic correlations in 904 cases* 2006;66:906-910.

159. Foroozan R, Buono LM, Savino PJ, Sergott RC. Acute demyelinating optic neuritis. *Curr Opin Ophthalmol* 2002;13:375-380.

160. Evangelou N, Konz D, Esiri MM, Smith S, Palace J, Matthews PM. Size-selective neuronal changes in the anterior optic pathways suggest a differential susceptibility to injury in multiple sclerosis. *Brain* 2001;124:1813-1820.

161. Fisher JB, Jacobs DA, Markowitz CE, et al. Relation of Visual Function to Retinal Nerve Fiber Layer Thickness in Multiple Sclerosis. *Ophthalmology* 2006;113:324-332.

162. Prieto R, Pascual JM, Barrios L. Optic Chiasm Distortions Caused by Craniopharyngiomas: Clinical and Magnetic Resonance Imaging Correlation and Influence on Visual Outcome. *World Neurosurg* 2015;83:500-529.

163. Fraser JA, Newman NJ, Biousse V. Chapter 8 - Disorders of the optic tract, radiation, and occipital lobe. In: Kennard C, Leigh RJ (eds), *Handb Clin Neurol*: Elsevier; 2011:205-221.

164. Moseman CP, Shelton S. Permanent Blindness as a Complication of Pregnancy-Induced Hypertension. *Obstet Gynecol* 2002;100:943-945.

165. Mathew T, D'Souza D, Nadig R, Sarma G. Bilateral lateral geniculate body hemorrhagic infarction: A rare cause of acute bilateral painless vision loss in female patients. *Neurol India* 2016;64:160-162.

166. Silva RR, Penna ER, Fusão EF, et al. An Unusual Cause of Blindness: Infarction in the Bilateral Lateral Geniculate Bodies. *J Stroke Cerebrovasc Dis* 2014;23:1736-1737.

167. Shang Q, Ma C-Y, Lv NAN, et al. Clinical study of cerebral palsy in 408 children with periventricular leukomalacia. *Exp Ther Med* 2015;9:1336-1344.

168. Kakaletsis N. *Bilateral Cortical Blindness due to Bilateral Occipital Infarcts without Anosognosia*; 2016.

169. Galluppi F, Pruneau D, Chavas J, et al. A stimulation platform for optogenetic and bionic vision restoration. *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*; 2017:1-4.

170. Fine I, Boynton GM. Pulse trains to percepts: the challenge of creating a perceptually intelligible world with sight recovery technologies. *Philosophical Transactions of the Royal Society B: Biological Sciences* 2015;370.

171. Li M, Yan Y, Wu K, et al. Penetrative Optic Nerve-Based Visual Prosthesis Research. In: Gabel VP (ed), *Artificial Vision: A Clinical Guide.* Cham: Springer International Publishing; 2017:165-176.

172. Veraart C, Raftopoulos C, Mortimer JT, et al. Visual sensations produced by optic nerve stimulation using an implanted self-sizing spiral cuff electrode. *Brain Res* 1998;813:181-186.

173. Veraart C, Grill WM, Mortimer JT. Selective control of muscle activation with a multipolar nerve cuff electrode. *IEEE Trans Biomed Eng* 1993;40:640-653.

174. Fitzgibbon T, Taylor SF. Retinotopy of the human retinal nerve fibre layer and optic nerve head. *The Journal of Comparative Neurology* 1996;375:238-251.

175. Miocinovic S, Somayajula S, Chitnis S, Vitek JL. History, applications, and mechanisms of deep brain stimulation. *JAMA Neurology* 2013;70:163-171.

176. Pezaris JS, Reid RC. Demonstration of artificial visual percepts generated through thalamic microstimulation. *Proceedings of the National Academy of Sciences* 2007;104:7670-7675.

177. Azzopardi P, Cowey A. The overrepresentation of the fovea and adjacent retina in the striate cortex and dorsal lateral geniculate nucleus of the macaque monkey. *Neuroscience* 1996;72:627-639.

178. Schiller PH, Malpeli JG. Functional specificity of lateral geniculate nucleus laminae of the rhesus monkey. *J Neurophysiol* 1978;41:788-797.

179. Luco C, Hoppe A, Schweitzer M, Vicuña X, Fantin A. Visual field defects in vascular lesions of the lateral geniculate body. *Journal of Neurology, Neurosurgery & Psychiatry* 1992;55:12-15.

180. Wieshmann UC, Symms MR, Clark CA, et al. Wallerian Degeneration in the Optic Radiation After Temporal Lobectomy Demonstrated In Vivo with Diffusion Tensor Imaging. *Epilepsia* 1999;40:1155-1158.

181. Stoerig P, Kleinschmidt A, Frahm J. *No visual responses in denervated V1: High-resolution functional magnetic resonance imaging of a blindsight patient*; 1998:21-25.

182. Plant GT, Kermode AG, Turano G, et al. Symptomatic retrochiasmal lesions in multiple sclerosis. *Clinical features, visual evoked potentials, and magnetic resonance imaging* 1992;42:68-68.

183. Wandell BA, Dumoulin SO, Brewer AA. Visual Field Maps in Human Cortex. *Neuron* 2007;56:366-383.

184. DeYoe EA, Carman GJ, Bandettini P, et al. Mapping striate and extrastriate visual areas in human cerebral cortex. *Proceedings of the National Academy of Sciences* 1996;93:2382-2386.

185. Kaido T, Hoshida T, Taoka T, Sakaki T. Retinotopy with coordinates of lateral occipital cortex in humans. *J Neurosurg* 2004;101:114-118.

186. Larsson J, Heeger DJ. Two retinotopic visual areas in human lateral occipital cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 2006;26:13128-13142.

187. Van Essen DC. A Population-Average, Landmark- and Surface-based (PALS) atlas of human cerebral cortex. *Neuroimage* 2005;28:635-662.

188. Abdollahi RO, Kolster H, Glasser MF, et al. Correspondences between retinotopic areas and myelin maps in human visual cortex. *Neuroimage* 2014;99:509-524.

189. Dobelle WH, Turkel J, Henderson DC, Evans JR. Mapping the representation of the visual field by electrical stimulation of human visual cortex. *Am J Ophthalmol* 1979;88:727-735.

190. Dobelle WH. Artificial vision for the blind by connecting a television camera to the visual cortex. *ASAIO J* 2000;46:3-9.

191. Lewis PM, Ayton LN, Guymer RH, et al. Advances in implantable bionic devices for blindness: a review. *ANZ J Surg* 2016;86:654-659.

192. Chow AY, Chow VY, Packo KH, Pollack JS, Peyman GA, Schuchard R. The artificial silicon retina microchip for the treatment of vision loss from retinitis pigmentosa. *Arch Ophthalmol* 2004;122:460-469.

193. Chow AY, Bittner AK, Pardue MT. The Artificial Silicon Retina in Retinitis Pigmentosa Patients (An American Ophthalmological Association Thesis). *Trans Am Ophthalmol Soc* 2010;108:120-154.

194. Ahuja AK, Dorn JD, Caspi A, et al. Blind subjects implanted with the Argus II retinal prosthesis are able to improve performance in a spatial-motor task. *Br J Ophthalmol* 2011;95:539-543.

195. Barry MP, Dagnelie G, Group AIS. Use of the Argus II retinal prosthesis to improve visual guidance of fine hand movements. *Invest Ophthalmol Vis Sci* 2012;53:5095-5101.

196. Dorn JD, Ahuja AK, Caspi A, et al. The Detection of Motion by Blind Subjects With the Epiretinal 60-Electrode (Argus II) Retinal Prosthesis. *JAMA Ophthalmol* 2013;131:183-189.

197. Ho AC, Humayun MS, Dorn JD, et al. Long-Term Results from an Epiretinal Prosthesis to Restore Sight to the Blind. *Ophthalmology* 2015;122:1547-1554.

198. da Cruz L, Coley BF, Dorn J, et al. The Argus II epiretinal prosthesis system allows letter and word reading and long-term function in patients with profound vision loss. *Br J Ophthalmol* 2013;97:632-636.

199. Zrenner E, Bartz-Schmidt KU, Benav H, et al. Subretinal electronic chips allow blind patients to read letters and combine them to words. *Proc Biol Sci* 2011;278:1489-1497.

200. Zrenner E, Bartz-Schmidt KU, Besch D, et al. The Subretinal Implant ALPHA: Implantation and Functional Results. In: Gabel VP (ed), *Artificial Vision: A Clinical Guide.* Cham: Springer International Publishing; 2017:65-83.

201. Edwards TL, Cottriall CL, Xue K, et al. Assessment of the Electronic Retinal Implant Alpha AMS in Restoring Vision to Blind Patients with End-Stage Retinitis Pigmentosa. *Ophthalmology* 2017.

202. Keserü M, Feucht M, Bornfeld N, et al. Acute electrical stimulation of the human retina with an epiretinal electrode array. *Acta Ophthalmol (Copenh)* 2012;90:e1-e8.

203. Richard G, Keserue M, Feucht M, Post N, Hornig R. Visual Perception After Long-Term Implantation of a Retinal Implant. *Invest Ophthalmol Vis Sci* 2008;49:1786-1786.

204. Hornig R, Zehnder T, Velikay-Parel M, Laube T, Feucht M, Richard G. The IMI retinal implant system. *Artificial sight*: Springer; 2007:111-128.

205. Eckmiller R, Neumann D, Baruth O. Tunable retina encoders for retina implants: why and how. *J Neural Eng* 2005;2:S91-S104.

206. Hornig R, Dapper M, Le Joliff E, et al. Pixium Vision: First Clinical Results and Innovative Developments. In: Gabel VP (ed), *Artificial Vision: A Clinical Guide.* Cham: Springer International Publishing; 2017:99-113.

207. Ivastinovic D, Langmann G, Nemetz W, Hornig R, Richard G, Velikay-Parel M. Clinical stability of a new method for fixation and explanation of epiretinal implants. *Acta Ophthalmol (Copenh)* 2010;88:e285-e286.

208. Posch C, Matolin D, Wohlgenannt R. A QVGA 143 dB Dynamic Range Frame-Free PWM Image Sensor With Lossless Pixel-Level Video Compression and Time-Domain CDS. *IEEE Journal of Solid-State Circuits* 2011;46:259-275.

209. Hemami P, Jacobs M. Pixium seeks to address short Iris II lifespan. London: Edison Investment Research; 2017.

210. Greenberg R. Updates on the Argus II Retinal Prosthesis and Orion Visual Cortical System Studies. *The Eye and the Chip.* Detroit, MI: Detroit Institute of Ophthalmology; 2017:38-39.

211. Second Sight Announces Record Number of Argus II Retinal Prosthesis Systems Implants and Completes First-in-Human Orion Cortical Implant. San Francisco, CA: Business Wire; 2018.

212. Palanker DV, Vankov A, Huie P, Fishman HA, Marmor MF, Blumenkranz MS. Can a Self-powered Retinal Prosthesis Support 100,000 Pixels in the Macula? *Invest Ophthalmol Vis Sci* 2003;44:5067-5067.

213. Lorach H, Palanker D. High Resolution Photovoltaic Subretinal Prosthesis for Restoration of Sight. In: Gabel VP (ed), *Artificial Vision: A Clinical Guide.* Cham: Springer International Publishing; 2017:115-124.

214. Fujikado T, Morimoto T, Kanda H, et al. Evaluation of phosphenes elicited by extraocular stimulation in normals and by suprachoroidal-transretinal stimulation in patients with retinitis pigmentosa. *Graefe's Archive for Clinical and Experimental Ophthalmology* 2007;245:1411-1419.

215. Fujikado T, Kamei M, Sakaguchi H, et al. Testing of Semichronically Implanted Retinal Prosthesis by Suprachoroidal-Transretinal Stimulation in Patients with Retinitis Pigmentosa. *Invest Ophthalmol Vis Sci* 2011;52:4726-4733.

216. Fujikado T. Retinal Prosthesis by Suprachoroidal-Transretinal Stimulation (STS), Japanese Approach. In: Gabel VP (ed), *Artificial Vision: A Clinical Guide.* Cham: Springer International Publishing; 2017:139-150.

217. Terasawa Y. Update of Suprachoroidal-Transretinal Stimulation (STS) Retinal Prosthesis: Efforts for the Practical Use and the Development of Next Generation Technology. *The Eye and the Chip.* Detroit, MI: Detroit Institute of Ophthalmology; 2017:66-67.

218. Ayton LN, Suaning GJ, Lovell NH, et al. Suprachoroidal Retinal Prostheses. In: Gabel VP (ed), *Artificial Vision: A Clinical Guide.* Cham: Springer International Publishing; 2017:125-138.

219. Troyk PR. The Intracortical Visual Prosthesis Project. In: Gabel VP (ed), *Artificial Vision: A Clinical Guide.* Cham: Springer International Publishing; 2017:203-214.

220. Pavan-Langston D. *Manual of ocular diagnosis and therapy*: Lippincott Williams & Wilkins; 2008.

221. Horton JC, Hoyt WF. The representation of the visual field in human striate cortex. A revision of the classic Holmes map. *Arch Ophthalmol* 1991;109:816-824.

222. Sabbah N, Authie CN, Sanda N, Mohand-Said S, Sahel JA, Safran AB. Importance of eye position on spatial localization in blind subjects wearing an Argus II retinal prosthesis. *Invest Ophthalmol Vis Sci* 2014;55:8259-8266.

223. Brindley G, Lewin W. The sensations produced by electrical stimulation of the visual cortex. *J Physiol (Lond)* 1968;196:479-493.

224. Stronks HC, Dagnelie G. The functional performance of the Argus II retinal prosthesis. *Expert Rev Med Devices* 2014;11:23-30.

225. Luo YH, Zhong JJ, da Cruz L. The use of Argus(R) II retinal prosthesis by blind subjects to achieve localisation and prehension of objects in 3-dimensional space. *Graefes Arch Clin Exp Ophthalmol* 2014.

226. Chen SC, Suaning GJ, Morley JW, Lovell NH. Rehabilitation regimes based upon psychophysical studies of prosthetic vision. *J Neural Eng* 2009;6:035009.

227. Dagnelie G, Thompson RW, Barnett D, Zhang W. Simulated prosthetic vision: Perceptual and performance measures. *OSA/VSIA* 2001;43-46.

228. Dagnelie G, Walter M, Yang L. Playing checkers: Detection and eye-hand coordination in simulated prosthetic vision. *J Modern Optics* 2006;53:1325-1342.

229. Hayes JS, Yin VT, Piyathaisere D, Weiland JD, Humayun MS, Dagnelie G. Visually guided performance of simple tasks using simulated prosthetic vision. *Artif Organs* 2003;27:1016-1028.

230. Humayun MS, Dorn JD, da Cruz L, et al. Interim results from the international trial of Second Sight's visual prosthesis. *Ophthalmology* 2012;119:779-788.

231. Kotecha A, Zhong J, Stewart D, da Cruz L. The Argus II prosthesis facilitates reaching and grasping tasks: a case series. *BMC Ophthalmol* 2014;14:71.

232. Srivastava NR, Troyk PR, Dagnelie G. Detection, eye-hand coordination and virtual mobility performance in simulated vision for a cortical visual prosthesis device. *J Neural Eng* 2009;6:035008.

233. Stingl K, Bartz-Schmidt KU, Besch D, et al. Artificial vision with wirelessly powered subretinal electronic implant alpha-IMS. *Proceedings Biological sciences / The Royal Society* 2013;280:20130077.

234. Luo YH-L, Zhong JJ, da Cruz L. The use of Argus® II retinal prosthesis by blind subjects to achieve localisation and prehension of objects in 3-dimensional space. *Graefe's Archive for Clinical and Experimental Ophthalmology* 2015;253:1907-1914.

235. Caspi A, Roy A, Wuyyuru V, et al. Eye movement control in the Argus II retinal-prosthesis enables reduced head movement and better localization precision. *Invest Ophthalmol Vis Sci* 2018;59:792-802.

236. Dagnelie G, Caspi A, Barry MP, Gibson P, Seifert G, Roy A. Thermal imaging prototype enhances person identification and warm object localization by Argus II wearers. *Invest Ophthalmol Vis Sci* 2016;57:5167-5167.

Biography

Michael Patrick Barry was born in Point Pleasant, New Jersey, USA in 1987. He attended Brick Township public schools and was very active in Boy Scout Troop 38 of Brick, NJ. He earned his Eagle Scout rank in 2003. Intrigued by the brain and the field of neuroprosthetics, Michael chose to attend Johns Hopkins University and major in neuroscience after graduating from Brick Township High School in 2005.

As part of his undergraduate studies, Michael joined the lab of Gislin Dagnelie, Ph.D. in 2007. Michael initially worked under Veronika Mueller on simulations of prosthetic vision. After enrolling in the university's B.A./M.S. program for neuroscience in 2009, Michael gained the opportunity to work with subjects participating in the Argus II Retinal Stimulation System Feasibility Study. Michael's project for the B.A./M.S. program asked subjects to perform a hand-eye coordination task that was originally developed for prosthetic vision simulations. Michael graduated with a B.A./M.S. degree from Johns Hopkins in 2010, and the results of the study were published in 2012.

Seeking to continue research in prosthetic vision, Michael enrolled in the biomedical engineering Ph.D. program at Johns Hopkins in 2010. He continued to work under Dr. Dagnelie and investigated previously unexplored aspects of prosthetic visuomotor coordination. Michael presented research regularly at annual meetings of the Association for Research in Vision and Ophthalmology (ARVO) starting in 2010. He also presented at The Eye and The Chip World Research Congresses in 2010 and later years. Michael published several articles and contributed to book chapters on the topics of prosthetic vision and low-vision rehabilitation. Michael was awarded a 2-year fellowship from the Johns Hopkins Visual Neuroscience Training Program in

2014. During the 2017 ARVO Annual Meeting, Michael was honored with the Envision-Atwell Award by the Low Vision Research Group.

As part of his Ph.D. program, Michael served as a teaching assistant for Introduction to Rehabilitation Engineering under Scott M. Paul, M.D. in 2014. Michael was later the head teaching assistant for Systems Bioengineering II Lab under Eileen B. Haase, Ph.D. in 2015. Michael also served as the head instructor of the Johns Hopkins Taekwondo Club until 2017, and continues to serve as an Assistant Scoutmaster for Boy Scout Troop 38.