

UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

Towards a Smart Bionic Eye: AI-powered artificial vision for the treatment of incurable blindness.

Permalink

<https://escholarship.org/uc/item/26s3b5jk>

Journal

Journal of Neural Engineering, 19(6)

Authors

Sanchez-Garcia, Melani
Beyeler, Michael

Publication Date

2022-12-07

DOI

10.1088/1741-2552/aca69d

Peer reviewed



HHS Public Access

Author manuscript

J Neural Eng. Author manuscript; available in PMC 2023 September 19.

Published in final edited form as:

J Neural Eng. ; 19(6): . doi:10.1088/1741-2552/aca69d.

Towards a *Smart Bionic Eye*: AI-Powered Artificial Vision for the Treatment of Incurable Blindness

Michael Beyeler,

Department of Computer Science, Department of Psychological & Brain Sciences, University of California, Santa Barbara, CA, USA

Melani Sanchez-Garcia

Department of Computer Science, University of California, Santa Barbara, CA, USA

Abstract

Objective.—How can we return a functional form of sight to people who are living with incurable blindness? Despite recent advances in the development of visual neuroprostheses, the quality of current prosthetic vision is still rudimentary and does not differ much across different device technologies.

Approach.—Rather than aiming to represent the visual scene as naturally as possible, a *Smart Bionic Eye* could provide visual augmentations through the means of artificial intelligence (AI)-based scene understanding, tailored to specific real-world tasks that are known to affect the quality of life of people who are blind, such as face recognition, outdoor navigation, and self-care.

Main results.—Complementary to existing research aiming to restore natural vision, we propose a patient-centered approach to incorporate deep learning-based visual augmentations into the next generation of devices.

Significance.—The ability of a visual prosthesis to support everyday tasks might make the difference between abandoned technology and a widely adopted next-generation neuroprosthetic device.

Keywords

visual prosthesis; artificial vision; artificial intelligence; computer vision

1. Introduction

How can we return a functional form of sight to people who are living with incurable blindness? Few disabilities affect human life more than the loss of the ability to see. Although recent advances in gene and stem cell therapies (e.g., Russell et al., 2017, da Cruz et al., 2018; for a recent review see McGregor, 2019) as well as retinal sheet transplants (e.g., Foik et al., 2018, Gasparini et al., 2019; for a recent commentary see Beyeler, 2019) are showing great promise as near-future treatment options for end-stage retinal degeneration, and some affected individuals can be treated with surgery or medication,

there are currently no effective treatments for many people blinded by severe degeneration or damage to the retina, the optic nerve, or cortex. In such cases, an electronic visual prosthesis (*bionic eye*) may be the only option (Fernandez, 2018; Roska and Sahel, 2018). Analogous to cochlear implants, these devices electrically stimulate surviving cells in the visual pathway to evoke visual percepts (*phosphenes*). Whereas there is only one regulatory-approved gene therapy (Luxturna), three visual prostheses have been commercialized over the years (Second Sight's Argus II, Retina Implant AG's Alpha-AMS, and Pixium Vision's IRIS II). Existing devices generally provide an improved ability to localize high-contrast objects and to perform basic orientation & mobility tasks (Geruschat et al., 2012; Karapanos et al., 2021).

However, the prosthetic vision generated by current retinal implants is still rudimentary and does not differ much across different device technologies (Erickson-Davis and Korzybska, 2021). Analogous to the first generation of cochlear implants, these devices have relied on straightforward signal processing and encoding schemes, assuming that each electrode in the array can be thought of as a "pixel" in an image (Dagnelie et al., 2007; Chen et al., 2009; Perez-Yus et al., 2017; Sanchez-Garcia et al., 2019); to generate a complex visual experience, one then simply needs to turn on the right combination of pixels. In contrast, current prosthesis users report seeing highly distorted phosphenes, which vary in shape across subjects as well as electrodes and often fail to assemble into more complex percepts (Wilke et al., 2011; Beyeler et al., 2019; Beauchamp et al., 2020; Erickson-Davis and Korzybska, 2021; Fernández et al., 2021). In the case of epiretinal implants, these distortions are largely due to inadvertent activation of passing axon fibers (Rizzo et al., 2003; Beyeler et al., 2019), but other device technologies based on electrical stimulation of visual cortex or optogenetics may face related issues. On the one hand, optogenetic prostheses may cause perceptual distortions due to differences in temporal dynamics between the optogenetic molecules and normal photopigments (Fine and Boynton, 2015). On the other hand, although there is a long history of patients reporting punctate percepts (sometimes described as "a star in the sky") in response to single-electrode stimulation of the visual cortex (Dobelle and Mladejovsky, 1974; Evans et al., 1979; Dobelle, 2000; Bosking et al., 2017), more recent work has highlighted that the percepts resulting from multi-electrode stimulation cannot be explained by a summative model based on single-electrode phosphenes (Beauchamp et al., 2020; Barry et al., 2020; Fernández et al., 2021).

While much work has focused on either making use of these documented distortions (Srivastava et al., 2009; Kiral-Kornek et al., 2013; Beyeler et al., 2019; Bruce and Beyeler, 2022) or finding ways to avoid them (Vilkhu et al., 2021; de Ruyter van Steveninck et al., 2022; Granley et al., 2022), these often theoretical insights have yet to be incorporated into a new generation of implantable technology.

2. Towards a *Smart Bionic Eye*

Rather than aiming to one day restore *natural* vision with visual prostheses (which may remain elusive until we fully understand the neural code of vision), we might be better off thinking about how to create practical and useful *artificial* vision now. Specifically, a visual prosthesis has the potential to provide visual augmentations through the means

of artificial intelligence (AI) based scene understanding (see Fig. 1), tailored to specific real-world tasks that are known to affect the quality of life of people who are blind (*e.g.*, wayfinding & navigation, face recognition, self-care). With recent breakthroughs in deep learning-based computer vision and AI, it is timely to consider how this work may best complement existing lines of animal and human behavioral research to inform the design of a next-generation visual prosthesis.

Instead of aiming to represent the visual scene as naturally as possible, a *Smart Bionic Eye* could locate the misplaced keys in the living room (Fig. 1, “Visual search”), read out medication labels (“Screen reader”), inform a user about people’s gestures and facial expressions (“Conversation”) during social interactions, or warn of nearby obstacles and outline safe paths (“Navigation”) when the user is going for a walk. Such a device could take inspiration from existing low vision aids (Htike et al., 2021), which do not promise any kind of sight restoration, but increasingly rely on AI to deliver functionality at a practical level (*e.g.*, Microsoft’s Seeing AI and Google Lookout are using computer vision to identify packaged food, and screen readers to read visually captured text aloud).

Indeed, we are not the first to point out that computer vision (and more generally: deep learning-based AI) may have an important role to play in visual prosthesis design (Barnes, 2012; Islam et al., 2019). A variety of studies have used simulations of prosthetic vision to demonstrate the benefit of simplifying the visual scene; for instance, by enhancing certain regions of interest (Boyle et al., 2008; Al-Atabany et al., 2010), highlighting visually salient information (Parikh et al., 2010; Li et al., 2018), segmenting important objects (Horne et al., 2016; Sanchez-Garcia et al., 2019, 2020b; Han et al., 2021), or segmenting nearby obstacles (McCarthy and Barnes, 2014; Rasla and Beyeler, 2022). However, although these studies are valuable in that they provide insights and specific hypotheses about the role of image processing and stimulus optimization for prosthetic vision, most of them were based on hypothetical future devices, did not involve prosthesis patients, or relied on overly simplified simulations that assumed phosphenes to be small, isolated, and independent light sources. It is therefore unclear how these findings would translate to real prosthesis patients. Only a handful studies have validated their computer vision algorithm on sighted subjects viewing prosthetic vision simulations (*e.g.*, McCarthy et al., 2014, Sanchez-Garcia et al., 2020b, Han et al., 2021), and even fewer have tested their setup with real prosthesis patients (two notable examples: He et al., 2019; Sadeghi et al., 2021).

However, with recent advances in computer vision and AI, the time is now to re-visit these ideas. It is only through the advent of deep learning that we can extract depth from a single image (without the need for extra sensors and bulky peripherals), that we can segment objects according to semantic labels, or that we can converse with an AI that understands our intention. In addition, the rapid development of deep learning-specific hardware (*e.g.*, Intel’s Neural Compute Stick) may soon allow these models to be deployed in real time in an energy-efficient way.

Ultimately, the ability of a visual prosthesis to support everyday tasks might make the difference between abandoned technology and a widely adopted next-generation neuroprosthetic device. Indeed, when Retina Implant AG (maker of the Alpha-IMS/AMS

subretinal implants) dissolved in March 2019, they cited their device not leading to “the concrete benefit in everyday life of those affected”[†] as one of the main reasons for shutting down.

2.1. The Scientific Challenge

How do we arrive at a *Smart Bionic Eye*? Achieving this ambitious goal will certainly require the engineering of next-generation visual prostheses with large electrode counts (Ferlauto et al., 2018; Shah and Chichilnisky, 2020; Chen et al., 2020) and the development of sophisticated AI systems. However, the challenge is less about dreaming up new computer vision algorithms and more about identifying the design principles and visual cues that are best suited to augment the visual scene in a way that supports behavioral performance for a potentially heterogeneous end-user demographic. For example, humans are able to flexibly adapt their visual navigation strategies depending on the visual cues that are available to them—in texture-rich environments they might use optic flow, but in texture-scarce environments they might rely on the perceived location of the goal, together with extraretinal information about their head and eye position (Turano et al., 2005). Furthermore, these strategies change under central and peripheral vision loss (Turano et al., 2001). How do we know which visual navigation cues are best suited for visual prosthesis patients?

Another concern is that the vision tests typically used in clinics and psychophysics laboratories (*e.g.*, perimetry, acuity, contrast sensitivity, orientation discrimination) are not designed to test the ability of prosthetic devices to restore vision (Peli, 2020). The main reason for this is the nature of the multi-alternative forced choice (MAFC) paradigm that is typically used to administer these tests. As such, they may not measure what the researchers intended, either because nuisance variables may provide spurious cues that can be learned in repeated training or because the tests can be passed without form vision (Peli, 2020). Consequently, superior performance on these tests does not necessarily imply sight restoration.

2.2. The Proposed Solution

To address these challenges, we propose to a patient-centered approach to incorporating AI-powered visual augmentations into the next generation of implanatable technology.

Most prosthesis designs share a common set of components: a camera to capture images, generally mounted on glasses; a video processing unit (VPU) that transform the visual scene into patterns of electrical stimulation and transmits this information through a radio-frequency link to the implanted device, and an electrode array implanted somewhere along the visual pathway.

The conventional approach to stimulus encoding, as implemented by previously commercialized devices such as Alpha-AMS and Argus II, is typically very simple, assuming a linear relationship between the gray level of a pixel in the captured image and the stimulating amplitude (Fig. 2A). Several studies have already proposed more

[†] <https://web.archive.org/web/20200805082212/https://www.retina-implant.de/en>

sophisticated stimulus encoding strategies to recreate a desired neural activity pattern over a given temporal window (Shah et al., 2019; Spencer et al., 2019; Ghaffari et al., 2021). However, none of these approaches are able to predict the perceptual consequences of the resulting neural activity.

We thus suggest an iterative workflow that begins and ends with the patient (see Fig. 2B). In line with research practices in the human-computer interaction (HCI) community, the first step is to identify the information needs of the end user through a series of qualitative and quantitative studies. This may involve low vision users navigating a virtual environment to (*e.g.*) avoid obstacles or reach a goal location. Their struggles and challenges may then inform the visual cues that are required to perform the task (Hoogsteen et al., 2022), which may lead to task-specific visual augmentation strategies. These strategies can be refined using qualitative feedback from the end user (*i.e.*, in which way do they prefer the information to be presented?) as well as their behavioral performance (*i.e.*, which strategies are most effective?). Finally, strategies that perform well in the training environment can be tested on real prosthesis patients. Below we expand on these ideas.

2.2.1. Patient-Centered Design—As pointed out by Htike et al. (2020) and Erickson-Davis and Korzybska (2021), the majority of research on visual prostheses (and more generally: low vision aids) has focused on the technical aspects rather than the usability of these devices. One promising development has been the Functional Low-Vision Observer Rated Assessment (FLORA), a tool to provide a subjective assessment to capture the functional visual ability and well-being of visual prosthesis patients (Geruschat et al., 2015). While it is encouraging to see increasing adoption of FLORA by the community (Geruschat et al., 2016; Karapanos et al., 2021), in practice it is often employed as an external validation tool that constitutes the very last step of the design process—a proof of concept, so to speak. However, if the proof of concept fails, researchers must start over and try again until they have found a better way to improve FLORA performance of their subjects.

This is in stark contrast to research practices of the HCI community, which typically aims to incorporate end users in the decision making and development during every step of the design process (Rubin and Chisnell, 2011; Lee et al., 2017). In particular, patient-centered design (PCD) is a methodology that aims to make systems usable and useful by first-and-foremost focusing on the needs and requirements of the patient (Reis et al., 2011; Light, 2019). Using a combination of clinical and technical tests, feedback and questionnaires, PCD can inform what potential end users may want out of a visual prosthesis, where and how they would use it, and what features they would consider essential. These tests may be conducted during each stage of the design process to ensure that development proceeds with the user as the center of focus (Rubin and Chisnell, 2011).

While this feedback may not be the solution to all problems related to the optimal encoding of visual information, it may represent an important first step towards developing more usable prosthetic devices that may complement existing lines of research that focus on prototyping with animal models or simulation systems. In a recent systematic review (Kasowski et al., 2022), we showed that although there is no shortage of publications that demonstrate a proof-of-concept augmentation strategy, less emphasis has been placed on

understanding the usability of their proposed technology. Involving appropriate end users in all stages of the design process may ultimately improve the effectiveness and accessibility of the technology as well as user satisfaction (Schicktanz et al., 2015).

2.2.2. Virtual Prototyping—Due to the unique requirements of working with bionic eye recipients (*e.g.*, constant assistance, increased setup time, travel cost), experimentation with new stimulation strategies remains time-consuming and expensive.

In the interim, a more cost-effective and increasingly popular alternative might be to rely on an immersive virtual reality (VR) prototype based on simulated prosthetic vision (SPV) (Zapf et al., 2014; Sanchez-Garcia et al., 2020a; Thorn et al., 2020; Kasowski and Beyeler, 2022). Here, the classical method relies on sighted subjects wearing a VR head-mounted display (HMD), who are then deprived of natural viewing and only perceive phosphenes displayed in the HMD. This allows sighted participants to “see” through the eyes of the bionic eye user, taking into account their head and/or eye movements as they explore a virtual environment. The visual scene can then be manipulated according to any desired image processing or visual augmentation strategy (Han et al., 2021).

In order for simulation results to translate to real prosthesis patients, simulations should rely on psychophysically validated phosphene models and employ a restricted field of view that necessitates head scanning (Kasowski and Beyeler, 2022). In addition, sighted participants in SPV studies are often sampled from the university’s undergraduate population (for practical reasons). Their age, navigational affordances, and experience with low vision may therefore be drastically different from real bionic eye users, who tend to not only be older and prolific cane users but also receive extensive vision rehabilitation training. For instance, Williams et al. (2014) compared sighted and blind navigation and found that both groups understand navigation differently, leading sighted people to struggle in guiding blind companions. Furthermore, blind people use a combination of devices and technology to complement their existing orientation and mobility skills (Williams et al., 2014), which may lead to a wide variety of navigation styles (Ahmetovic et al., 2019; Htike et al., 2020). An important step towards designing more usable visual prosthetics may thus be to recruit age-appropriate participants for SPV studies.

If done right, the use of a VR prototype may drastically speed up the development process by testing theoretical predictions in high-throughput experiments, the best of which can be validated and improved upon in an iterative process with the bionic eye recipient in the loop (Kasowski et al., 2021).

2.2.3. Visual Augmentations to Support Real-World Tasks—Most visual prostheses are equipped with an external video processing unit (VPU) capable of applying simple image processing techniques to the video feed in real time. For instance, edge detection and contrast maximization are already routinely used in current devices. In the near future, these techniques may include deep learning-based algorithms aimed at improving a patient’s scene understanding.

Based on this premise, researchers have developed various image optimization strategies, and assessed their performance by having sighted observers conduct daily visual tasks under SPV (Dagnelie et al., 2007; Al-Atabany et al., 2010; Li et al., 2018; McCarthy et al., 2014). Simulation allows a wide range of computer vision systems to be developed and tested without requiring implanted devices. SPV studies suggest that one benefit of image processing may be to provide an importance mapping that can aid scene understanding; that is, to enhance certain image features or regions of interest, at the expense of discarding less important or distracting information (Boyle et al., 2008; Al-Atabany et al., 2010; Horne et al., 2016; Sanchez-Garcia et al., 2019). This limited compensation may be significant to retinal prosthesis patients carrying out visual tasks in daily life.

Examples of suitable strategies are shown in Fig. 3. Such strategies may include simplifying the scene by segmenting objects of interest from background clutter (Fig. 3A), highlighting nearby obstacles by substituting relative depth for intensity (Fig. 3 B), helping a user orient themselves in the room by highlighting structural edges of indoor environments (Fig. 3C). Plenty of these ideas have already been tested in isolation (*e.g.*, Al-Atabany et al., 2010, Parikh et al., 2010, McCarthy and Barnes 2014, Horne et al., 2016), but more research is needed to compare these approaches side-by-side (Han et al., 2021) and to test their ability to support real-world tasks (He et al., 2019; Sadeghi et al., 2021)

A so-far unexplored application domain concerns the use of visual question answering (VQA) to help a user retrieve misplaced items or orient themselves in their environment (Fig. 3D). VQA models (*e.g.*, Antol et al., 2015) are able to give a visual answer to a verbal question; for example, in response to the question “How many giraffes are drinking water?” and a given image, the network would respond by drawing bounding boxes around all the giraffes drinking from the water hole (but not the other ones, even if they are standing *by* the water hole). In the context of the *Smart Bionic Eye*, VQA models would allow a user to ask questions such as “Where did I put my keys again?”, and the system would respond by segmenting the keys in the prosthetic image while the user is looking around the room (see also Fig. 1).

Other concrete examples to support practical tasks might include 1) an *outdoor navigation mode*, where we may need to test the utility of highlighting nearby obstacles, highlighting the goal location, or outlining structural edges to let a user orient themselves in the environment, and 2) a *conversation mode*, where we may need to test the utility of highlighting different facial features to allow for face discrimination, highlighting the person that is currently speaking to determine whether they are addressing the user or someone else, or notifying the user of people entering or leaving the room. Importantly, these ideas should constitute only the beginning of a conversation with potential end users, such that the proposed solution can be iteratively refined based on both qualitative feedback from real patients and quantitative measures from virtual patients with the VR prototype.

It is easy to see how the above deep learning techniques could become an integral part of the *Smart Bionic Eye* once they reach a certain maturity that allows them to be used in unstructured environments. In the future, these visual augmentations could be combined with GPS to give directions, warn users of impending dangers in their immediate

surroundings, or even extend the range of “visible” light with the use of an infrared sensor (Sadeghi et al., 2021). Once the quality of the generated artificial vision experience reaches a certain threshold, there are a lot of exciting avenues to pursue.

2.3. Challenges & Limitations

Despite its potential, development of a *Smart Bionic Eye* faces a number of challenges and limitations, which we briefly address below.

2.3.1. Risks & Benefits—At the core of the question about whether to develop and implant a *Smart Bionic Eye* lies a risk/benefit assessment. Indeed, the AI-powered algorithms outlined above could also be used as input to other low-vision devices such as smart glasses and sensory substitution devices, which do not necessitate risky and invasive surgery. Future patients thinking about whether to implant should therefore not only consider device safety and efficacy data in their decision, but should also be informed about less-invasive alternatives that may deliver similar benefits.

That being said, one advantage that a *Smart Bionic Eye* could offer over nonvisual alternatives is a combination of both a conventional “natural vision” mode next to a number of “artificial vision” modes designed to support everyday tasks. Such a device (though invasive and expensive) might thus be superior to other accessibility aids such as smartphone apps and sensory substitution devices, because it could directly tap into the visual cortex of a blind user to make them see. On the other hand, one might also consider a next-generation device to combine the benefits of prosthetic vision with other sensory augmentations (Kvansakul et al., 2020).

2.3.2. Neural Code of Vision—A major outstanding challenge is translating electrode stimulation into a code that the brain can understand. Interactions between the device electronics and the underlying neurophysiology can lead to perceptual distortions that severely limit the quality of the generated visual experience (Fine and Boynton, 2015; Beyeler et al., 2019; Erickson-Davis and Korzybska, 2021). One possibility is thus that we must first address fundamental questions about the neural code of vision (Abbasi and Rizzo, 2021) and (the lack of) cortical plasticity in adult visual cortex (Beyeler et al., 2017), before we can explore AI-based visual augmentations.

However, since the goal is not primarily to create *natural* vision, it suffices that phosphene characteristics are distinct and stable over time, which is the case for current implants (Luo et al., 2016; Fernández et al., 2021). In addition, there often exists a numeric or symbolic forward model, constrained by empirical data, that can predict a neuronal or ideally perceptual response to the applied stimulus (Bosking et al., 2017; Beyeler et al., 2019). To find the stimulus that will elicit a desired response, one essentially needs to find the inverse of the forward model, which can be achieved in a number of ways (Spencer et al., 2019; Fauvel and Chalk, 2022; Granley et al., 2022).

2.3.3. Robustness & Safety—It can be downright dangerous to allow computer vision algorithms to operate in the real world without people in the loop. These AI systems can make serious mistakes that no sane human would make (Hole and Ahmad, 2021). For

example, it is possible to make subtle changes to images and objects that fool vision-based AI systems into misclassifying objects. This can have grave consequences if the system is relied upon to warn of impending dangers, such as an approaching car, where a false negative could be fatal.

However, this issue is not unique to the *Smart Bionic Eye*, but affects applications ranging from self-driving cars to remote sensing and medical imaging. While more work is needed to improve the robustness of vision-based AI systems in real-world scenarios, potential solutions may range from techniques to improve model performance under naturally-induced image corruptions and alterations (Drenkow et al., 2021) to human-machine partnership (Patel et al., 2019; Fauvel and Chalk, 2022).

2.3.4. Engineering—Even if the stimulus encoding problem and safety issues are solved, there remains the question of how to fit a sophisticated AI system on a low-power, portable “edge device” such as a VPU.

Although still an active field of research, a potential solution may take the form of a serverless cloud service (Zhang et al., 2021), as is currently being developed for Internet of Things (IoT) solutions, or deep learning-specific neuromorphic hardware, such as Intel’s Neural Compute Stick. While the latter has the potential to dramatically improve the latency, robustness, and power consumption compared to traditional computers, new computer vision algorithms are needed to process the unconventional output of neuromorphic sensors to unlock their potential (Gallego et al., 2022; Sanchez-Garcia et al., 2022). In addition, since people who are blind tend to spend a lot of time indoors (Jeamwathanachai et al., 2019), it is not outlandish to assume that a *Smart Bionic Eye* could be shipped with a central desktop computer that would handle most of the computationally expensive processing while communicating wirelessly with the external glasses of the implant.

3. Conclusion

In this letter, we propose to complement existing lines of bionic vision research with a patient-centered approach that considers the possibility of a visual prosthesis to function as an AI-powered visual aid. This *Smart Bionic Eye* would harness recent developments in deep learning-based computer vision and AI to provide useful visual augmentations for everyday tasks.

To enable such a technology, we first need to address fundamental questions at the intersection of neuroscience, engineering, and human-computer interaction to better understand how visual prostheses interact with the human visual system to shape perception (Beyeler et al., 2017; Abbasi and Rizzo, 2021) and to identify visual augmentation strategies that best support specific real-world tasks (Han et al., 2021). This advance in technology could improve the ability of a visual prosthesis to support everyday tasks and lead to a successful next-generation neuroprosthetic device.

References

- Abbasi B and Rizzo JF (2021). Advances in Neuroscience, Not Devices, Will Determine the Effectiveness of Visual Prostheses. *Seminars in Ophthalmology*, 0(0):1–8. Publisher: Taylor & Francis eprint: 10.1080/08820538.2021.1887902.
- Ahmetovic D, Guerreiro J, Ohn-Bar E, Kitani KM, and Asakawa C (2019). Impact of Expertise on Interaction Preferences for Navigation Assistance of Visually Impaired Individuals. In *Proceedings of the 16th International Web for All Conference, W4A '19*, pages 1–9, New York, NY, USA. Association for Computing Machinery.
- Al-Atabany WI, Tong T, and Degenaar PA (2010). Improved content aware scene retargeting for retinitis pigmentosa patients. *Biomed Eng Online*, 9:52. [PubMed: 20846440]
- Antol S, Agrawal A, Lu J, Mitchell M, Batra D, Zitnick CL, and Parikh D (2015). VQA: Visual Question Answering. pages 2425–2433.
- Barnes N (2012). The role of computer vision in prosthetic vision. *Image and Vision Computing*, 30(8):478–479.
- Barry MP, Armenta Salas M, Patel U, Wuyyuru V, Niketeghad S, Bosking WH, Yoshor D, Dorn JD, and Pouratian N (2020). Video-mode percepts are smaller than sums of single-electrode phosphenes with the Orion[®] visual cortical prosthesis. *Investigative Ophthalmology & Visual Science*, 61(7):927.
- Beauchamp MS, Oswalt D, Sun P, Foster BL, Magnotti JF, Niketeghad S, Pouratian N, Bosking WH, and Yoshor D (2020). Dynamic Stimulation of Visual Cortex Produces Form Vision in Sighted and Blind Humans. *Cell*, 181(4):774–783.e5. Publisher: Elsevier. [PubMed: 32413298]
- Beyeler M (2019). Commentary: Detailed Visual Cortical Responses Generated by Retinal Sheet Transplants in Rats With Severe Retinal Degeneration. *Frontiers in Neuroscience*, 13.
- Beyeler M, Nanduri D, Weiland JD, Rokem A, Boynton GM, and Fine I (2019). A model of ganglion axon pathways accounts for percepts elicited by retinal implants. *Scientific Reports*, 9(1):1–16. [PubMed: 30626917]
- Beyeler M, Rokem A, Boynton GM, and Fine I (2017). Learning to see again: biological constraints on cortical plasticity and the implications for sight restoration technologies. *J Neural Eng*, 14(5):051003. [PubMed: 28612755]
- Bosking WH, Sun P, Ozker M, Pei X, Foster BL, Beauchamp MS, and Yoshor D (2017). Saturation in phosphene size with increasing current levels delivered to human visual cortex. *Journal of Neuroscience*.
- Boyle JR, Maeder AJ, and Boles WW (2008). Region-of-interest processing for electronic visual prostheses. *Journal of Electronic Imaging*, 17(1):013002. Publisher: International Society for Optics and Photonics.
- Bruce A and Beyeler M (2022). Greedy Optimization of Electrode Arrangement for Epiretinal Prostheses. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VII*, pages 594–603, Berlin, Heidelberg. Springer-Verlag.
- Chen SC, Suaning GJ, Morley JW, and Lovell NH (2009). Simulating prosthetic vision: I. Visual models of phosphenes. *Vision Research*, 49(12):1493–506. [PubMed: 19504749]
- Chen X, Wang F, Fernandez E, and Roelfsema PR (2020). Shape perception via a high-channel-count neuroprosthesis in monkey visual cortex. *Science*, 370(6521):1191–1196. Publisher: American Association for the Advancement of Science Section: Research Article. [PubMed: 33273097]
- da Cruz L, Fynes K, Georgiadis O, Kerby J, Luo YH, Ahmado A, Vernon A, Daniels JT, Nommiste B, Hasan SM, Gooljar SB, Carr A-JF, Vugler A, Ramsden CM, Bictash M, Fenster M, Steer J, Harbinson T, Wilbrey A, Tufail A, Feng G, Whitlock M, Robson AG, Holder GE, Sagoo MS, Loudon PT, Whiting P, and Coffey PJ (2018). Phase 1 clinical study of an embryonic stem cell-derived retinal pigment epithelium patch in age-related macular degeneration. *Nature Biotechnology*, 36(4):328–337.
- Dagnelie G, Keane P, Narla V, Yang L, Weiland J, and Humayun M (2007). Real and virtual mobility performance in simulated prosthetic vision. *J Neural Eng*, 4(1):S92–101. [PubMed: 17325421]

- de Ruyter van Steveninck J, Güçlü U, van Wezel R, and van Gerven M (2022). End-to-end optimization of prosthetic vision. *Journal of Vision*, 22(2):20.
- Dobelle WH (2000). Artificial Vision for the Blind by Connecting a Television Camera to the Visual Cortex. *ASAIO Journal*, 46(1):3–9. [PubMed: 10667705]
- Dobelle WH and Mladejovsky MG (1974). Phosphenes produced by electrical stimulation of human occipital cortex, and their application to the development of a prosthesis for the blind. *The Journal of Physiology*, 243(2):553–576. _eprint: 10.1113/jphysiol.1974.sp010766. [PubMed: 4449074]
- Drenkow N, Sani N, Shpitser I, and Unberath M (2021). Robustness in Deep Learning for Computer Vision: Mind the gap? Technical Report arXiv:2112.00639, arXiv. arXiv:2112.00639 [cs] type: article.
- Erickson-Davis C and Korzybska H (2021). What do blind people “see” with retinal prostheses? Observations and qualitative reports of epiretinal implant users. *PLOS ONE*, 16(2):e0229189. Publisher: Public Library of Science. [PubMed: 33566851]
- Evans JR, Gordon J, Abramov I, Mladejovsky MG, and Dobelle WH (1979). Brightness of phosphenes elicited by electrical stimulation of human visual cortex. *Sensory Processes*, 3(1):82–94. [PubMed: 515743]
- Fauvel T and Chalk M (2022). Human-in-the-loop optimization of visual prosthetic stimulation. *Journal of Neural Engineering*.
- Ferlauto L, Leccardi MJIA, Chenais NAL, Gilliéron SCA, Vagni P, Bevilacqua M, Wolfensberger TJ, Sivula K, and Ghezzi D (2018). Design and validation of a foldable and photovoltaic wide-field epiretinal prosthesis. *Nature Communications*, 9(1):1–15.
- Fernandez E (2018). Development of visual Neuroprostheses: trends and challenges. *Bioelectronic Medicine*, 4(1):12. [PubMed: 32232088]
- Fernández E, Alfaro A, Soto-Sánchez C, González-López P, Ortega AML, Peña S, Grima MD, Rodil A, Gómez B, Chen X, Roelfsema PR, Rolston JD, Davis TS, and Normann RA (2021). Visual percepts evoked with an Intracortical 96-channel microelectrode array inserted in human occipital cortex. *The Journal of Clinical Investigation*. Publisher: American Society for Clinical Investigation.
- Fine I and Boynton GM (2015). Pulse trains to percepts: the challenge of creating a perceptually intelligible world with sight recovery technologies. *Philos Trans R Soc Lond B Biol Sci*, 370(1677):20140208. [PubMed: 26240423]
- Foik AT, Lean GA, Scholl LR, McLelland BT, Mathur A, Aramant RB, Seiler MJ, and Lyon DC (2018). Detailed Visual Cortical Responses Generated by Retinal Sheet Transplants in Rats with Severe Retinal Degeneration. *Journal of Neuroscience*, 38(50):10709–10724. Publisher: Society for Neuroscience Section: Research Articles. [PubMed: 30396913]
- Gallego G, Delbrück T, Orchard G, Bartolozzi C, Taba B, Censi A, Leutenegger S, Davison AJ, Conrath J, Daniilidis K, and Scaramuzza D (2022). Event-Based Vision: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):154–180. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. [PubMed: 32750812]
- Gasparini SJ, Llonch S, Borsch O, and Ader M (2019). Transplantation of photoreceptors into the degenerative retina: Current state and future perspectives. *Progress in Retinal and Eye Research*, 69:1–37. [PubMed: 30445193]
- Geruschat DR, Bittner AK, and Dagnelie G (2012). Orientation and Mobility Assessment in Retinal Prosthetic Clinical Trials. *Optometry and Vision Science*, 89(9):1308–1315. [PubMed: 22902422]
- Geruschat DR, Flax M, Tanna N, Bianchi M, Fisher A, Goldschmidt M, Fisher L, Dagnelie G, Deremeik J, Smith A, Anafloos F, and Dorn J (2015). FLORA™: Phase I development of a functional vision assessment for prosthetic vision users. *Clinical and Experimental Optometry*, 98(4):342–347. [PubMed: 25675964]
- Geruschat DR, Richards TP, Arditi A, da Cruz L, Dagnelie G, Dorn JD, Duncan JL, Ho AC, Olmos de Koo LC, Sahel J, Stanga PE, Thumann G, Wang V, and Greenberg RJ (2016). An analysis of observer-rated functional vision in patients implanted with the Argus II Retinal Prosthesis System at three years. *Clinical & Experimental Optometry*, 99(3):227–232. [PubMed: 26804484]
- Ghaffari DH, Chang Y-C, Mirzakhali E, and Weiland JD (2021). Closed-loop Optimization of Retinal Ganglion Cell Responses to Epiretinal Stimulation: A Computational Study. In 2021

- 10th International IEEE/EMBS Conference on Neural Engineering (NER), pages 597–600. ISSN: 1948-3554.
- Godard C, Aodha OM, Firman M, and Brostow G (2019). Digging Into Self-Supervised Monocular Depth Estimation. In 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 3827–3837. ISSN: 2380-7504.
- Granley J, Relic L, and Beyeler M (2022). A Hybrid Neural Autoencoder for Sensory Neuroprostheses and Its Applications in Bionic Vision. Technical Report arXiv:2205.13623, arXiv: arXiv:2205.13623 [cs] type: article.
- Han N, Srivastava S, Xu A, Klein D, and Beyeler M (2021). Deep Learning–Based Scene Simplification for Bionic Vision. In Augmented Humans Conference 2021, AHs’21, pages 45–54, New York, NY, USA. Association for Computing Machinery.
- He Y, Huang NT, Caspi A, Roy A, and Montezuma SR (2019). Trade-Off Between Field-of-View and Resolution in the Thermal-Integrated Argus II System. *Translational Vision Science & Technology*, 8(4):29.
- Hole KJ and Ahmad S (2021). A thousand brains: toward biologically constrained AI. *SN Applied Sciences*, 3(8):743.
- Hoogsteen KM, Szpiro S, Kreiman G, and Peli E (2022). Beyond the Cane: Describing Urban Scenes to Blind People for Mobility Tasks. *ACM Transactions on Accessible Computing*. Just Accepted.
- Horne L, Alvarez J, McCarthy C, Salzmann M, and Barnes N (2016). Semantic labeling for prosthetic vision. *Computer Vision and Image Understanding*, 149:113–125.
- Htike HM, Margrain H, T., Lai Y-K, and Eslambolchilar P (2021). Augmented Reality Glasses as an Orientation and Mobility Aid for People with Low Vision: a Feasibility Study of Experiences and Requirements. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, number 729, pages 1–15. Association for Computing Machinery, New York, NY, USA.
- Htike HM, Margrain TH, Lai Y-K, and Eslambolchilar P (2020). Ability of Head-Mounted Display Technology to Improve Mobility in People With Low Vision: A Systematic Review. *Translational Vision Science & Technology*, 9(10).
- Islam MM, Sheikh Sadi M, Zamli KZ, and Ahmed MM (2019). Developing Walking Assistants for Visually Impaired People: A Review. *IEEE Sensors Journal*, 19(8):2814–2828. Conference Name: IEEE Sensors Journal.
- Jeamwathanachai W, Wald M, and Wills G (2019). Indoor navigation by blind people: Behaviors and challenges in unfamiliar spaces and buildings. *British Journal of Visual Impairment*, 37(2):140–153. Publisher: SAGE Publications Ltd.
- Karapanos L, Abbott CJ, Ayton LN, Kolic M, McGuinness MB, Baglin EK, Titchener SA, Kvensakul J, Johnson D, Kentler WG, Barnes N, Nayagam DAX, Allen PJ, and Petoe MA (2021). Functional Vision in the Real-World Environment With a Second-Generation (44-Channel) Suprachoroidal Retinal Prosthesis. *Translational Vision Science & Technology*, 10(10):7–7. Publisher: The Association for Research in Vision and Ophthalmology.
- Kasowski J and Beyeler M (2022). Immersive Virtual Reality Simulations of Bionic Vision. In Augmented Humans 2022, AHs 2022, pages 82–93, New York, NY, USA. Association for Computing Machinery.
- Kasowski J, Johnson BA, Neydavood R, Akkaraju A, and Beyeler M (2022). A Systematic Review of Extended Reality (XR) for Understanding and Augmenting Vision Loss. arXiv:2109.04995 [cs].
- Kasowski J, Wu N, and Beyeler M (2021). Towards Immersive Virtual Reality Simulations of Bionic Vision. In Augmented Humans Conference 2021, AHs’21, pages 313–315, New York, NY, USA. Association for Computing Machinery.
- Kiral-Kornek FI, Savage CO, O’Sullivan-Greene E, Burkitt AN, and Grayden DB (2013). Embracing the irregular: A patient-specific image processing strategy for visual prostheses. In 2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pages 3563–3566. ISSN: 1558-4615.
- Kvensakul J, Hamilton L, Ayton LN, McCarthy C, and Petoe MA (2020). Sensory augmentation to aid training with retinal prostheses. *Journal of Neural Engineering*, 17(4):045001. [PubMed: 32554868]

- Lee J, Wickens C, Liu Y, and Boyle L (2017). Designing for People: An introduction to human factors engineering.
- Li H, Su X, Wang J, Kan H, Han T, Zeng Y, and Chai X (2018). Image processing strategies based on saliency segmentation for object recognition under simulated prosthetic vision. *Artificial Intelligence in Medicine*, 84:64–78. [PubMed: 29129481]
- Light G (2019). User-Centered Design Strategies for Clinical Brain-Computer Interface Assistive Technology Devices. Walden Dissertations and Doctoral Studies.
- Lin T-Y, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P, and Zitnick CL (2014). Microsoft COCO: Common Objects in Context. In Fleet D, Pajdla T, Schiele B, and Tuytelaars T, editors, *Computer Vision – ECCV 2014*, Lecture Notes in Computer Science, pages 740–755, Cham. Springer International Publishing.
- Luo YH, Zhong JJ, Clemo M, and da Cruz L (2016). Long-term Repeatability and Reproducibility of Phosphene Characteristics in Chronically Implanted Argus(R) II Retinal Prosthesis Subjects. *Am J Ophthalmol*.
- McCarthy C and Barnes N (2014). Importance weighted image enhancement for prosthetic vision: An augmentation framework. In 2014 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pages 45–51.
- McCarthy C, Walker JG, Lieby P, Scott A, and Barnes N (2014). Mobility and low contrast trip hazard avoidance using augmented depth. *Journal of Neural Engineering*, 12(1):016003. Publisher: IOP Publishing. [PubMed: 25427135]
- McGregor JE (2019). Restoring vision at the fovea. *Current Opinion in Behavioral Sciences*, 30:210–216. [PubMed: 33553510]
- Parikh N, Itti L, and Weiland J (2010). Saliency-based image processing for retinal prostheses. *Journal of Neural Engineering*, 7(1):016006. Publisher: IOP Publishing.
- Patel BN, Rosenberg L, Willcox G, Baltaxe D, Lyons M, Irvin J, Rajpurkar P, Amrhein T, Gupta R, Halabi S, Langlotz C, Lo E, Mammarrappallil J, Mariano AJ, Riley G, Seekins J, Shen L, Zucker E, and Lungren MP (2019). Human-machine partnership with artificial intelligence for chest radiograph diagnosis. *npj Digital Medicine*, 2(1):1–10. Number: 1 Publisher: Nature Publishing Group. [PubMed: 31304351]
- Peli E (2020). Testing Vision Is Not Testing For Vision. *Translational Vision Science & Technology*, 9(13):32–32. Publisher: The Association for Research in Vision and Ophthalmology.
- Perez-Yus A, Bermudez-Cameo J, Lopez-Nicolas G, and Guerrero JJ (2017). Depth and Motion Cues With Phosphene Patterns for Prosthetic Vision. pages 1516–1525.
- Rasla A and Beyeler M (2022). The Relative Importance of Depth Cues and Semantic Edges for Indoor Mobility Using Simulated Prosthetic Vision in Immersive Virtual Reality. arXiv:2208.05066 [cs].
- Reis CI, Freire CS, Fernández J, and Monguet JM (2011). Patient Centered Design: Challenges and Lessons Learned from Working with Health Professionals and Schizophrenic Patients in e-Therapy Contexts. In Cruz-Cunha MM, Varajão J, Powell P, and Martinho R, editors, *ENTERprise Information Systems, Communications in Computer and Information Science*, pages 1–10, Berlin, Heidelberg. Springer.
- Rizzo JF, Wyatt J, Loewenstein J, Kelly S, and Shire D (2003). Perceptual efficacy of electrical stimulation of human retina with a microelectrode array during short-term surgical trials. *Invest Ophthalmol Vis Sci*, 44(12):5362–9. [PubMed: 14638739]
- Roska B and Sahel J-A (2018). Restoring vision. *Nature*, 557(7705):359–367. [PubMed: 29769667]
- Rubin J and Chisnell D (2011). *Handbook of Usability Testing: How to Plan, Design, and Conduct Effective Tests*. John Wiley & Sons. Google-Books-ID: 1e1MmVzMb0C.
- Russell S, Bennett J, Wellman JA, Chung DC, Yu Z-F, Tillman A, Wittes J, Pappas J, Elci O, McCague S, Cross D, Marshall KA, Walshire J, Kehoe TL, Reichert H, Davis M, Raffini L, George LA, Hudson FP, Dingfield L, Zhu X, Haller JA, Sohn EH, Mahajan VB, Pfeifer W, Weckmann M, Johnson C, Gewaily D, Drack A, Stone E, Wachtel K, Simonelli F, Leroy BP, Wright JF, High KA, and Maguire AM (2017). Efficacy and safety of voretigene neparvovec (AAV2-hRPE65v2) in patients with RPE65-mediated inherited retinal dystrophy: a randomised, controlled, open-label, phase 3 trial. *The Lancet*, 390(10097):849–860. Publisher: Elsevier.

- Sadeghi R, Kartha A, Barry MP, Bradley C, Gibson P, Caspi A, Roy A, and Dagnelie G (2021). Glow in the dark: Using a heat-sensitive camera for blind individuals with prosthetic vision. *Vision Research*, 184:23–29. [PubMed: 33780753]
- Sanchez-Garcia M, Chauhan T, Cottureau BR, and Beyeler M (2022). Efficient visual object representation using a biologically plausible spike-latency code and winner-take-all inhibition. Technical Report arXiv:2205.10338, arXiv. arXiv:2205.10338 [cs] type: article.
- Sanchez-Garcia M, Martinez-Cantin R, Bermudez-Cameo J, and Guerrero-Campo JJ (2020a). Influence of field of view in visual prostheses design: Analysis with a VR system. *Journal of Neural Engineering*.
- Sanchez-Garcia M, Martinez-Cantin R, and Guerrero JJ (2019). Indoor Scenes Understanding for Visual Prosthesis with Fully Convolutional Networks. In VISIGRAPP.
- Sanchez-Garcia M, Martinez-Cantin R, and Guerrero JJ (2020b). Semantic and structural image segmentation for prosthetic vision. *PLOS ONE*, 15(1):e0227677. [PubMed: 31995568]
- Schick Tanz S, Amelung T, and Rieger JW (2015). Qualitative assessment of patients' attitudes and expectations toward BCIs and implications for future technology development. *Frontiers in Systems Neuroscience*, 9.
- Shah NP and Chichilnisky EJ (2020). Computational challenges and opportunities for a bi-directional artificial retina. *Journal of Neural Engineering*, 17(5):055002. Publisher: IOP Publishing. [PubMed: 33089827]
- Shah NP, Madugula S, Grosberg L, Mena G, Tandon P, Hottowy P, Sher A, Litke A, Mitra S, and Chichilnisky E (2019). Optimization of Electrical Stimulation for a High-Fidelity Artificial Retina. In 2019 9th International IEEE/EMBS Conference on Neural Engineering (NER), pages 714–718. ISSN: 1948-3554.
- Spencer MJ, Kameneva T, Grayden DB, Meffin H, and Burkitt AN (2019). Global activity shaping strategies for a retinal implant. *Journal of Neural Engineering*, 16(2):026008. Publisher: IOP Publishing. [PubMed: 30523981]
- Srivastava NR, Troyk PR, and Dagnelie G (2009). Detection, eye–hand coordination and virtual mobility performance in simulated vision for a cortical visual prosthesis device. *Journal of Neural Engineering*, 6(3):035008. [PubMed: 19458397]
- Thorn JT, Migliorini E, and Ghezzi D (2020). Virtual reality simulation of epiretinal stimulation highlights the relevance of the visual angle in prosthetic vision. *Journal of Neural Engineering*. Publisher: IOP Publishing.
- Turano KA, Geruschat DR, Baker FH, Stahl JW, and Shapiro MD (2001). Direction of Gaze while Walking a Simple Route: Persons with Normal Vision and Persons with Retinitis Pigmentosa. *Optometry and Vision Science*, 78(9):667–675. [PubMed: 11587201]
- Turano KA, Yu D, Hao L, and Hicks JC (2005). Optic-flow and egocentric-direction strategies in walking: Central vs peripheral visual field. *Vision Research*, 45(25):3117–3132. [PubMed: 16084556]
- Vilkhu RS, Madugula SS, Grosberg LE, Gogliettino AR, Hottowy P, Dabrowski W, Sher A, Litke AM, Mitra S, and Chichilnisky EJ (2021). Spatially patterned bi-electrode epiretinal stimulation for axon avoidance at cellular resolution. *Journal of Neural Engineering*, 18(6):066007. Publisher: IOP Publishing.
- Wilke RGH, Moghadam GK, Lovell NH, Suaning GJ, and Dokos S (2011). Electric crosstalk impairs spatial resolution of multi-electrode arrays in retinal implants. *Journal of Neural Engineering*, 8(4):046016. [PubMed: 21673395]
- Williams MA, Galbraith C, Kane SK, and Hurst A (2014). "just let the cane hit it": how the blind and sighted see navigation differently. In Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility, ASSETS '14, pages 217–224, New York, NY, USA. Association for Computing Machinery.
- Wu Yuxin, Kirillov Alexander, Massa Francisco, Lo Wan-Yen, and Girshick Ross (2019). Detectron2.
- Zapf MP, Matteucci PB, Lovell NH, Zheng S, and Suaning GJ (2014). Towards photorealistic and immersive virtual-reality environments for simulated prosthetic vision: integrating recent breakthroughs in consumer hardware and software. In Annual International Conference of the

IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference, volume 2014, pages 2597–2600.

Zhang M, Krintz C, and Wolski R (2021). Edge-adaptable serverless acceleration for machine learning Internet of Things applications. *Software: Practice and Experience*, 51(9):1852–1867. [eprint: 10.1002/spe.2944](https://doi.org/10.1002/spe.2944).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Smart Bionic Eye Concept

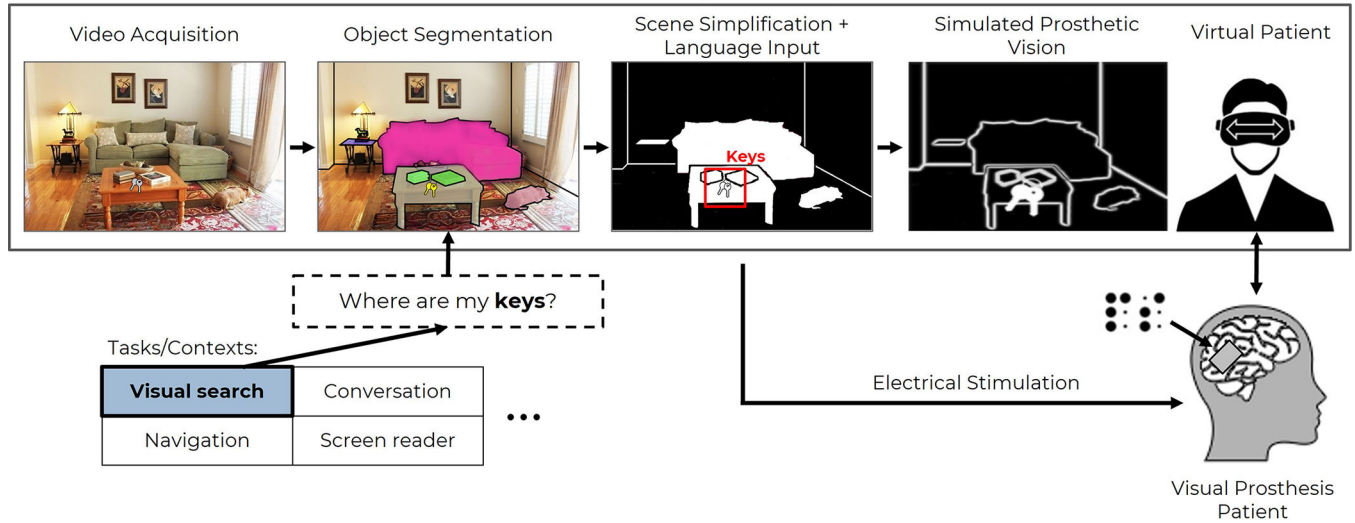


Figure 1. *Smart Bionic Eye.* A visual prosthesis has the potential to provide visual augmentations through the means of artificial intelligence (AI) based scene understanding (here shown for visual search). For example, a user may verbally instruct the *Smart Bionic Eye* to locate misplaced keys, and the system would respond visually by segmenting the keys in the prosthetic image while the user is looking around the room (room image reprinted under CC-BY from Lin et al., 2014). To guide the development of such a device, we propose to develop a virtual reality prototype supported by simulated prosthetic vision. Figure reprinted under CC-BY from <https://doi.org/10.6084/m9.figshare.20092640.v1>.

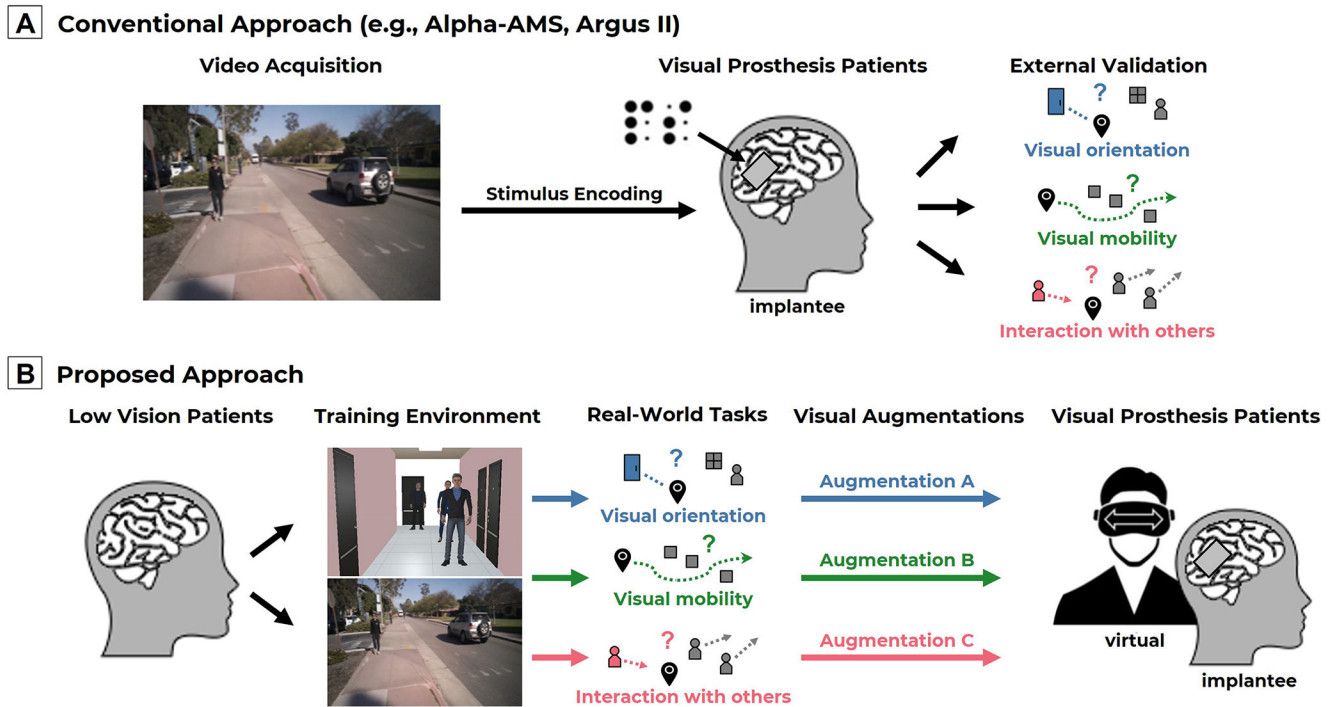


Figure 2. Stimulation strategies for visual prostheses that use an external camera. *A)* In the conventional approach implemented by previously commercialized devices such as Alpha-AMS and Argus II, a fixed and simple (e.g., linear) mapping is used to translate the grayscale value of a pixel in each video frame to a current amplitude of the corresponding electrode in the implant. The same encoding is used for all possible use cases. *B)* In the proposed approach, visual augmentation modes are task-dependent and informed by qualitative feedback as well as behavioral performance of virtual and real prosthesis patients on real-world tasks. The user is able to switch between modes on demand.




Figure 3.

Deep learning–based visual augmentations to support scene understanding. *A*) Segmenting objects of interest from background clutter using detectron2 (Wu, Yuxin et al., 2019). *B*) Substituting relative depth as sensed from single images for intensity using monodepth2 (Godard et al., 2019). *C*) Detecting structural edges of indoor environments (Sanchez-Garcia et al., 2020b). *D*) Visual question answering, where a deep neural network responds to “How many giraffes are drinking water?” visually by drawing bounding boxes around all giraffes by the water hole. (Antol et al., 2015).