Greedy Optimization of Electrode Arrangement for Epiretinal Prostheses

Ashley Bruce[1] and Michael Beyeler[1,2(✉)]

[1] Department of Computer Science, University of California, Santa Barbara, CA 93106, USA
{ashleybruce,mbeyeler}@ucsb.edu
[2] Department of Psychological and Brain Sciences, University of California, Santa Barbara, CA 93106, USA

Abstract. Visual neuroprostheses are the only FDA-approved technology for the treatment of retinal degenerative blindness. Although recent work has demonstrated a systematic relationship between electrode location and the shape of the elicited visual percept, this knowledge has yet to be incorporated into retinal prosthesis design, where electrodes are typically arranged on either a rectangular or hexagonal grid. Here we optimize the intraocular placement of epiretinal electrodes using dictionary learning. Importantly, the optimization process is informed by a previously established and psychophysically validated model of simulated prosthetic vision. We systematically evaluate three different electrode placement strategies across a wide range of possible phosphene shapes and recommend electrode arrangements that maximize visual subfield coverage. In the near future, our work may guide the prototyping of next-generation neuroprostheses.

AQ1

Keywords: Retinal prosthesis · Implant design · Dictionary selection

1 Introduction

Current visual neuroprostheses consist of a microelectrode array (MEA) implanted into the eye or brain that is used to electrically stimulate surviving cells in the visual system in an effort to elicit visual percepts ("phosphenes"). Current epiretinal implant users perceive highly distorted percepts, which vary in shape not just across subjects, but also across electrodes [7,10], and may be caused by incidental stimulation of passing nerve fiber bundles (NFBs) in the retina [4,13]. However, this knowledge has yet to be incorporated into prosthesis design.

To address this challenge, we make the following contributions:

1. We optimize the intraocular placement of epiretinal electrodes using dictionary learning. Importantly, this optimization process is informed by a previously established and psychophysically validated model of simulated prosthetic vision [1,4].

L. Wang (Eds.): MICCAI 2022, LNCS 13437, pp. 1–10, 2022.
https://doi.org/10.1007/978-3-031-16449-1_57

2. We systematically evaluate three different electrode placement strategies across a wide range of possible phosphene shapes and recommend electrode arrangements that maximize visual subfield coverage.

2 Related Work

Sensory neuroprostheses such as retinal and cochlear implants are emerging as a promising technology to restore lost sensory function. These devices bypass the natural sensory transduction mechanism and provide direct electrical stimulation of (retinal or auditory) nerve fibers that the brain interprets as a (visual or auditory) percept. However, perceptual distortions can occur when multiple electrodes stimulate the same neural pathways [2,15], and electrode placement has been shown to impact the quality of visual and hearing outcomes [4,8].

However, this information has yet to be incorporated into neuroprosthetic device design and surgical placement, with most patients having less-than-optimal MEA placement [3,6]. Whereas previous studies have optimized the shape of individual electrodes [12], which electrodes to activate in order to produce a desired visual response [5,9,14], or the overall implant placement [3], we

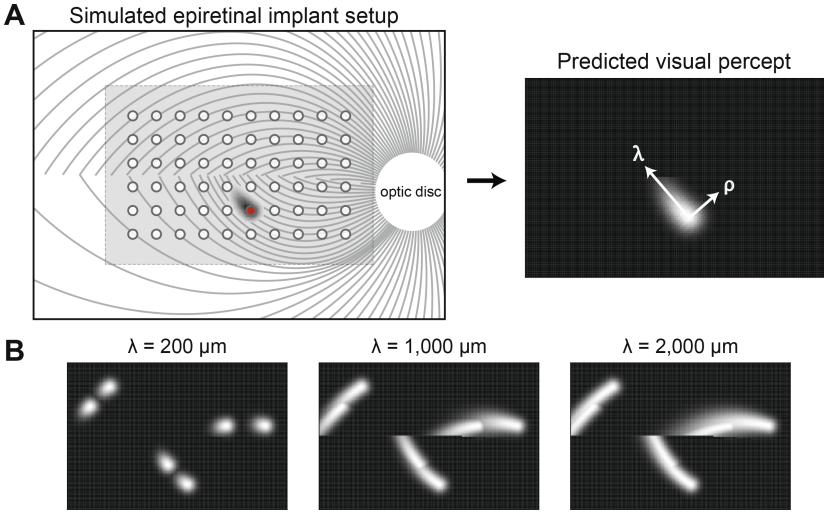

Fig. 1. A) Axon map model. *Left*: Electrical stimulation (red disc) of a NFB (gray lines) leads to tissue activation (dark-gray shaded region) elongated along the NFB trajectory away from the optic disc (white circle). The light-gray shaded region indicates the visual subfield that is being simulated. *Right*: The resulting visual percept appears elongated as well; its shape can be described by two parameters, λ (spatial extent along the NFB trajectory) and ρ (spatial extent perpendicular to the NFB). **B)** As λ increases, percepts become more elongated and start to overlap. (Color figure online)

are unaware of any studies that have attempted to optimize the placement of individual electrodes within an implant.

In the case of retinal implants, recent work has demonstrated that phosphene shape strongly depends on the retinal location of the stimulating electrode [4]. Because retinal ganglion cells (RGCs) send their axons on highly stereotyped pathways to the optic nerve, an electrode that stimulates nearby axonal fibers would be expected to antidromically activate RGC bodies located peripheral to the point of stimulation, leading to percepts that appear elongated in the direction of the underlying NFB trajectory (Fig. 1A, *right*). Using a simulated map of NFBs, Reference [4] was thus able to accurately predict phosphene shape for various users of the Argus Retinal Prosthesis System (Second Sight Medical Products, Inc.), by assuming that an axon's sensitivity to electrical stimulation:

i. decays exponentially with decay constant ρ as a function of distance from the stimulation site,
ii. decays exponentially with decay constant λ as a function of distance from the cell body, measured as axon path length.

As can be seen in Fig. 1B, electrodes near the horizontal meridian are predicted to elicit circular percepts, while other electrodes are predicted to produce elongated percepts that will differ in angle based on whether they fall above or below the horizontal meridian. In addition, the values of ρ and λ dictate the size and elongation of elicited phosphenes, respectively, which may drastically affect visual outcomes. Specifically, if two electrodes happen to activate the same NFB, they might not generate two distinct phosphenes (Fig. 1B, *right*).

Instead of arranging electrodes such that they efficiently tile the *retinal surface*, a better approach might thus be to arrange electrodes such that the *elicited percepts* effectively tile the *visual field*. In the following, we will demonstrate that this is equivalent to finding the smallest set of electrodes that cover a desired portion of the visual field (here termed a visual subfield). However, this strategy can be readily applied wherever there is a topological mapping from stimulus space to perceptual space (*e.g.*, visual, auditory, tactile stimulation). Incorporating this knowledge into implant design could therefore be indispensable to the success of future visual neuroprostheses.

3 Methods

3.1 Phosphene Model

Let $\mathcal{E} = \{e_1, \ldots, e_N\}$ be the set of N electrodes in a MEA, where the i-th electrode $e_i = (x_i, y_i, r_i)$ is described by its location on the retinal surface (x_i, y_i) and its radius $r_i > 0$. For example, Argus II can be described by $\mathcal{E}_{\text{ArgusII}}$, where $N = |\mathcal{E}_{\text{ArgusII}}| = 60$, $r_i = 122.5\,\mu\text{m}$ $\forall i$, and (x_i, y_i) are spaced $575\,\mu\text{m}$ apart on a rectangular grid. We do not assume any particular ordering of $\mathcal{E}$.

Furthermore, let $\mathcal{S} = \{(s_1, \ldots, s_k)\}$ be the set of stimuli sent to $k \leq N$ electrodes in the MEA, where the i-th stimulus $s_i = (e_i, a_i)$ is described by an

electrode $e_i \in \mathcal{E}$ and its corresponding activation function a_i. In practice, a_i may be a biphasic pulse train of a given duration, pulse amplitude, and pulse frequency. However, for the purpose of this study, we limited ourselves to spatial activation values $a_i \in \mathbb{R}_{\geq 0}$, which did not contain a temporal component.

A phosphene model $\mathcal{M}$ then takes a set of stimuli $\mathcal{S}$ as input and outputs a visual percept $p \in \mathbb{R}_{\geq 0}^{H \times W}$, which is a height ($H$) $\times$ width (W) grayscale image. In general, $\mathcal{M}$ is a nonlinear function of $\mathcal{S}$ and depends on subject-specific parameters θ_{subject} such as ρ and λ; thus $p = \mathcal{M}(\mathcal{S}; \theta_{\text{subject}})$.

For the purpose of this study, we used *pulse2percept* 0.8.0, a Python-based simulation framework for bionic vision [1] that provides an open-source implementation of the axon map model [4], as described in the previous section. Constrained by electrophysiological and psychophysical data, this model predicts what a bionic eye user should "see" for any given set of stimuli $\mathcal{S}$.

3.2 Dictionary Selection

Problem Formulation. Let $\mathcal{T}$ now be the set of all nonoverlapping epiretinal electrodes, and $\mathcal{D}$ be a subset of those; that is, $\mathcal{D} \subset \mathcal{T}$, where $|\mathcal{D}| \ll |\mathcal{T}|$. In the dictionary selection problem, we are interested in finding the dictionary $\mathcal{D}^*$ that maximizes a utility function F (see next subsection):

$$\mathcal{D}^* = \underset{|\mathcal{D}| \leq k}{\arg\max}\, F(\mathcal{D}), \tag{1}$$

where k is a constraint on the number of electrodes that the dictionary can be composed of.

This optimization problem presents combinatorial challenges, as we have to find the set $\mathcal{D}^*$ out of exponentially many options in $\mathcal{T}$. However, we will only consider utility functions F with the following properties:

i. The empty set has zero utility; that is, $F(\emptyset) = 0$.
ii. F increases monotonically; that is, whenever $\mathcal{D} \subseteq \mathcal{D}'$, then $F(\mathcal{D}) \leq F(\mathcal{D}')$.
iii. F is approximately submodular; that is, there exists an ε such that whenever $\mathcal{D} \subseteq \mathcal{D}' \subseteq \mathcal{T}$ and an electrode $e \in \mathcal{T} \setminus \mathcal{D}'$, it holds that $F(\mathcal{D} \cup \{e\}) - F(\mathcal{D}) \geq F(\mathcal{D}' \cup \{e\}) - F(\mathcal{D}') - \varepsilon$. This property implies that adding a new electrode e to a larger dictionary $\mathcal{D}'$ helps at most ε more than adding e to a subset $\mathcal{D} \subseteq \mathcal{D}'$.

For utility functions with the above properties, Nemhauser et al. [11] proved that a simple greedy algorithm that starts with the empty set $\mathcal{D}_0 = \emptyset$, and at every iteration i adds the element

$$d_i = \underset{d \in \mathcal{T} \setminus \mathcal{D}}{\arg\max}\, F(\mathcal{D}_{i-1} \cup \{d\}), \tag{2}$$

where $\mathcal{D}_i = \{d_1, \ldots, d_i\}$, is able to obtain a near-optimal solution.

Utility Function. Ideally, the utility function F would directly assess the quality of the generated artificial vision. As a first step towards such a quality measure, we considered the ability of a set of electrodes to lead to phosphenes that cover a specific visual subfield (i.e., the gray shaded region in Fig. 1). We would thus activate every electrode in $\mathcal{D}$, represented by $\mathcal{S}_\mathcal{D}$, and calculate the percept $p_\mathcal{D} = \mathcal{M}(\mathcal{S}_\mathcal{D}; \theta_{\text{subject}})$. Then F was given as the visual subfield coverage; that is:

$$F(\mathcal{D}) = \sum_{w=1}^{w=W} \sum_{h=1}^{h=H} p_{hw} \geq \varepsilon, \tag{3}$$

where W and H were the width and height of the percept, respectively, $\varepsilon = 0.1 \max_{\forall h,w}(p_{hw})$, and $F \in [0, HW]$.

Dictionary Selection Strategies. To find $\mathcal{D}^*$, we considered three different strategies:

- **No Overlap:** At each iteration i, d_i was chosen according to Eq. 2.
- **Not Too Close:** To consider manufacturing constraints, we modified Eq. 2 above to enforce that all electrodes were placed at least $c = 112.5\,\mu\text{m}$ apart from each other:

$$d_i = \underset{d \in \mathcal{T} \setminus \mathcal{D} \text{ s.t. } ||d, d_j||_2 \geq c \; \forall d_j \in \mathcal{D} \setminus d}{\arg\max} F(\mathcal{D}_{i-1} \cup \{d\}). \tag{4}$$

- **Max Pairwise Distance:** As electrical crosstalk is one of the main causes of impaired spatial resolution in retinal implants [15], we further modified Eq. 4 to place electrodes as far away from each other as possible:

$$d_i = \underset{d \in \mathcal{T} \setminus \mathcal{D} \text{ s.t. } ||d, d_j||_2 \geq c \; \forall d_j \in \mathcal{D} \setminus d}{\arg\max} F(\mathcal{D}_{i-1} \cup \{d\}) + \alpha \sum_{d_i \in \mathcal{D} \setminus d} ||d, d_i||_2, \tag{5}$$

  where $\alpha = 1 \times 10^{-4}$ was a scaling factor.

Implementation Details. For the sake of feasibility, we limited $\mathcal{T}$ to electrodes placed on a finely spaced search grid, $x_i \in [-3000, 3000]\,\mu\text{m}$ and $y_i \in [-2000, 2000]\,\mu\text{m}$, sampled at $112.5\,\mu\text{m}$ (i.e., the radius of an Argus II electrode). This led to a manageable set size ($|\mathcal{T}| \approx 2000$) while still allowing electrodes to be placed right next to each other (if desirable). To find the electrode d_i at each iteration in Eqs. 2, 4, and 5 above, we thus performed a grid search.

Stopping Criteria. The dictionary search was stopped when at least one of the following criteria were met:

- visual subfield coverage reached 99 %,
- the utility score F did not improve by $\geq 1 \times 10^{-6}$ on two consecutive runs,
- no more viable electrode locations were available (i.e., $\mathcal{T} \setminus \mathcal{D}$ s.t. $||d, d_i||_2 \geq c \; \forall d_i \in \mathcal{D} \setminus d = \emptyset$).

4 Results

4.1 Visual Subfield Coverage

The results of the greedy dictionary selection are shown in Fig. 2. For all three dictionary selection strategies, the number of electrodes required to cover at least 99 % of the visual subfield was inversely proportional to ρ and λ. As expected, the largest number was achieved with the smallest, most compact phosphene shape ($\rho = 100\,\mu\text{m}, \lambda = 200\,\mu\text{m}$). The required electrode number dropped rapidly with increasing ρ and λ, indicating that for large phosphenes, a prototype implant such as Argus I (4×4 electrodes) might be sufficient to cover the whole subfield.

It is interesting to note that, for any given ρ and λ combination, the Max Pairwise Distance strategy required a smaller number of electrodes than the No Overlap strategy. However, using the Not Too Close strategy, smaller ρ and λ combinations were no longer able to reach full coverage, as electrodes could no longer be placed too close to each other or to the visual subfield boundary. This issue was amplified with the Max Pairwise Distance strategy, with which coverage dropped for most ρ and λ combinations to 95%.

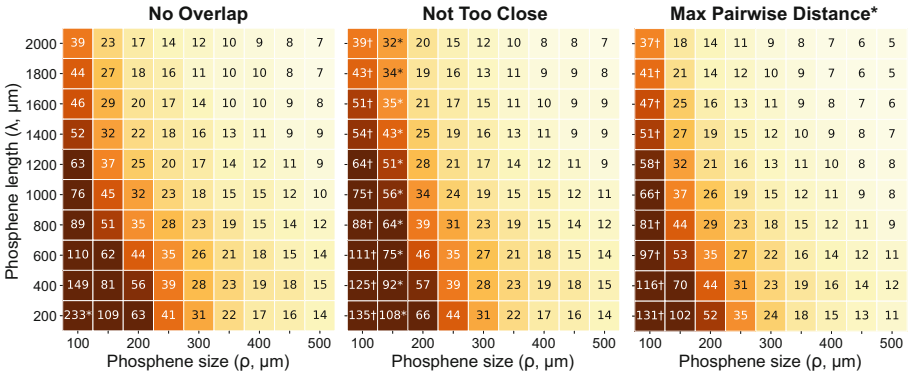

Fig. 2. Number of electrodes needed to reach maximum visual subfield coverage for the three dictionary selection strategies (†: 80−95% coverage, ∗: 95−99% coverage). Max Pairwise Distance always had ∗, unless otherwise noted by †.

With larger ρ and λ values, an implant required less than ten electrodes to cover the visual subfield. However, such a small number was not necessarily desirable, as it also reduced the number of distinct phosphenes that the implant can produce. Rather than focusing on visual subfield coverage alone, one might therefore ask what kinds of electrode arrangements the three dictionary selection strategies yield and what the resulting percepts look like.

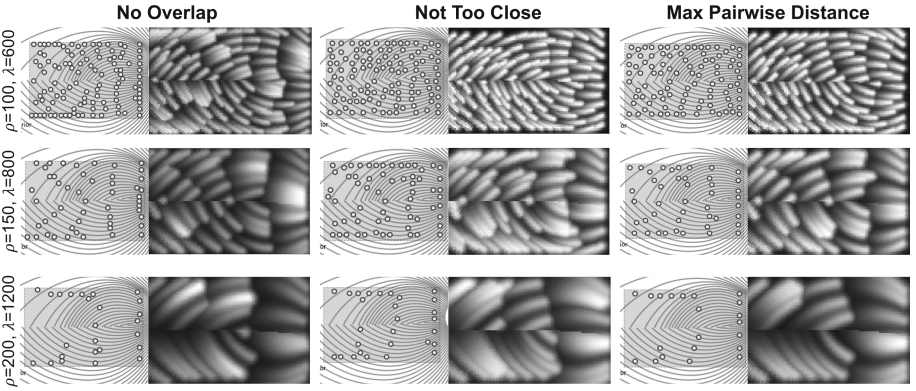

Fig. 3. Representative examples of the final electrode arrangements generated with the three dictionary selection strategies for different ρ and λ combinations. The left half of each panel shows the retinal location of all electrodes (small circles) in the implant, and the shaded region indicates the visual subfield (compare to Fig. 1). The visual percept that results from simultaneously activating every electrode in the implant is shown in the right half of each panel.

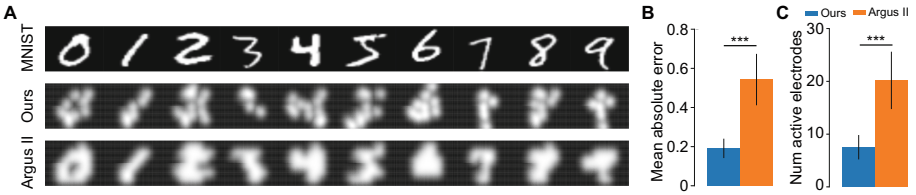

Fig. 4. A) Representative samples from the MNIST database (*top row*) represented with an optimized 60-electrode implant ("No Overlap", *middle row*) and Argus II (*bottom row*). **B)** Mean absolute error for 100 randomly selected MNIST digits (vertical bars: standard error of the mean, ***: $p < .001$). **C)** Number of active electrodes needed to represent the 100 digits. All simulations had $\rho = 200\,\mu\text{m}, \lambda = 400\,\mu\text{m}$.

4.2 Electrode Arrangement

Example electrode arrangements suggested by the three dictionary selection strategies are shown in Fig. 3. Here it is evident that, as λ increased, electrodes were preferentially placed on the top, bottom, and right boundaries of the visual subfield. This placement would often lead to the longest streaks, thus yielding the largest coverage. For small ρ values, electrodes aggregated mainly on NFBs, spaced λ apart, so that the streaks generated by different electrodes tiled the visual subfield. As ρ increased, electrodes tended to migrate away from the border and more inward, due to the outward spread of the generated percept.

4.3 Comparison with Argus II

To assess whether the optimized electrode arrangements provided an improvement over a rectangular MEA, we compared our results to Argus II. We thus modified our experiment such that the dictionary consisted of at most 60 electrodes (i.e., same as Argus II). We then considered the ability of the optimized implant to represent handwritten digits from the MNIST database (Fig. 4A). We found that the optimized implant not only led to smaller errors between predicted percepts and ground-truth digits (paired t-test, $p < .001$; Fig. 4B) but also required less active electrodes overall (paired t-test, $p < .001$; Fig. 4C).

Furthermore, we calculated the visual subfield coverage achieved with these 60 electrodes and compared the result to Argus II (Table 1). Given the fixed number of electrodes, our electrode selection strategies were always able to cover a larger portion of the visual subfield than Argus II.

Overall these results suggest that a rectangular grid might not be the best electrode arrangement for epiretinal prostheses.

Table 1. Visual subfield coverage achieved with 60 electrodes using different electrode placement strategies. AII: Argus II. NTC: Not Too Close. MPD: Max Pairwise Distance. Asterisk (*): Search terminated before utilizing 60 electrodes.

ρ	100 µm			150 µm			200 µm		
λ	AII	NTC	MPD	AII	NTC	MPD	AII	NTC	MPD
200 µm	39.2	43.2	42.9	74.1	75.6	75.9	86.2	>99*	95.2
400 µm	55.0	66.8	66.8	80.6	91.3	91.4	87.5	>99*	95.1
600 µm	62.4	81.8	81.3	81.9	96.9	95.3	88.4	>99*	95.2
800 µm	66.7	88.8	88.8	82.8	97.7	95.2	89.0	>99*	95.4
1000 µm	69.4	92.5	92.2	83.5	98.1*	95.2	89.6	>99*	95.8
1200 µm	71.9	92.6	92.5	84.2	98.0*	95.3	90.1	>99*	95.7
1400 µm	72.4	93.9*	93.8	84.6	98.1*	95.27	90.4	>99*	95.6
1600 µm	73.4	93.9*	93.8	85.0	98.5*	95.41	90.7	>99*	95.8
1800 µm	74.1	94.0*	93.9	85.4	98.6*	95.16	91.0	>99*	95.7
2000 µm	74.7	94.4*	94.4	85.9	98.3*	95.08	91.4	>99*	95.2

5 Conclusion

We report epiretinal electrode arrangements that maximize visual subfield coverage as discovered by dictionary learning. We were able to obtain nearly full coverage across different ρ and λ values for a given implant size and electrode radius, and report the number of electrodes necessary to reach this coverage. Future work should extend the dictionary selection strategy to other implant

sizes and explore all possible stimuli, not just the ones that may result in sufficient visual subfield coverage.

This preliminary study is a first step towards the use of computer simulations in the prototyping of novel neuroprostheses. To the best of our knowledge, this is the first study using a psychophysically validated phosphene model to optimize electrode arrangement. Even though we have focused on a specific implant technology, our strategy can be readily applied wherever there is a topological mapping from stimulus space to perceptual space (e.g., visual, auditory, tactile). This means that our approach could be extended to other neuromodulation technologies that include (but are not limited to) other electronic prostheses and optogenetic technologies. In the near future, our work may therefore guide the prototyping of next-generation epiretinal prostheses.

Acknowledgements. This work was supported by the National Institutes of Health (NIH R00 EY-029329 to MB). The authors would like to thank Madori Spiker and Alex Rasla for valuable contributions to an earlier version of the project.

References

1. Beyeler, M., Boynton, G.M., Fine, I., Rokem, A.: Pulse2percept: a Python-based simulation framework for bionic vision. In: Huff, K., Lippa, D., Niederhut, D., Pacer, M. (eds.) Proceedings of the 16th Science in Python Conference, pp. 81–88 (2017). https://doi.org/10.25080/shinma-7f4c6e7-00c
2. Beyeler, M., Rokem, A., Boynton, G.M., Fine, I.: Learning to see again: biological constraints on cortical plasticity and the implications for sight restoration technologies. J. Neural Eng. **14**(5), 051003 (2017). https://doi.org/10.1088/1741-2552/aa795e
3. Beyeler, M., Boynton, G.M., Fine, I., Rokem, A.: Model-based recommendations for optimal surgical placement of epiretinal implants. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11768, pp. 394–402. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32254-0_44
4. Beyeler, M., Nanduri, D., Weiland, J.D., Rokem, A., Boynton, G.M., Fine, I.: A model of ganglion axon pathways accounts for percepts elicited by retinal implants. Sci. Rep. **9**(1), 1–16 (2019). https://doi.org/10.1038/s41598-019-45416-4
5. Bratu, E., Dwyer, R., Noble, J.: A graph-based method for optimal active electrode selection in cochlear implants. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12263, pp. 34–43. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59716-0_4
6. Chakravorti, S., et al.: Further evidence of the relationship between cochlear implant electrode positioning and hearing outcomes. Otol. Neurotol. **40**(5), 617–624 (2019). https://doi.org/10.1097/MAO.0000000000002204
7. Erickson-Davis, C., Korzybska, H.: What do blind people "see" with retinal prostheses? Observations and qualitative reports of epiretinal implant users. PLOS ONE **16**(2), e0229189 (2021). https://doi.org/10.1371/journal.pone.0229189. publisher: Public Library of Science
8. Finley, C.C., et al.: Role of electrode placement as a contributor to variability in cochlear implant outcomes. Otol. Neurotol. **29**(7), 920–928 (2008). https://doi.org/10.1097/MAO.0b013e318184f492

9. Granley, J., Relic, L., Beyeler, M.: A hybrid neural autoencoder for sensory Neuroprostheses and Its Applications in Bionic Vision (2022). http://arxiv.org/abs/2205.13623,arXiv:2205.13623 [cs]
10. Luo, Y.H.L., Zhong, J.J., Clemo, M., da Cruz, L.: Long-term repeatability and reproducibility of phosphene characteristics in chronically implanted Argus II retinal prosthesis subjects. Am. J. Ophthalmol. **170**, 100–109 (2016). https://doi.org/10.1016/j.ajo.2016.07.021
11. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An analysis of approximations for maximizing submodular set functions-I. Math. Program. **14**(1), 265–294 (1978). https://doi.org/10.1007/BF01588971
12. Rattay, F., Resatz, S.: Effective electrode configuration for selective stimulation with inner eye prostheses. IEEE Trans. Biomed. Eng. **51**(9), 1659–1664 (2004). https://doi.org/10.1109/TBME.2004.828044
13. Rizzo, J.F., Wyatt, J., Loewenstein, J., Kelly, S., Shire, D.: Perceptual efficacy of electrical stimulation of human retina with a microelectrode array during short-term surgical trials. Invest. Ophthalmol. Vis. Sci. **44**(12), 5362–5369 (2003). https://doi.org/10.1167/iovs.02-0817
14. Shah, N.P., et al.: Optimization of electrical stimulation for a high-fidelity artificial retina. In: 2019 9th International IEEE/EMBS Conference on Neural Engineering (NER), pp. 714–718 (2019). https://doi.org/10.1109/NER.2019.8716987, ISSN: 1948-3546
15. Wilke, R.G.H., Moghadam, G.K., Lovell, N.H., Suaning, G.J., Dokos, S.: Electric crosstalk impairs spatial resolution of multi-electrode arrays in retinal implants. J. Neural Eng. **8**(4), 046016 (2011). https://doi.org/10.1088/1741-2560/8/4/046016