

Optimization of Electrical Stimulation for a High-Fidelity Artificial Retina

Nishal P. Shah¹, Sasidhar Madugula¹, Lauren Grosberg¹, Gonzalo Mena², Pulkit Tandon¹, Pawel Hottowy³, Alexander Sher⁴, Alan Litke⁴, Subhasish Mitra¹ and E.J. Chichilnisky¹

Abstract—Retinal prostheses aim to restore visual perception in patients blinded by photoreceptor degeneration, by stimulating surviving retinal ganglion cells (RGCs), causing them to send artificial visual signals to the brain. Present-day devices produce limited vision, in part due to indiscriminate and simultaneous activation of many RGCs of different types that normally signal asynchronously. To improve artificial vision, we propose a closed-loop, cellular-resolution device that automatically identifies the types and properties of nearby RGCs, calibrates its stimulation to produce a dictionary of achievable RGC activity patterns, and then uses this dictionary to optimize stimulation patterns based on the incoming visual image. To test this concept, we use a high-density multi-electrode array as a lab prototype, and deliver a rapid sequence of electrical stimuli from the dictionary, progressively assembling a visual image within the visual integration time of the brain. Greedily minimizing the error between the visual stimulus and a linear reconstruction (as a surrogate for perception) yields a real-time algorithm with an efficiency of 96% relative to optimum. This framework also provides insights for developing efficient hardware. For example, using only the most effective 50% of electrodes minimally affects performance, suggesting that an adaptive device configured using measured properties of the patient’s retina may permit efficiency with accuracy.

I. INTRODUCTION

Future neural interfaces will be used to treat a wide range of incurable diseases, from sensory and motor degeneration to psychiatric disorders. A prominent present-day example is a retinal prosthesis, which aims to stimulate the retinal ganglion cells (RGCs) to restore visual perception in patients blinded by photoreceptor degeneration (e.g macular degeneration or retinitis pigmentosa). Even though retinal prostheses have been successfully implanted in patients, the elicited perception has limited fidelity [1], [2], [3]. To improve perception, multiple technological advances are being pursued, including higher density electrode arrays [4], [5].

However, most of the technological advances being pursued do not address a fundamental problem. To produce high fidelity vision, future devices may need to accurately reproduce the diverse and specific firing pattern of 20 different RGC types [6] that asynchronously convey distinct visual features to the brain. When a single stimulation electrode simultaneously activates multiple nearby RGCs irrespective of their type, reproducing the asynchronous neural code is impossible [7]. To address this problem, it will be necessary for a future high resolution device to operate in a closed-loop manner - identifying the RGCs of different types in a

patient’s retina, measuring their responses to electrical stimulation, and using this information to optimize the artificial visual signal. Even though prior studies have used closed-loop experimental methods to optimize response selectivity [8] or perception in existing low resolution devices [9], it remains unclear how a future closed-loop device would function in real time to precisely reproduce the retinal code at cellular resolution.

Here, we present an approach to this problem using a laboratory prototype of a future high-resolution, closed-loop device. The lab prototype consists of a multi-electrode array preparation for stimulating and recording from isolated primate retina. Our solution for faithfully reproducing natural visual sensations relies on three main pillars: a) minimizing the error between expected perception and the target visual stimulus, rather than attempting to match evoked spikes exactly; b) developing and exploiting a calibrated *dictionary* of achievable RGC response patterns, and c) exploiting temporal dithering – sequential stimulation at a rate faster than the integration time of visual perception – to accumulate an optimal visual image without having to calibrate the device for all spatial patterns of stimulation.

In this paper, we use experimental data to show that the approach significantly improves the predicted perception of a visual image. Subsequently, we show that the algorithmic framework can be used to guide the design of more efficient hardware and electrical stimulation patterns in a future high-resolution device.

II. GREEDY DICTIONARY BASED STIMULATION

In this section, we formalize the problem and present a simple, real-time algorithm to optimize stimulation.

A. Functional components of the prosthesis

In a healthy retina, light from a visual stimulus is encoded in neural responses, which the brain reconstructs to perceive the outside world (Figure 1, top row). In a blind patient, we aim to replace the function of degraded neural circuitry by implanting a chip that encodes the visual stimulus in electrical stimulation patterns that directly elicit responses in RGCs (Figure 1, bottom row). Our goal is to provide the patterns of stimulation that produce the highest-fidelity artificial vision. This is accomplished in three steps. First, a model is developed describing how the patient perceives the artificial image, assuming that the brain optimally exploits the signals transmitted by the retina. Second, electrical stimulation and recording are used to develop a calibrated dictionary of RGC

¹ Stanford University, ² Harvard University, ³ AGH University of Science and Technology, ⁴ University of California, Santa Cruz

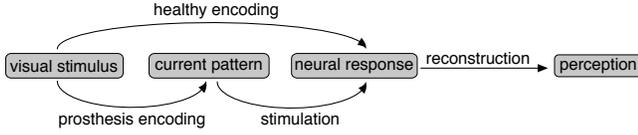


Fig. 1. **Functional components of a prosthesis.** The diagram indicates the steps of visual encoding in a healthy retina and with a retinal prosthesis, and the postulated reconstruction that gives rise to perception.

activity patterns that are possible with the specific implanted device. Finally, after calibration, the device operates in run-time by processing the incoming visual image and selecting a sequence of calibrated stimulation patterns that produces the most accurate inferred perception, using greedy optimization.

B. Notation

- S : Target visual stimulus (dimensions: $\# \text{ pixels} \times 1$).
- R : Neural responses (dimensions: $\# \text{ cells} \times 1$).
- A : Reconstruction matrix (dimensions : $\# \text{ pixels} \times \# \text{ cells}$). Columns of A correspond to the reconstruction filters of different cells.
- $\mathcal{D} = \{(e_d, a_d, p_d)\}_{d=1}^{d=|\mathcal{D}|}$: Stimulation dictionary, each element d consists of stimulation electrodes (e_d), current amplitudes (a_d) and response probability vector ($p_d \in [0, 1]^{\# \text{ cells} \times 1}$). Response for cell i , $R_d^i \sim \text{Bernoulli}(p_d^i)$ when current pattern d is stimulated.

C. Assumptions

- S is a static frame of the visual stimulus. We assume that a dynamic stimulus can be approximated with a sequence of static frames.
- A sequence of retinal responses, produced in rapid succession by electrical stimulation, is summed by the brain to produce a single perceived image. Hence, the reconstructed stimulus for the sequence R_1, \dots, R_T is given by $R = \sum_{t=1}^{t=T} R_t$ (Figure 2).
- Responses generated by a stimulation pattern, d_t , are independent of the responses to the previous stimulation patterns, d_1, \dots, d_{t-1} , as long as no stimulus is delivered to a cell during its spiking refractory period.
- The difference between the visual stimulus and achieved perception to be minimized is the mean-squared error (MSE): $\rho(S, \hat{S}) = \|S - \hat{S}\|_2^2$
- The perceived visual image is approximated by linearly reconstructing neural responses : $\hat{S} = AR$, with A an optimal linear filter.

In a later section, we discuss the limitations of these assumptions, and how to overcome them.

D. Greedy algorithm

The problem reduces to choosing a sequence of dictionary elements d_1, \dots, d_T to minimize the expected error in the perceived visual image ρ as follows:

$$\mathbb{E}\rho(S, \hat{S}) = \mathbb{E}\|S - AR\|_2^2 \quad (1)$$

$$= \mathbb{E}\|S - A(R_1 + \dots + R_T)\|_2^2 \quad (2)$$

A simple algorithm for minimizing ρ in a dynamic and unpredictable sequence of incoming images is to choose the dictionary elements greedily, maximizing the reduction in error at each time step. The stimulation choice at time step t is given by:

$$d_t = \underset{d \in \mathcal{D}}{\text{argmin}} \mathbb{E} \|S - A(\sum_{l=1}^{l=t-1} R_l + R_d)\|_2^2 \quad (3)$$

where the expectation is over R_1, \dots, R_{t-1} , the responses elicited in previous time steps (with $R_l \sim \text{Bernoulli}(p_{d_l}) \forall l \in \{0, \dots, t-1\}$).

For the choices of linear reconstruction and mean-squared error, the greedy algorithm can be implemented efficiently by decomposing the objective into bias and variance components (Equation 5). Because the responses at each time step are assumed to be independent of preceding responses, the total variance decomposes as the sum of the variance in all time steps (Equation 6):

$$\mathbb{E} \|S - A(\sum_{l=1}^{l=t-1} R_l + R_d)\|_2^2 \quad (4)$$

$$= \|S - \mathbb{E}[A(\sum_{l=1}^{l=t-1} R_l + R_d)]\|_2^2 +$$

$$\sum_{i \in \text{pixels}} \text{var}(A_i(\sum_{l=1}^{l=t-1} R_l + R_d)) \quad (5)$$

$$= \|S - (\tilde{S}_{t-1} + Ap_d)\|_2^2 + v_d + \tilde{V}_{t-1} \quad (6)$$

where A_i is the i th row of A ; Ap_{d_t} and v_d are the contributions to mean and variance in perception for the d th dictionary element; $\tilde{S}_t (= \tilde{S}_{t-1} + Ap_{d_t})$ and $\tilde{V}_t (= \tilde{V}_{t-1} + v_{d_t})$ are the cumulative mean and variance in perception due to stimulation patterns chosen in time steps $\{1, \dots, t\}$. To satisfy the independence assumption, it is necessary to avoid stimulating a cell during its refractory period. To accomplish this, the dictionary in Equation (3) was restricted to stimulation patterns that do not activate cells that were already targeted (firing probability greater than 0.1) in the preceding 5ms.

III. APPLICATION TO RECORDED DATA

A. Data

We used our lab prototype to evaluate the stimulus reconstruction associated with the greedy stimulation algorithm. A high density multi-electrode array (512-channels, $60 \mu\text{m}$ pitch, 20 kHz) was used to stimulate and record from several hundred RGCs over a $\sim 2\text{mm} \times 2\text{mm}$ region of isolated primate retina [10]. To calibrate response properties for the lab prototype, we use the fact that we can identify RGC types using visually evoked responses in our recordings from healthy retina. Responses to a 30 min white noise visual stimulus (80×40 pixel grid, refresh rate 120 Hz) were used to identify the cells, their types, and their locations [11], [12]. We restrict the subsequent analysis to two numerically dominant RGC types - ON parasol (116 cells) and

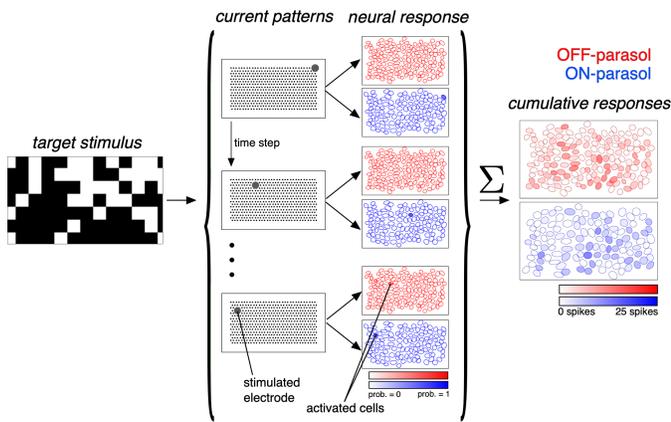


Fig. 2. **Fast, sequential stimulation to reproduce the neural code of the retina.** Temporal dithering is achieved by exploiting the assumed slow integration time of brain: the inferred visual image is assumed to depend only on the total number of spikes in each cell generated within a temporal integration window. Based on the visual stimulus, a rapid sequence of current patterns is delivered, with the time between successive stimuli much smaller than this temporal integration window).

OFF parasol (150 cells) cells, which are sampled efficiently in these experiments and formed nearly complete mosaics covering the region recorded. The stimulus reconstruction filter was first estimated from the continuous white noise stimulus, followed by a scaling based on the response to flashes of static white noise frames, averaged across different recordings. Note that although we identified the cell types and reconstruction filters from light-evoked responses for convenience, we have shown previously that for the future clinical application, the distinct cell types can be identified from spontaneous electrical activity [13].

The dictionary was estimated by stimulating each of the 512 electrodes individually, with 40 different current amplitudes (250 nA - 4 μ A), 27 times each. Response probabilities for each stimulation pattern were identified after removing electrical artifacts using custom spike sorting software [14]. Dictionary elements that stimulated cells along their axons, with somas off the electrode array, were removed due to their unknown receptive field locations and thus uncertain contribution to stimulus reconstruction [7]. Finally, only dictionary elements that activated at least one cell with probability ≥ 0.01 were retained, resulting in 4,233 dictionary elements. A single dictionary element that does not activate any cell ($p_d = 0$) was added to allow the greedy algorithm to avoid stimulation when any real stimulation pattern would increase error.

B. Algorithm progressively captures perception

The stimulation sequence chosen by the greedy algorithm progressively reconstructed the visual stimulus. This was seen by sampling responses according to the RGC response probabilities associated with the selected dictionary elements, and then reconstructing the stimulus linearly from these responses. The algorithm stimulated a spatially diverse sequence of electrodes, and the reconstructed stimulus converged toward a smoothed version of the target (Figure

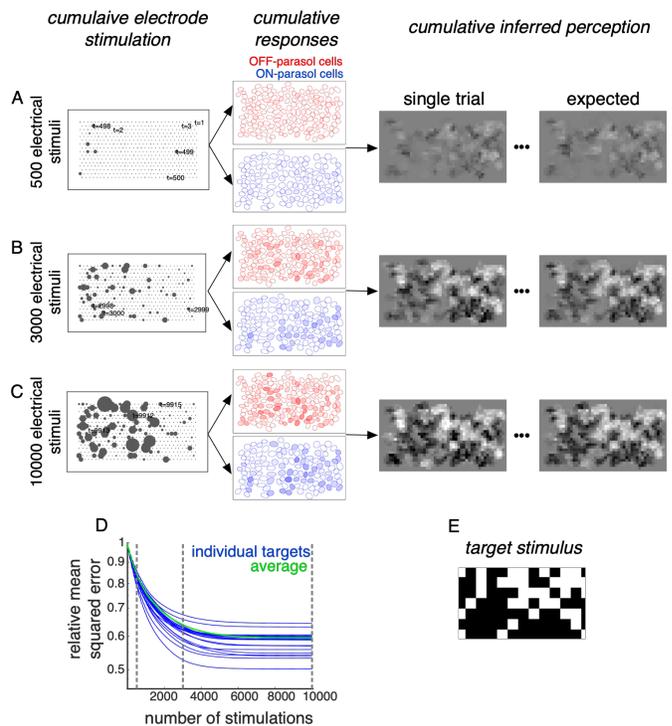


Fig. 3. **Stimulus reconstruction achieved using the greedy algorithm.** (A-C) Responses to electrical stimulation patterns, accumulated over time, produces inferred perception (stochastic; single trial and expected value shown). (D) The difference between target and reconstructed stimulus declines as a function of the number of stimulation patterns. (E) The target visual stimulus which the electrical stimulation is intended to reproduce.

3A). Since responses were generated stochastically according to the measured firing probabilities, variability across trials is expected in the reconstructed stimulus. However, the reconstruction in a single trial was similar to the expected reconstruction across trials (Figure 3(A), right two columns). Across 20 distinct visual targets, the error monotonically decreased with the number of stimulation patterns delivered, but saturated after $\sim 4,000$, because the reduction in bias was not sufficient to counter the increase in variance .

Note that the 4,000 stimulation patterns required for asymptotic performance must be delivered within a single integration time window, i.e. the time over which the brain integrates RGC inputs, in order to produce a single perceived image. For the 10 kHz stimulation frequency used here, this nominally requires a 400 ms integration time, much longer than likely integration times in the brain. Although the true integration time of visual circuitry of the brain is not known, a more realistic estimate would be tens of ms, as suggested by sampling rates appropriate for psychophysical experiments [15] and by the refresh rates used in cinematography [16]. To achieve 4,000 stimulation patterns in this shorter time would therefore require presenting several pulses during the same time step, and/or increasing the sampling frequency of the device. Note that the latter change would also influence the exclusion of stimulation patterns during the refractory period.

C. Near optimal perception achieved by greedy algorithm

In this section, we test whether the achieved reconstruction was limited by either a) the greedy nature of the algorithm, or b) the limited dictionary of single-electrode stimulation patterns. To answer this question, the error between achieved perception and the target stimulus was compared to two different lower bounds:

- 1) **Lower bound with any algorithm and fixed dictionary:** Let $w \in \mathbb{Z}_+^{|\mathcal{D}|\times 1}$ be the total number of times each dictionary element is chosen during the course of stimulation and let $D \in [0, 1]^{\#\text{cells}\times|\mathcal{D}|}$ be the dictionary matrix, with d th column equal to the probability vector p_d . The stimulus reconstruction $A \sum R_t$ has mean ADw and variance $V^T w$, where $V \in \mathbb{R}_+^{|\mathcal{D}|\times 1}$ is a vector of variances of dictionary elements. Minimizing the expected mean-squared error between the target and estimated perception with the positive integer constraint on w is difficult. However, a lower bound can be obtained by relaxing the integer spike count constraint and solving the following convex optimization problem:

$$\text{minimize}_{w \geq 0} [\|S - ADw\|_2^2 + V^T w] \quad (7)$$

- 2) **Lower bound of error with any algorithm and any dictionary:** Let $q \in \mathbb{Z}_+^{\#\text{cells}\times 1}$ denote the total number of spikes for each cell during the course of stimulation ($q = \sum_t R_t$). Assume that we have a "perfect" dictionary, such that an arbitrary number of spikes can be evoked in all cells using the entries in the dictionary. Then, a lower bound is estimated by solving the following convex optimization problem:

$$\text{minimize}_{q \geq 0} [\|S - Aq\|_2^2] \quad (8)$$

The stimulus reconstruction achieved by the greedy algorithm was very similar to that achieved by the optimal algorithm (Figure 4B, C for an example target, Figure 4E for summary across retinas), capturing $\sim 50\%$ and $\sim 48\%$ of the target stimulus variance respectively. Moreover, the error was $\sim 35\%$ when an ideal stimulation dictionary was available (Figure 4 (D, E)). In both cases, the upper bound derived by discretizing the optimal solution did not change these observations. Hence, the performance of the greedy algorithm was very close that of an optimal algorithm for the single electrode stimulation dictionary, and the improvement in performance requires a better stimulation dictionary.

D. Guiding hardware design

The results indicate that efficient hardware can be designed by limiting stimulation to the most frequently stimulated electrodes. Across 20 different target images, the distribution of stimulated electrodes was spatially non-uniform. More frequently chosen electrodes had larger number of axons nearby (Figure 5(A)). These electrodes typically stimulated multiple cells simultaneously, indicating that the algorithm exploits non-selective stimulation patterns. This suggests

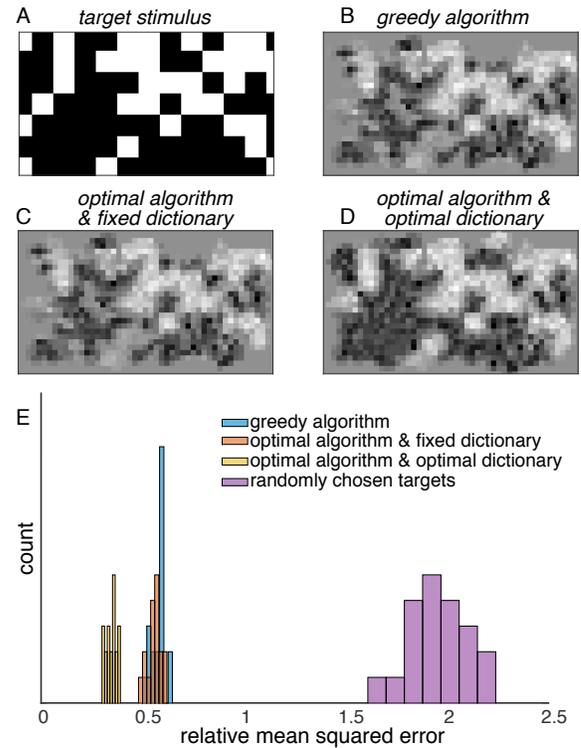


Fig. 4. **Quantitative evaluation of greedy algorithm.** (A) Target visual stimulus. (B) Expected perception (stimulus reconstruction) using greedy stimulation. (C) Reconstruction using an optimal algorithm and the fixed dictionary. (D) Reconstruction using an optimal algorithm and optimal dictionary. (E) Histogram of relative mean squared error (normalized by magnitude of target) across 20 different examples.

that previous approaches to optimal stimulation based on maximizing selectivity are not always the most effective [17].

To understand if this finding can guide efficient hardware design, the dictionary was restricted to the most frequently chosen electrodes. The greedy algorithm was applied with the restricted dictionary, and reconstruction error was averaged across 20 random targets. A minimal increase in error was observed even after reducing the number of available electrodes by 50% (Figure 5 (B) individual targets and Figure 5 (C, D, E, F) for an example target image). On reducing the number of available electrodes further, the error increased gradually. This increase was not due to the greedy stimulation choices, as a lower bound for the best algorithm showed similar behavior (Figure 5(B), green curve). These results suggest that an efficient device could maintain a reduced dictionary in which $\sim 50\%$ of the stimulating units are turned off, reducing memory access and static power. Note that all stimulating units are required during the initial calibration phase to select the most frequently used subset of electrodes - similar to previously developed methods in cochlear implants [18].

IV. CONCLUSIONS

We present a novel framework for optimizing cellular-resolution electrical stimulation in a closed-loop visual prosthesis. This framework enables adaptation of the device

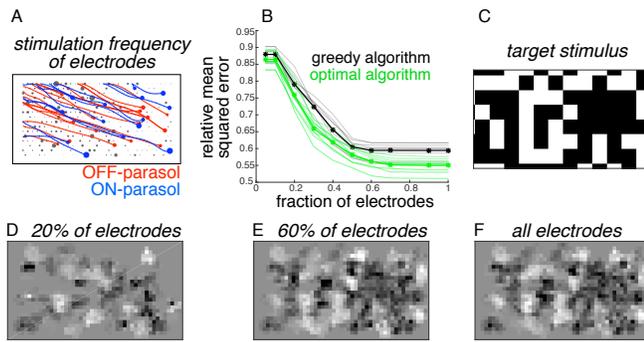


Fig. 5. **Subsampling electrodes with minimal effect on stimulus reconstruction.** (A) Frequency of stimulating different electrodes (size of circles), overlaid with axons (lines) and somas (colored circles) inferred from electrophysiological imaging. (B) Reconstruction error as a function of the fraction of electrodes included in the dictionary for multiple targets (black, thin lines) and average over 20 targets (black, thick line). Lower bound on error of any algorithm for the subsampled dictionaries for individual targets (green, thin lines) and averaged across targets (green, thick line). (C) Target stimulus. (D-F) Reconstruction using the dictionary with most frequently used 20%, 60% and 100% of electrodes, respectively.

encoding to a patient’s measured electrical (dictionary) and visual (reconstruction) response properties.

For a linear stimulus reconstruction and a single electrode stimulation dictionary, the greedy algorithm was surprisingly close to optimal. Further work will be needed to assess the quality of the elicited image with more accurate stimulation algorithms, nonlinear reconstruction methods, perceptual similarity measures, and dictionaries with multi-electrode stimulation patterns. Also, the temporal integration of retinal signals in the brain, and its impact on stimulation frequency, requires further investigation.

The modular nature of the greedy algorithm may help with design decisions for a future advanced device. For example, the impact of different electrical stimulation patterns such as current steering [17], [19], [20] and local return [21], [22] can be directly compared by evaluating the quality of the final inferred visual image.

Finally, applying the algorithm to experimental data can lead to concrete insights for building efficient hardware, such as exclusion of unimportant electrodes to save power. This approach can therefore be useful for designing a sensory neural interface, based on a rigorous experimental prototype.

REFERENCES

- [1] M. S. Humayun, J. D. Dorn, L. Da Cruz, G. Dagnelie, J.-A. Sahel, P. E. Stanga, A. V. Cideciyan, J. L. Duncan, D. Elliott, E. Filley *et al.*, “Interim results from the international trial of second sight’s visual prosthesis,” *Ophthalmology*, vol. 119, no. 4, pp. 779–788, 2012.
- [2] K. Stingl, K.-U. Bartz-Schmidt, F. Gekeler, A. Kusnyerik, H. Sachs, and E. Zrenner, “Functional outcome in subretinal electronic implants depends on foveal eccentricity,” *Investigative ophthalmology & visual science*, vol. 54, no. 12, pp. 7658–7665, 2013.
- [3] K. Stingl, K. U. Bartz-Schmidt, D. Besch, A. Braun, A. Bruckmann, F. Gekeler, U. Greppmaier, S. Hipp, G. Hördörfer, C. Kernstock *et al.*, “Artificial vision with wirelessly powered subretinal electronic implant alpha-ims,” *Proc. R. Soc. B*, vol. 280, no. 1757, p. 20130077, 2013.
- [4] H. Lorach, G. Goetz, R. Smith, X. Lei, Y. Mandel, T. Kamins, K. Mathieson, P. Huie, J. Harris, A. Sher *et al.*, “Photovoltaic restoration of sight with high visual acuity,” *Nature medicine*, vol. 21, no. 5, p. 476, 2015.

- [5] C. Sekirnjak, P. Hottowy, A. Sher, W. Dabrowski, A. M. Litke, and E. Chichilnisky, “High-resolution electrical stimulation of primate retina for epiretinal implant design,” *Journal of Neuroscience*, vol. 28, no. 17, pp. 4446–4456, 2008.
- [6] D. Dacey, “20 origins of perception: Retinal ganglion cell diversity and the creation of parallel visual pathways,” 2004.
- [7] L. E. Grosberg, K. Ganesan, G. A. Goetz, S. S. Madugula, N. Bhaskhar, V. Fan, P. Li, P. Hottowy, W. Dabrowski, A. Sher *et al.*, “Activation of ganglion cells and axon bundles using epiretinal electrical stimulation,” *Journal of neurophysiology*, vol. 118, no. 3, pp. 1457–1471, 2017.
- [8] T. Guo, C. Y. Yang, D. Tsai, M. Muralidharan, G. J. Suaning, J. W. Morley, S. Dokos, and N. H. Lovell, “Closed-loop efficient searching of optimal electrical stimulation parameters for preferential excitation of retinal ganglion cells,” *Frontiers in neuroscience*, vol. 12, p. 168, 2018.
- [9] M. Beyeler, D. Nanduri, J. D. Weiland, A. Rokem, G. M. Boynton, and I. Fine, “Optimizing stimulation protocols for prosthetic vision based on retinal anatomy,” *Annual Meeting of the Organization of Computational Neuroscience (CNS)*, 2018.
- [10] P. Hottowy, A. Skoczeń, D. E. Gunning, S. Kachiguine, K. Mathieson, A. Sher, P. Wiacek, A. M. Litke, and W. Dabrowski, “Properties and application of a multichannel integrated circuit for low-artifact, patterned electrical stimulation of neural tissue,” *Journal of neural engineering*, vol. 9, no. 6, p. 066005, 2012.
- [11] E. Chichilnisky and R. S. Kalmar, “Functional asymmetries in on and off ganglion cells of primate retina,” *Journal of Neuroscience*, vol. 22, no. 7, pp. 2737–2747, 2002.
- [12] E. Frchette, A. Sher, M. Grivich, D. Petrusca, A. Litke, and E. Chichilnisky, “Fidelity of the ensemble code for visual motion in primate retina,” *Journal of neurophysiology*, vol. 94, no. 1, pp. 119–135, 2005.
- [13] E. Richard, G. A. Goetz, and E. Chichilnisky, “Recognizing retinal ganglion cells in the dark,” in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Curran Associates, Inc., 2015, pp. 2476–2484. [Online]. Available: <http://papers.nips.cc/paper/5652-recognizing-retinal-ganglion-cells-in-the-dark.pdf>
- [14] G. E. Mena, L. E. Grosberg, S. Madugula, P. Hottowy, A. Litke, J. Cunningham, E. Chichilnisky, and L. Paninski, “Electrical stimulus artifact cancellation and neural spike detection on large multi-electrode arrays,” *PLoS computational biology*, vol. 13, no. 11, p. e1005842, 2017.
- [15] C. Rashbass, “The visibility of transient changes of luminance,” *The Journal of Physiology*, vol. 210, no. 1, pp. 165–186, 1970.
- [16] R. Efron, “Conservation of temporal information by perceptual systems,” *Perception & Psychophysics*, vol. 14, no. 3, pp. 518–530, 1973.
- [17] L. H. Jepson, P. Hottowy, K. Mathieson, D. E. Gunning, W. Dabrowski, A. M. Litke, and E. Chichilnisky, “Spatially patterned electrical stimulation to enhance resolution of retinal prostheses,” *Journal of Neuroscience*, vol. 34, no. 14, pp. 4871–4881, 2014.
- [18] H. J. McDermott, C. M. McKay, and A. E. Vandali, “A new portable sound processor for the university of melbourne/nucleus limited multi-electrode cochlear implant,” *The Journal of the Acoustical Society of America*, vol. 91, no. 6, pp. 3367–3371, 1992.
- [19] P. B. Matteucci, S. C. Chen, D. Tsai, C. W. Dodds, S. Dokos, J. W. Morley, N. H. Lovell, and G. J. Suaning, “Current steering in retinal stimulation via a quasimonopolar stimulation paradigm,” *Investigative ophthalmology & visual science*, vol. 54, no. 6, pp. 4307–4320, 2013.
- [20] G. Dumm, J. B. Fallon, C. E. Williams, and M. N. Shivasani, “Virtual electrodes by current steering in retinal prostheses,” *Investigative ophthalmology & visual science*, vol. 55, no. 12, pp. 8077–8085, 2014.
- [21] M. Abramian, N. H. Lovell, J. W. Morley, G. J. Suaning, and S. Dokos, “Activation of retinal ganglion cells following epiretinal electrical stimulation with hexagonally arranged bipolar electrodes,” *Journal of neural engineering*, vol. 8, no. 3, p. 035004, 2011.
- [22] V. H. Fan, L. E. Grosberg, S. S. Madugula, P. Hottowy, W. Dabrowski, A. Sher, A. M. Litke, and E. J. Chichilnisky, “Epiretinal stimulation with local returns enhances selectivity at cellular resolution,” *Journal of Neural Engineering*, 2018. [Online]. Available: <http://iopscience.iop.org/10.1088/1741-2552/aaee1>