

Biologically plausible phosphene simulation for the differentiable optimization of visual cortical prostheses

Maureen van der Grinten^{1,a}, Jaap de Ruyter van Steveninck^{1,a}✉, Antonio Lozano^{1,b}, Laura Pijnacker^a, Bodo Rückauer^a, Pieter Roelfsema^b, Marcel van Gerven^a, Richard van Wezel^a, Umut Güçlü^a, and Yağmur Güçlütürk^a

^aDonders Institute for Brain, Cognition and Behaviour, Nijmegen, the Netherlands

^bNetherlands Institute for Neuroscience, Amsterdam, the Netherlands

¹These authors have contributed equally to this work

Blindness affects millions of people around the world, and is expected to become increasingly prevalent in the years to come. For some blind individuals, a promising solution to restore a form of vision are cortical visual prostheses, which convert camera input to electrical stimulation of the cortex to bypass part of the impaired visual system. Due to the constrained number of electrodes that can be implanted, the artificially induced visual percept (a pattern of localized light flashes, or ‘phosphenes’) is of limited resolution, and a great portion of the field’s research attention is devoted to optimizing the efficacy, efficiency, and practical usefulness of the encoding of visual information. A commonly exploited method is the non-invasive functional evaluation in sighted subjects or with computational models by making use of simulated prosthetic vision (SPV) pipelines. Although the SPV literature has provided us with some fundamental insights, an important drawback that researchers and clinicians may encounter is the lack of realism in the simulation of cortical prosthetic vision, which limits the validity for real-life applications. Moreover, none of the existing simulators address the specific practical requirements for the electrical stimulation parameters. In this study, we developed a PyTorch-based, fast and fully differentiable phosphene simulator. Our simulator transforms specific electrode stimulation patterns into biologically plausible representations of the artificial visual percepts that the prosthesis wearer is expected to see. The simulator integrates a wide range of both classical and recent clinical results with neurophysiological evidence in humans and non-human primates. The implemented pipeline includes a model of the retinotopic organisation and cortical magnification of the visual cortex. Moreover, the quantitative effect of stimulation strength, duration, and frequency on phosphene size and brightness as well as the temporal characteristics of phosphenes are incorporated in the simulator. Our results demonstrate the suitability of the simulator for both computational applications such as end-to-end deep learning-based prosthetic vision optimization as well as behavioural experiments. The modular approach of our work makes it ideal for further integrating new insights in artificial vision as well as for hypothesis testing. In summary, we present an open-source, fully differentiable, biologically plausible phosphene simulator as a tool for computational, clinical and behavioural neuroscientists working on visual neuroprosthetics.

simulated prosthetic vision | cortical stimulation | phosphene vision | bionic vision | visual cortex | deep learning-based optimization | visual prosthesis | blindness | computational modelling | deep learning | neurotechnology |

Correspondence: jaap.deruyter@donders.ru.nl

Introduction

Globally, as per 2020, an estimated 43.3 million people were blind (1). For some cases of blindness, visual prosthetics may provide a promising solution. These devices aim to restore a rudimentary form of vision by interacting with the visual system using electrical stimulation (2–4). In particular, our work concerns prosthetic devices that target the primary visual cortex (V1). Despite recent advances in the field, more research is required before cortical prosthesis will become clinically available. Besides research into the improvement of the safety and durability of cortical implants (5, 6), a great portion of the research attention is devoted to optimizing the efficacy, efficiency, and practical usefulness of the prosthetic percepts. The artificially induced visual percepts consist of patterns of localized light flashes (‘phosphenes’) with a limited resolution. To achieve a functional level of vision, scene-processing is required to condense complex visual information from the surroundings in an intelligible pattern of phosphenes (7–13). Many studies employ a simulated prosthetic vision (SPV) paradigm to non-invasively evaluate the functional quality of the prosthetic vision with the help of sighted subjects (9, 14–16) or through ‘end-to-end’ approaches, using in silico models (7, 12, 17). Although the aforementioned SPV literature has provided us with some fundamental insights, an important drawback is the lack of realism and biological plausibility of the simulations. At this stage of the development, given the steadily expanding empirical literature on cortically-induced phosphene vision, it is both feasible and desirable to have a more phenomenologically accurate model of cortical prosthetic vision. Such an accurate simulator has already been developed for retinal prostheses (15), which has formed an inspiration for our work on simulation of cortical prosthetic vision. Thus, in this current work, we propose a realistic, biophysically-grounded computational model for the simulation of cortical prosthetic vision. Our simulator integrates empirical findings and quantitative models from the literature on cortical stimulation in V1. The elements that are modeled in our simulator include cortical magnification, current-dependent spread of activation and charge-dependent activation thresholds. Furthermore, our simulator models the effects of specific stimulation

parameters, accounting for temporal dynamics. A schematic overview of the pipeline and some example outputs are displayed in Figure 1. Our simulator runs in real-time, is open-source, and makes use of fully differentiable functions which allow for gradient-based optimization of phosphene encoding models. This design enables both simulations with sighted participants, as well as end-to-end optimization in machine-learning frameworks, thus fitting the needs of fundamental, clinical and computational scientists working on neuroprosthetic vision.

A. Background and related work.

A.1. Cortical prostheses.

Early attempts by Brindley and Lewin, and Dobbelle successfully reported to reliably induce the perception of phosphenes (described as localized round flashes of light) via electrical stimulation of the cortical surface (18, 19). More recent pre-clinical studies demonstrate promising results concerning the safety and efficacy of long-term stimulation in the primary visual cortex, either via surface electrodes (20, 21) or with intracortical electrodes (6, 22, 23). Other studies that performed V1 stimulation in sighted subjects (24, 25) and non-human primates (5, 26) have shown similar success. Some milestones include the implantation of over 1000 electrodes in a monkey's visual cortex (5), and the testing of a preliminary artificial vision system that presents visual information from the surroundings to a blind subject using a penetrating multi-electrode array in the visual cortex (6). Taken together, the previous literature provides strong evidence for the clinical potential of cortical prostheses for the blind.

A.2. Perceptual reports on cortical prosthetic vision.

In our simulator, we integrate empirical findings and quantitative models from existing literature on electrical stimulation in the visual cortex. Stimulation in V1 with intracortical electrodes is estimated to activate tens to thousands of neurons (27), resulting in the perception of often 'featureless' white dots of light with a circular shape (6, 18, 20, 22, 23). Due to the cortical magnification (the foveal information is represented by a relatively large surface area in visual cortex), the size of the phosphene is strongly correlated with its eccentricity (24, 25). Furthermore, phosphene size, stimulation threshold (defined as the minimum current to reliably produce a visible phosphene) and brightness are reported to be dependent on stimulation parameters such as the pulse width, train length, amplitude and frequency of stimulation (6, 20, 22–25). To account for these effects, we integrated and adapted previously proposed quantitative models that estimate the charge-dependent activation level of cortical tissue (24, 25, 28–31). Furthermore, our simulator includes a model of the temporal dynamics, observed by (23), accounting for response-attenuation after prolonged or repeated stimulation, as well as the delayed 'offset' of phosphene perception.

A.3. Simulated prosthetic vision.

A wide range of previous studies has employed SPV with sighted subjects to non-invasively investigate the usefulness

of prosthetic vision in everyday tasks, such as mobility (13, 14, 16, 32, 33), hand-eye coordination (34), reading (16, 35) or face recognition (36, 37). Several studies have examined the effect of the number of phosphenes, spacing between phosphenes and the visual angle over which the phosphenes are spread (e.g., (14, 34, 38–40)). The results of these studies vary widely, which could be explained by the difference in the implemented tasks, or, more importantly, by the differences in the simulation of phosphene vision. Most of the aforementioned studies used highly simplified phosphene simulations, with equally-sized phosphenes that were uniformly distributed over the visual field (informally referred to as the 'scoreboard model'). Furthermore, most studies assumed either full control over phosphene brightness or used quantized levels of brightness (e.g. "on" / "off"), but did not provide a description of the associated electrical stimulation parameters. Even studies that have explicitly made steps towards more realistic phosphene simulations, by taking into account cortical magnification or using visuotopic maps (34, 41, 42), did not model temporal dynamics or provide a description of the parameters used for electrical stimulation. Some recent studies developed descriptive models of the phosphene size or brightness as a function of the stimulation parameters (24, 25). Another very recent study has developed a deep-learning based model for predicting a realistic phosphene percept for single stimulating electrodes (43). While these studies have made important contributions to improve our understanding of the effects of different stimulation parameters, they do not provide a full simulation model that can be used for the functional evaluation of cortical visual prosthetic systems. This is what we aim to achieve in the current study. Meanwhile, a realistic and biologically-plausible simulator has already been developed for retinal prosthetic vision (Pulse2Percept, (15)), which takes into account the axonal spread of activation along ganglion cells and temporal nonlinearities to construct plausible simulations of stimulation patterns. Even though scientists increasingly realize that more realistic models of phosphene vision are required to narrow the gap between simulations and reality (13, 44), a biophysically-grounded simulation model for the functional evaluation of cortical prosthetic vision remains to be developed. Realistic SPV can aid technological developments by allowing neuroscientists, clinicians and engineers to test the perceptual effects of changes in stimulation protocols, and subsequently select stimulation parameters that yield the desired phosphene percepts without the need for extensive testing in blind volunteers. A realistic simulator could also be used as support in the rehabilitation process, assisting clinicians and caregivers in identifying potential problematic situations and adapt preprocessing or stimulation protocols accordingly (44).

A.4. Deep learning-based optimization of prosthetic vision.

SPV is often employed to develop, optimize and test encoding strategies for capturing our complex visual surroundings in an informative phosphene representation. Numerous scene-processing methods have been proposed in the literature, ranging from basic edge detection or contour-detection




Fig. 1. Left: schematic illustration of our simulator pipeline. Our simulator is initialized with electrode locations on a visuotopic map of the visual cortex. Each frame, the simulator takes a set of stimulation parameters that for each electrode specify the amplitude, pulse width, and frequency of electrical stimulation. Based on the electrode locations on the cortical map and the stimulation parameters, the phosphene characteristics are estimated and for each phosphene the effects are rendered on a map of the visual field. Finally, the phosphene renderings are summed to obtain the resulting simulated prosthetic percept. Right: Example renderings after initializing the simulator with four 10×10 electrode arrays (size: 5×5 millimetres) placed in the left hemisphere. The output is visualized for 166ms pulse trains with stimulation amplitudes of 40, 80, and $120 \mu A$, a pulse width of $170 \mu s$, and a frequency of $300 Hz$. In these example frames, we can observe the effects of cortical magnification, thresholds for activation, current-dependent spread (size) and proportion (brightness) of cortical tissue activation.

algorithms (45–47) to more intelligent deep learning-based approaches, which can be tailored to meet task-specific demands (7, 8, 10, 13, 36, 48, 49). The advantage of deep learning-based methods is clear: more intelligent and flexible extraction of useful information in camera input leads to less noise or unimportant information in the low-resolution phosphene representation, allowing for more successful completion of tasks. Some recent studies demonstrated that the simulation of prosthetic vision can even be incorporated directly into the deep learning-based optimization pipeline for end-to-end optimization (7, 12, 17). With end-to-end optimization, the image processing can be tailored to an individual user or specific tasks. Here, the usefulness of the prosthetic vision is evaluated by a computational agent, or decoder model, which assists in the direct optimization of the stimulation parameters required to optimally encode the present image. Note that one requirement of the optimization pipeline is that the simulator makes use of fully differentiable operations to convert the stimulation parameters to an image of phosphenes. In the current study, we have adapted and improved upon the framework of (7) to demonstrate the compatibility of our realistic simulator for end-to-end optimization. The currently proposed simulator can handle temporal sequences and can be compared to the aforementioned work; our experiments explore a more biologically grounded simulation of phosphene size and locations. Furthermore, instead of a more abstract or qualitative description of the required stimulation ('on' / 'off'), we included a biophysical model for predicting the perceptual effects of different stimulation parameters such as the current amplitude, the duration, the pulse width and the frequency. This opens new doors for optimization of the stimulation parameters in realistic ranges: although technological developments advance the state-of-the-art hardware capabilities rapidly, cortical prosthesis devices will be operating under energy constraints, due to both hardware limitations as well as safety limits regarding neurostimulation (50, 51). Deep learning methods trained in tandem with a biologically plausible phosphene simulator can be leveraged to produce constrained optimal stimulation paradigms that take into account these limitations, allowing

for safe and viable stimulation protocols to be developed.

Methods

Our simulator is implemented in Python, using the PyTorch deep learning library (52). The simulator makes use of differentiable functions which, given the entire set of phosphenes and their modelled properties, calculate the brightness of each pixel in the output image in parallel. This architecture makes our model memory intensive, but allows for fast computations that can be executed on a GPU. Each frame, the simulator maps electrical stimulation parameters (stimulation current, pulse width and frequency) to a predicted phosphene perception, taking into account the stimulation history. In the sections below, we discuss the different components of the simulator model, followed by a description of some showcase experiments. Our simulator can be imported as a python package and the source code is available on GitHub¹.

B. Visuotopic mapping.

Our simulator can be flexibly initialized with different modes for the phosphene topological mapping. In the most basic mode, a custom list of phosphene locations within the visual field can be provided by the user, expressed in polar coordinates. Alternatively, the phosphene locations can be calculated from a list of electrode locations on a flattened cortical map of V1. For mapping phosphene locations from the cortical electrode locations to the neuroprosthesis user's visual field, our simulator uses the reverse wedge-dipole visuotopic model of V1, proposed by (53). This model maps a complex polar coordinate $z = re^{i\theta}$ in the visual field, to a cortical location w in one hemisphere of V1, following the visuotopic relationship

$$w = k(\log(re^{i\alpha\theta} + a) - \log(re^{i\alpha\theta} + b)). \quad (1)$$

Here r and $-\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}$ are the eccentricity and azimuth of the point in the visual field, the parameter α controls the shear

¹The source code of our simulator can be retrieved from Github: <https://github.com/neuralcodinglab/dynaphos>. The latest stable release can be installed using pip: `$ pip install dynaphos`

of the ‘wedge map’, k is a scaling factor that scales the mapping to realistic proportions in cortical distance (millimeters), and a and b are parameters that control the singularities of the dipole model. For the mapping from cortical coordinates to phosphene location, we use the inverse of equation 1, which is given by

$$z = \Lambda^{-1} \left(\frac{ab(e^{\frac{w}{k}} - 1)}{b - ae^{\frac{w}{k}}} \right) \quad (2)$$

for the inverse shearing operation

$$\Lambda^{-1}(re^{i\theta}) = re^{i\frac{\theta}{\alpha}}. \quad (3)$$

The visuotopic model also provides us with the cortical magnification M , which defines the relative amount of cortical tissue that is involved in processing of visual information, depending on the eccentricity in the visual field. The cortical magnification is given by the derivative of equation 1 along the horizontal meridian:

$$M = \frac{k(b-a)}{(r+a)(r+b)}. \quad (4)$$

Here, M is given in millimetres of cortical surface per degree of visual angle. The parameters of the models are configurable. The default values are specified in section G. Note that in our simulation software, we provide the option of substituting equations 1-4, with other estimates described such as the mono- or dipole model in (53, 54).

C. Phosphene size.

Based on a model by (28), the phosphene size (in degrees),

$$P = \frac{D}{M} \quad (5)$$

is obtained via an estimation of the current spread from the stimulating electrodes, where

$$D = 2\sqrt{\frac{I}{K}} \quad (6)$$

is the diameter of the activated cortical tissue (in mm), for stimulation current I (in μA) and excitability constant K (in $\mu A/mm^2$). Note that the cortical magnification factor M is obtained in equation 4. The default value for K is specified in section G. In our simulation software, we provide the option to substitute equation 6 with an estimate by (25). Based on verbal descriptions (6, 22, 23), phosphenes are shown as Gaussian blobs with two standard deviations set equal to the phosphene size P , such that 95% of the Gaussian falls within the predicted phosphene size.

D. Phosphene brightness.

In our simulator, the brightness and detection threshold of each phosphene are based on a model of the intracortical tissue activation in response to electrical stimulation with biphasic square pulse trains. Our model assumes brightness and detection thresholds of phosphene perception to be primarily correlated with the deposited charge, and accounts

for the relative inefficiency of higher stimulation frequencies, longer pulse-widths, and longer train durations, as found in (6, 20, 24). Our simulator models the combined effects of these stimulation parameters as follows: First, we subtract from the stimulation amplitude I_{stim} a leak current I_0 , which represents the ineffective component of the stimulation input, and a memory trace Q (further explained in section F) that accounts for the decreased neural excitability after prior stimulation. I_0 is set equal to the rheobase current (the absolute threshold for continuous stimulation at infinite duration), following prior literature on the strength-duration relationship of neural tissue activation for isolated single-pulse trials (29). To calculate the effective stimulation current of trains of pulses, the remaining current amplitude is multiplied with the duty cycle of the stimulation signal ($Pw \cdot f$, the fraction of one period in which the signal is active), such that

$$I_{eff} = \max\left(0, (I_{stim} - I_0 - Q) \cdot Pw \cdot f\right) \quad (7)$$

for pulse width Pw and frequency f . Then, the cortical tissue activation is estimated by integrating the effective input current over multiple frames, using a leaky integrator model. For each frame with duration Δt , the estimated cortical activation is updated as

$$A_t = A_{t-\Delta t} + \Delta A \quad (8)$$

with

$$\Delta A = \left(-\frac{A_{t-\Delta t}}{\tau_{act}} + I_{eff} \cdot d \right) \cdot \Delta t. \quad (9)$$

Here, τ_{act} is the time constant of the activation decay in seconds and $d \in (0, 1]$ is a parameter that scales the duration of the stimulation relative to the frame duration. By default, d is set to 1 to simulate a stimulation duration equal to the frame duration, where the total pulse train duration is controlled with the number of successive frames in which stimulation is provided to the simulator. Finally, if the cortical activation reaches the detection threshold (explained in section E), the phosphene is activated with a brightness equal to the sigmoidal activation

$$\frac{1}{1 + e^{-\lambda(A - A_{50})}}, \quad (10)$$

where λ is the slope of the sigmoidal curve and A_{50} is the value of A for which the phosphene reaches half the maximum brightness.

E. Stimulation threshold.

Our simulator uses a thresholding model based on psychometric data from (6). Phosphenes are only generated when the cortical tissue activation (explained in section D) reaches the activation threshold A_{thr} , which is obtained for each electrode separately upon initialization of the simulator. To introduce a degree of variability between electrodes, A_{thr} is sampled from the normal distribution

$$\mathcal{N}(Th_{50}, \sigma^2). \quad (11)$$

The default values of the 50% probability threshold Th_{50} , and the standard deviation σ are fit on data from (6) and can be found in section G. Note that, by default, the detection thresholds remain constant after initialization. However, in accordance to the user requirements, the values can be flexibly adjusted or re-initialized manually.

F. Temporal dynamics.

Using a memory trace of the stimulation history, our simulator accounts for basic accommodation effects of brightness for prolonged or repeated stimulation, as described in prior work (23). Each frame, the memory trace is dynamically updated as follows:

$$Q_t = Q_{t-\Delta t} + \Delta Q \quad (12)$$

with

$$\Delta Q = \left(-\frac{Q_{t-\Delta t}}{\tau_{\text{trace}}} + I_{\text{eff}} \cdot \kappa \right) \Delta t. \quad (13)$$

Here, τ_{trace} is the time constant of the trace decay in seconds, and the parameter κ controls the input effect. Note that the memory trace is used for the phosphene brightness and not the phosphene size. As there is little experimental data on the temporal dynamics of phosphene size in relation to the accumulated charge, only the instantaneous current is used in the calculation of the phosphene size.

G. Parameter estimates.

By default, our model uses the parameters specified below. Unless stated otherwise, these parameter estimates were obtained by fitting our model to experimental data using the SciPy Python package, version 1.9.0 (55). More details on the comparison between the models' predictions and the experimental data can be found in the next section. Note that the parameter settings may strongly depend on the specific experimental conditions (such as the type of electrodes).

- In equations 1, 2, 3, 4 we use $a = 0.75$, $k = 17.3$, $b = 120$, and $\alpha = 0.95$, based on a fit by (53) on data of the human V1 from (56).
- In equation 6, the parameter K is set to $675 \mu A/mm^2$, following an estimate by (57), who measured the responses to intracortical stimulation in V1 at different current levels.
- In equation 7, we use a rheobase current $I_0 = 23.9 \mu A$ based on a fit on data from (6). Here, we used the strength-duration curve for tissue-excitability $Q_{\text{thr}} = I_0(c+t)$, with minimal input charge Q_{thr} , chronaxial parameter c and total stimulation duration t as described in (29).
- In equation 9, the parameter d is set equal to 1. The parameter τ_{act} is set equal to $0.111 s$ to reflect qualitative descriptions found in (23). Note: this parameter is not obtained by fitting to experimental data.
- In equation 10, we use a slope $\lambda = 19.2 \cdot 10^7$ and offset $A_{50} = 1.06 \cdot 10^{-6}$ for the brightness curve, based on a fit of our model on data by (6).

- The descriptive parameters of the distribution 11, are set to $Th_{50} = 9.14 \cdot 10^{-8}$ and $\sigma = 6.72 \cdot 10^{-8}$, based on a fit on psychometric data by (6).
- In equations 12 and 13, we use $\tau_{\text{trace}} = 1.97 \cdot 10^3 s$ and $\kappa = 14.0$. These values are based on a fit of our model to data from (23).

Experiments and Results

In this section we present computational experiments and comparisons with the literature to validate the performance, biological realism, and practical usability of our simulator.

H. Performance.

We tested the computational efficiency of our simulator, by converting a pre-processed example video² (1504 frames) into simulated phosphene images, for different numbers of phosphenes, and at varying image resolutions. The simulator was run on a CUDA-enabled graphics card (NVIDIA® A30) and each setting was run five times. The results are displayed in Figure 2. The lowest measured frame rate (10,000 phosphenes at a resolution of 256×256) was 28.7 frames per second. Note that the missing combinations in Figure 2 indicate that the required memory exceeded the capacity of our GPU, as the simulation of large numbers of phosphenes at high resolutions can be memory intensive.

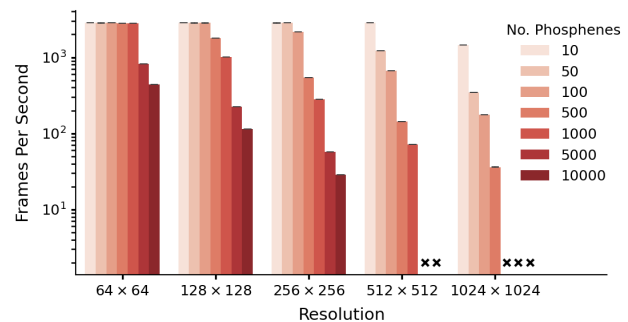


Fig. 2. Performance as a function of resolution and number of phosphenes. The data is based on 5 runs of 1540 frames per condition, with batch size equal to 1 frame. Simulation was run with an NVIDIA® A30 GPU. Crosses indicate missing conditions.

I. Biological plausibility.

Here, we report on experimental data obtained from the literature, and evaluate the capacity of our simulator of fitting these empirical data. Using equations 7-10, our simulator accurately reproduces the relative phosphene brightness that was reported in the previous study for different stimulation amplitudes ($R^2 = 0.950$, Figure 3). Figure 4 visualizes the effect of changing the stimulation parameters on the probability of phosphene perception, as estimated by our model. We compare our estimates with data reported by (6). As can be observed, our model accurately fits the reported effect of pulse width, frequency and train duration on the probability of phosphenes perception ($R^2 = 0.868$).

²The example video with the simulated phosphene output can be downloaded via [this link](#).




Fig. 3. Estimate of the relative phosphene brightness for different stimulation amplitudes. The simulator was provided with a stimulation train of $166ms$ with a pulse width of $170\mu s$ at a frequency of $300Hz$ (see equations 7-10). Left: the predicted peak brightness levels reproduced by our model (red) and psychometric data reported by (6) (light blue). Note that for stimulation amplitudes of $20.0\mu A$ and lower, the simulator generated no phosphenes as the threshold for activation was not reached. Right: the modeled tissue activation and brightness response over time. Values below the 50% threshold for the tissue activation and the corresponding brightness values are displayed with dashed lines.

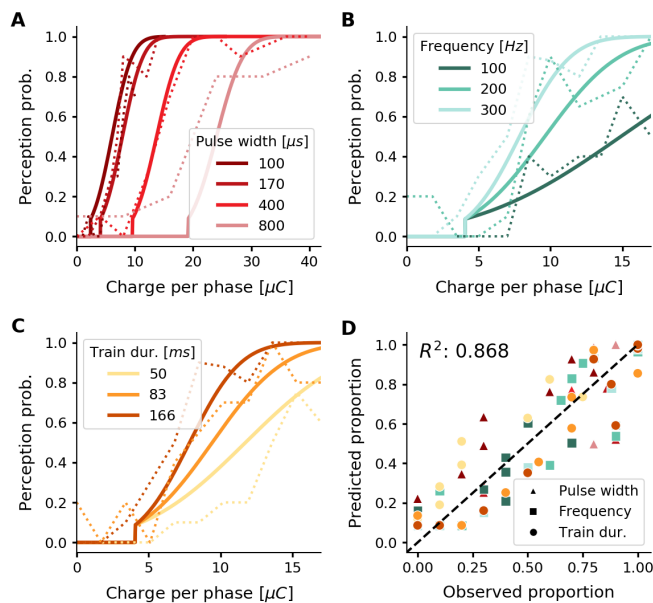


Fig. 4. (a-c) Psychometric curves (solid lines) overlaid on experimental data (dashed lines) (ref. 6, Fig. 2a, b). The predicted probability of phosphene perception is visualized as a function of charge per phase for (a) different pulse widths, (b) different frequencies, and (c) different train durations. Note that rather than the total charge per trial, we report the charge per phase to facilitate easy comparison with aforementioned experimental data. In panel (d) the probabilities of phosphene perception reproduced with our model are compared to the detection probabilities reported in (ref. 6, Fig. 2a, b). Colors conform to the conditions in panels a, b and c

Figure 5 displays the results of exploring our simulator's capacity of modeling temporal dynamics found in a previous published study by (23). For repeated stimulation at different timescales (intervals of 2 seconds, and intervals of 4 minutes), the measured brightness of a single phosphene is evaluated after fitting the memory trace parameters. The observed accommodation effects in our simulator are compared with the data from (23).

J. Usability in a deep learning SPV pipeline.

To validate that the simulator can conveniently be incorporated in a machine learning pipeline, we used our simulator in an existing SPV pipeline by (7) and performed several phosphene encoding optimization experiments, described be-




Fig. 5. Relative brightness of a phosphene in response to repeated stimulation, overlaid on experimental results by (23). The stimulation sequence consisted of 10 pulse trains at a four-second stimulation interval, followed by five pulse trains at an interval of four minutes to test recovery. Please notice the split x-axis with variable scaling.

low. In this pipeline, a convolutional neural network encoder is trained to process images or video frames and generate adequate electrode stimulation parameters. To train the encoder, a simulation of the prosthetic percept is generated by our differentiable phosphene simulator. This simulated percept is evaluated by a second convolutional neural network, the decoder, which decodes the simulated percept into a reconstruction of the original input image (Figure 6). The quality of the phosphene encoding is optimized by iteratively updating the network parameters of the encoder and decoder (simultaneously) using backpropagation of the reconstruction error. In addition to the reconstruction error, which measures similarity between the reconstruction and the input, we used a regularization term that measures similarity between the phosphenes and the input. For a more detailed description of the end-to-end optimization pipeline, see (7).

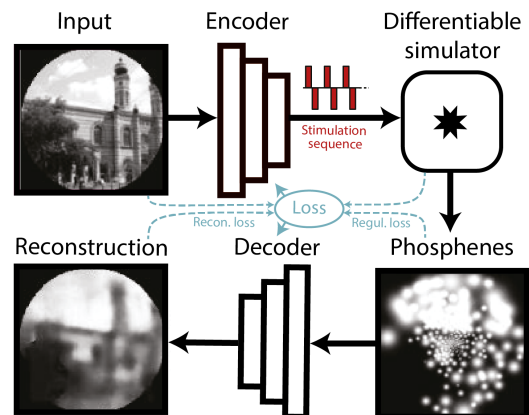


Fig. 6. Schematic illustration of the end-to-end machine-learning pipeline adapted from (7). A convolutional neural network encoder is trained to convert input images or video frames into a suitable electrical stimulation protocol. In the training procedure, our simulator generates a simulation of the expected prosthetic percept, which is evaluated by a second convolutional neural network that decodes a reconstruction of the input image. The quality of the encoding is iteratively optimized by updating the network parameters using back-propagation. Different loss terms can be used to constrain the phosphene encoding, such as the reconstruction error between the reconstruction and the input, a regularization loss between the phosphenes and the input, or a supervised loss term between the reconstructions and some ground-truth labeled data (not depicted here). Note that the internal parameters of the simulator (e.g. the estimated tissue activation) can also be used as loss terms.

J.1. Dynamic end-to-end encoding of videos.

In a further experiment, we explored the potential of using our simulator in a dynamic end-to-end encoding pipeline. We extended the previously published pipeline with 3D-

convolutions (as an additional temporal dimension) to enable encoding of subsequent video frames. The model was trained on a basic video dataset with moving white digits on a black background (Moving MNIST Dataset, (58)). We used video sequences of five frames. The framerate of the simulation was set at five frames per second. We used a combination of two even-weighted mean squared error (MSE) loss functions: the MSE loss between reconstruction and input, and the MSE loss between the simulated phosphene representation and the input. Figure 7 displays several frames after training for 45 epochs (for a total of 810,000 training examples). We can observe that the model has successfully learned to represent the original input frames in phosphene vision over time, and the decoder is able to approximately reconstruct the original input.



Fig. 7. Results of training the end-to-end pipeline on video sequences from the moving MNIST dataset (58). Columns indicate different frames. Top row: the input frames; middle row: the simulated phosphene; bottom row: the decoded reconstruction of the input. This figure is best viewed in digital format.

J.2. Constrained stimulation and naturalistic scenes.

In a second experiment, we trained the end-to-end model with a more challenging dataset containing complex images of naturalistic scenes (the ADE20K dataset (59)). In this experiment, we implemented the original pipeline described in (7) (experiment 4). The images were normalized and converted to grayscale, and we applied a circular mask such that the corners (outside the field covered by phosphenes) are ignored in the reconstruction task. The experiment consisted of three training runs, in which we tested different conditions: a free optimization condition, a constrained optimization condition, and a supervised boundary reconstruction condition. In the free optimization condition, the model was trained using an equally weighted combination of a MSE reconstruction loss between input and reconstruction, and a MSE regularization loss between the phosphenes and input images. After six epochs the model has learned to successfully find an optimal encoding strategy that can accurately represent the scene and allows the decoder to accurately reconstruct pixel intensities while qualitatively maintaining the image structure (see Figure 8). Importantly, note that the encoder has learned to encode brighter areas of the input picture by using large stimulation amplitudes (over $2000\mu A$). The encoding strategy found in such an unconstrained optimization scheme is not feasible for real-life applications. In practice, the electrical stimulation protocol will need to satisfy safety bounds and it will need to comply with technical requirements and limi-

tations of the stimulator hardware. For instance, rather than continuous stimulation intensities it is likely that the stimulator will allow for stimulation with only a number of (discrete) amplitudes. To evaluate whether our end-to-end pipeline can be harnessed to optimize the encoding in a constrained context, we performed a second training run (the constrained condition) where we reconfigured the encoder to output 10 discrete values between 0 and $128\mu A$. We used straight-through estimation with a smooth staircase function to estimate the gradients during backpropagation. We increased the training stability by adapting the relative weights of the reconstruction loss and the regularization loss (to 0.999 and 0.0001, respectively). The results of the safety-constrained training after six epochs are visualized in Figure 8. Note that overall, the resulting phosphenes are less bright and smaller due to the lower stimulation amplitudes. Nevertheless, the decoder is able to accurately reconstruct the original input. One limitation is that we did not test the subjective interpretability for human observers. As not all information in the scene is equally important, it may be informative to further constrain the phosphene representation to encode specific task-relevant features. In a third training run (the supervised boundary condition) we validated whether our simulator can be used in a supervised machine learning pipeline for the reconstruction of specific target features, such as the boundaries between objects. Instead of using the input image as a reference, now the MSE is used between the reconstruction and a ground truth target image and between the phosphene representation and the target image. The ground truth semantic boundary targets were obtained by performing canny edge detection and subsequent line thickening on the semantic segmentation labels provided with the dataset. The results after training for 16 epochs are visualized in Figure 8. Note that the model successfully converged to a sparse phosphene encoding that selectively represents the object boundaries.

Discussion

The aim of this study is to present a differentiable and biologically plausible phosphene simulator, which takes realistic ranges of stimulation parameters, and generates a phenomenologically accurate representation of phosphene vision. In order to achieve this, we have modeled and incorporated an extensive body of work regarding the psychophysics of phosphene perception. From the results presented in section I, we observe that our simulator is able to produce phosphene percepts that match the descriptions of phosphene vision that were gathered in basic and clinical visual neuroprosthetics studies that were conducted over the past decades. Running on a GPU, the simulator runs in real time, and as such, it could be used in clinical experiments with sighted volunteers. Furthermore, our proof-of-principle computational experiments, presented in section J demonstrate the suitability of our simulator for machine learning pipelines, aimed at improving the image processing and stimulation strategies. Here we discuss some implications of our findings.




Fig. 8. Results of training our simulator in an end-to-end pipeline on naturalistic images from the ADE20K dataset (59). In the constrained optimization condition and the supervised boundary reconstruction condition, the encoder was configured to output 10 discrete stimulation amplitudes within the safe range of stimulation (0 to $128\mu A$). The selected images represent the first three categories in the validation dataset ('Abbey', 'Access Road', 'Airbase'). Note that the brightness is enhanced in the phosphene images of the constrained optimization and the supervised boundary condition by 40%. The figure is best viewed in digital format.

K. Validation experiments.

K.1. Visuotopic mapping.

The results presented in Figure 1 illustrate the value of including a visuotopic model based on spread of cortical activation to realistically estimate phosphene locations and size. Some previous studies have used a model of cortical magnification (34, 60) or visuotopic mapping (41, 42) in their phosphene simulations. However, our simulator is the first to incorporate empirical models of the current spread in cortical tissue (24, 25, 28) to simulate the effects of stimulation current on the phosphene size. The accurate modelling of this biophysical relationship can help to increase the validity of simulation studies and brings fundamental SPV research closer to addressing questions regarding the practical real-life requirements of a visual prosthesis. Furthermore, the explicit link between the modeled area of cortical activation and the simulated phosphene sizes and locations makes our software very suitable for including new receptive field modeling results. Collaborative international projects such as the PRIME Resource Exchange (PRIME-RE) offer advanced tools which allow to fit probabilistic retinotopic maps generated from large samples to any individual NHP brain (61, 62) and it is currently possible to accurately predict human cortical receptive field mapping based on anatomical scans (63, 64), or other retinotopic mapping strategies that do not rely on visual input (65, 66). This opens new doors for future research into the functionality of visual prostheses with limited visual field coverage. Thanks to the machine learning compatibility, the model can also be used for the pre-operative optimization of the implant placement and design (67).

K.2. Threshold and brightness.

The results presented in Figures 3 and 4 indicate that our simulator closely models existing psychophysical data on the stimulation thresholds and phosphene brightness, for different electrical stimulation settings. Note that the effects found by (6) (that were modeled by us) are consistent with findings

by other studies, which report brighter phosphenes for higher stimulation strengths (23), and a lower stimulation efficiency (i.e. higher total charge thresholds) for longer pulse trains or higher pulse widths and frequencies (20). Moreover, our results are in line with other computational models for detection thresholds in the somatosensory cortex (68, 69). Our results indicate how a leaky integrator model and normally-distributed activation thresholds, provide a suitable approximation of the tissue activation in cortical prostheses. Note that alternative, more complex, models can possibly predict the psychometric data more accurately. However, most probably, this will entail a trade-off with the simplicity and modularity found in our current simulator. Future research may further improve our understanding of the neural processing underlying the conscious perception of phosphenes, possibly borrowing insights from the domain of natural vision. More elaborate theories on this matter have been developed and tested in (70). More specific limitations and suggestions for future adaptations are discussed in section M.

K.3. Temporal dynamics.

The results presented in Figure 5 reveal that the model accounts for experimental data on the accommodation in response to repeated stimulation in time periods up to 200 seconds. However, in contrast to the findings by (23), our simulator predicts a moderate recovery over the next 1000 seconds. Although we cannot provide an explanation for this difference, the modelled recovery largely stays within the 95% confidence interval of the experimental data. Similar to the other components of our simulator, the memory trace was chosen as a basic, yet effective, model of neural habituation. Possibly, a more complex, non-linear, model can more accurately fit the neurophysiological data, as discussed in section M. However, the presented model has the benefit of simplicity (there are only three parameters). Also, note that there is still some ambivalence in the clinical data. In contrast to the findings by (23), some studies have found accommodation over different time scales and sometimes no accommodation

at all (6, 19, 22, 71). More research is required for a better understanding of the neural response after repeated or prolonged stimulation.

K.4. Phosphene shape and appearance.

The appearance of phosphenes in our simulation (white, round, soft dots of light) are largely in line with previous reports on intracortical stimulation (6, 22, 23). However, more elongated shapes and more complex shapes have been reported as well (22, 24, 25). By using separately generated phosphene renderings, our simulator enables easy adjustments of the appearance of individual phosphenes. Additionally, we incorporated the possibility to change the default Gaussian blob appearance into Gabor patches with a specific frequency and orientation. Regarding the colour of phosphenes, there is still some ambivalence in reports, including descriptions of phosphene color ranging from black or white to different tones of color (23, 28, 72). Notably, increasing the stimulation amplitudes can lead the appearance to shift from colored to yellowish or white (23). This effect may be explained by the increased current spread for higher stimulation amplitudes, which is predicted to span multiple cortical columns coding for specific visual characteristics (e.g. orientation or colour), thus giving rise to phosphenes with amalgamated features (28). Currently, the limited amount of systematic data render it difficult to enable more accurate simulations of the variability in phosphene appearance.

L. End-to-end optimization.

L.1. Dynamic encoding.

The results presented in Figure 7 demonstrate that our proposed realistic phosphene simulator is well-suited for the dynamic optimization in an end-to-end architecture. Our proof-of-principle video-encoding experiments are the first to explicitly optimize the stimulation across the temporal domain. This provides a basis for the further exploration of computationally-optimized dynamic stimulation patterns. Dynamic optimization of the stimulation may be necessary to counteract unwanted effects such as response fading due to accommodation after repeated or prolonged stimulation (23), or delayed phosphene perception after stimulation onset. The inclusion of a realistic simulator in the optimization pipeline enables researchers to exploit the optimal combination of stimulation parameters to obtain precise control over the required perception. Moreover, besides acquiring optimal control over the transfer function from stimulation to phosphenes, dynamic phosphene encoding could also prove useful to expand the encoded information along the temporal domain (21). Although this was not in the scope of the current study, our software is well-suited for simulation experiments that further investigate dynamic stimulation. Note that there remain some challenging perceptual considerations for the design of useful dynamics stimulation patterns (for an excellent review on asynchronous stimulation in relation to flicker fusion, form vision and apparent motion perception, please see (73)).

L.2. Constrained, efficient stimulation for natural stimuli.

Our second optimization experiment addresses a more natural and realistic context. The results presented in Figure 8 demonstrate that our simulator is well-suited for the optimization of prosthetic vision to natural stimuli and that it can be configured to comply with constraints regarding the stimulation protocol. Note that the quality of the reconstructions for the constrained version of the encoder indicate that the model can still find an efficient information encoding strategy using a limited set of stimulation amplitudes (10 discrete values between 0 and 128 μ A). These results are in line with previous results on constrained end-to-end optimization, indicating that task-relevant information can be maximized under sparsity constraints (7). While in the current experiments the stimulation amplitude is maximized for the individual electrodes, future studies could investigate other sparsity constraints, such as a maximum total charge delivered per second across all electrodes. Ultimately, rather than an accurate overall portrayal of the visual surroundings, a visual prosthesis may need to prioritize task-relevant information. For this reason, in recent SPV research with sighted human observers much attention is devoted to semantic (boundary) segmentation for discriminating the important information from irrelevant background (10, 13, 48). Note that the explored image processing strategies in these behavioral studies are compatible with the automated optimization through an end-to-end machine learning pipeline. Our experiments exemplify how supervision targets obtained from semantic segmentation data can be adopted to promote task-relevant information in the phosphene representation. Furthermore, in addition to reconstruction of the input or labelled targets, another recent study experimented with different decoding tasks, including more interactive, goal-driven tasks in virtual game environments (17). Although these proof-of-principle results remain to be translated to real-world tasks and environments, they provide a valuable basis for further exploration. Ultimately, the development of task-relevant scene-processing algorithms will likely benefit both from computational optimization experiments as well as exploratory SPV studies with human observers.

L.3. Interpretability and perceptual correspondence.

Besides the encoding efficiency (characterized by the computational decodability of task-relevant information), it is important to consider the subjective interpretability of the simulated phosphene representation. From the results in Figure 8 it can be observed that the model has successfully learned to preserve correspondence between the phosphene representation and the input image in all of the training conditions. However, as a more formal analysis was outside the scope of this study, we do not further quantify the subjective interpretability. In our model the subjective interpretability was promoted through the regularization loss between the simulated phosphenes and the input image. Similarly, a recent study adapted an auto-encoder architecture designed to directly maximize the perceptual correspondence between a target representation and the simulated percept in a retinal prosthesis, using basic stimuli (12). The preservation of sub-

jective interpretability in an automated optimization pipeline remains a non-trivial challenge, especially when using natural stimuli. This may hold even more for cortical prostheses, as the distinct punctuate phosphenes are in nature very dissimilar from natural images, possibly hampering perceptual similarity metrics that rely on low-level feature correspondence. Regardless of the implementation, it will always be important to include human observers in the optimization cycle to ensure subjective interpretability for the end user (74, 75).

M. General limitations and future directions.

M.1. Performance and hardware.

There are some remaining practical limitations and challenges for future research. We identify three considerations for future research related to the the performance of our model and the required hardware for implementation in an experimental setup. Firstly, although our model runs in real-time and is faster than state-of-the art simulation for retinal prostheses (15), there is a trade-off between speed and the memory demand. Therefore, for higher resolutions and larger number of phosphenes, future experimental research may need to adapt a simplified version of our model - although most of the simulation conditions can be run easily with common graphical cards. Secondly, a suggestion for follow-up research, is to combine our simulator with the latest developments in mixed reality (XR) to enable immersive simulation in virtual environments. More specifically, a convenient direction would be the implementation of our simulator using the Cg shader programming language for graphics processing, which is used in 3D game engines like Unreal Engine, or Unity 3D, as previously demonstrated for epiretinal simulations by (38). Thirdly, and lastly, future studies could integrate our simulation software with eye-tracking technology. Due to the retinotopic organization of visual cortex, cortical stimulation leads to phosphenes that are not fixed to a position in space, but change after every eye movement. These effects should be taken into account for a faithful simulation of the experience of a prosthesis user. Furthermore, as has been demonstrated in previous SPV studies (34, 35), one could investigate the potential benefits of including an eye tracker in the prosthetic hardware to allow for sampling of the visual environment using eye movements. Note that testing such gaze-assisted processing does not require any changes to our simulator. It merely involves processing the gaze-centered image as opposed to the entire camera input.

M.2. Complexity and realism of the simulation.

There are some remaining challenges regarding the realistic simulation of the effects of neural stimulation. A complicating factor is that cortical neuroprostheses are still in the early stages of development. Neurostimulation hardware and stimulation protocols are continuously being improved (21), and clinical trials with cortical visual neuroprostheses are often limited to small numbers of blind volunteers (11, 76). Therefore, it is no surprise that the amount of data that is available at the present moment is limited, often open for multi-

ple interpretations, and sometimes contains apparent contradictory information. Notably, the trade-off between model complexity and accurate psychophysical predictions is a recurrent theme in the validation of the components implemented in our simulator. These factors play a role in some of the potential limitations of our current simulator. Here we name a few of the important limitations and some interesting directions for future research. Firstly, in our simulator, phosphenes are only rendered when the activation is above threshold and vice-versa. This might be an inaccurate depiction of the perceptual experience of an implant user, and in reality the distinction may be less strict. The conscious perception of phosphenes requires considerable training and the detection process is influenced by attention (6). Although our implementation is effective for modeling the psychometric data, alternative implementations could also be considered. The perceptual effect of different simulated phosphene threshold implementations for sighted subjects remains to be evaluated in future SPV work. Secondly, the leaky integrator and the memory trace that are implemented in our simulator might be an oversimplified model of tissue activation in the visual cortex. This means that some non-linear dynamics might be missed. Also, several studies reported that higher stimulation amplitudes may give rise to double phosphenes (18, 19, 23, 77), or a reduction of phosphene brightness (23). Furthermore, interactions between simultaneous stimulation of multiple electrodes can have an effect on the phosphene size and sometimes lead to unexpected percepts (6). Further clinical data could help to improve our understanding of these non-linear dynamics. A third limitation is that our simulator currently only models responses of V1 stimulation. Future studies could explore the possible extension of modeling micro-stimulation of higher visual areas, such as V2, V3, V4 or inferotemporal cortex. In previous NHP research, reliable phosphene thresholds could be obtained with the stimulation of in V1, V2, V3A, MT (78). Furthermore, IT stimulation has shown to bias face perception (79). Similar effects have been confirmed in human subjects, and previous work has demonstrated that electrical stimulation of higher order visual areas can elicit a range of feature-specific percepts (80–82). Our simulator could be extended with maps of higher cortical areas with clear retinotopy, and an interesting direction for future research will be the implementation of feature-specific percepts, including texture, shape and colour.

N. Conclusion.

We present a biologically plausible simulator of phosphene vision. This simulator models psychophysical and neurophysiological findings in a wide array of experimental results. Its phenomenologically accurate simulations allow for the optimisation of visual cortical prosthesis in a manner that drastically narrows the gap between simulation and reality, compared to previous studies of simulated phosphene vision. It can operate in real time, therefore being a viable option for behavioural experiments with sighted volunteers. Additionally, the PyTorch implementation and its differentiable nature makes it a good choice for machine learning approaches to study and optimize phosphene vision. The modular de-

sign of our simulator allows for straightforward adaptation of novel insights and improved models of cortical activation. In summary, our open-source, fully differentiable, biologically plausible phosphene simulator aims to provide an accessible bedrock software platform that fits the needs of fundamental, clinical and computational vision scientists working on cortical neuroprosthetic vision. With this work, we aspire to contribute to increasing the field's translational impact.

ACKNOWLEDGEMENTS

This work was supported by two grants (NESTOR, INTENSE) of the Dutch Organization for Scientific Research (NWO) and the European Union's Horizon 2020 research and innovation programme (under grant agreement No 899287). We thank Xing Chen for help with the compilation and reviewing of relevant literature regarding phosphene perception.

Bibliography

- Rupert R.A. Bourne, Jaimie D. Steinmetz, Seth Flaxman, Paul Svtil Briant, Hugh R. Taylor, Serge Resnikoff, Robert James Casson, Amir Abdoli, Eman Abu-Gharbieh, Ashkan Afshin, Hamid Ahmadieh, Yonas Akalu, Alehegn Aderaw Alamneh, Wondu Alemayehu, Ahmed Samir Alfaar, Wahid Alipour, Etsay Woldu Anbesu, Sofia Androudi, ..., and Theo Vos. Trends in prevalence of blindness and distance and near vision impairment over 30 years: An analysis for the Global Burden of Disease Study. *The Lancet Global Health*, 9(2): e130–e143, 2021. ISSN 2214-109X. doi: 10.1016/S2214-109X(20)30425-3.
- Eduardo Fernandez. Development of visual Neuroprostheses: trends and challenges. *Bioelectronic Medicine*, 4(1):12, dec 2018. ISSN 2332-8886. doi: 10.1186/s42234-018-0013-8.
- Edward Bloch, Yvonne Luo, and Lyndon da Cruz. Advances in retinal prosthesis systems. *Therapeutic Advances in Ophthalmology*, 11:2515841418817501, jan 2019. ISSN 2515-8414. doi: 10.1177/2515841418817501.
- Kamil Nowik, Ewa Langwińska-Wośko, Piotr Skopiński, Katarzyna E. Nowik, and Jacek P. Szaflik. Bionic eye review – An update. *Journal of Clinical Neuroscience*, 78:8–19, aug 2020. ISSN 09675868. doi: 10.1016/j.jocn.2020.05.041.
- Xing Chen, Feng Wang, Eduardo Fernandez, and Pieter R Roelfsema. Shape perception via a high-channel-count neuroprosthesis in monkey visual cortex. *Science*, 370(6521): 1191–1196, dec 2020. ISSN 0036-8075. doi: 10.1126/science.abd7435.
- Eduardo Fernández, Arantxa Alfaro, Cristina Soto-Sánchez, Pablo Gonzalez-Lopez, Antonio M. Lozano, Sebastian Peña, María Dolores Grima, Alfonso Rodil, Bernardeta Gómez, Xing Chen, Pieter R. Roelfsema, John D. Rolston, Tyler S. Davis, and Richard A. Normann. Visual percepts evoked with an intracortical 96-channel microelectrode array inserted in human occipital cortex. *Journal of Clinical Investigation*, 131(23), dec 2021. ISSN 1558-8238. doi: 10.1172/JCI1151331.
- Jaap de Ruyter van Steveninck, Umüt Güçlü, Richard van Wezel, and Marcel van Gerven. End-to-end optimization of prosthetic vision. *Journal of Vision*, 22(2):20, feb 2022. ISSN 1534-7362. doi: 10.1167/jov.22.2.20.
- Antonio Lozano, Juan Sebastián Suárez, Cristina Soto-Sánchez, Javier Garrigós, J. Javier Martínez-Alvarez, J. Manuel Ferrández, and Eduardo Fernández. NeuroLight: A Deep Learning Neural Interface for Cortical Visual Prostheses. *International Journal of Neural Systems*, 30(09):2050045, sep 2020. ISSN 0129-0657. doi: 10.1142/S0129065720500458.
- Richard A Normann, Bradley A Greger, Paul House, Samuel F Romero, Francisco Pelayo, and Eduardo Fernandez. Toward the development of a cortically based visual neuroprosthesis. *Journal of Neural Engineering*, 6(3):035001, jun 2009. ISSN 1741-2560. doi: 10.1088/1741-2560/6/3/035001.
- Melani Sanchez-Garcia, Ruben Martinez-Cantin, and Jose J. Guerrero. Semantic and structural image segmentation for prosthetic vision. *PLOS ONE*, 15(1):e0227677, jan 2020. ISSN 1932-6203. doi: 10.1371/journal.pone.0227677.
- Philip R Troyk. The Intracortical Visual Prosthesis Project. In *Artificial Vision*, pages 203–214. Springer International Publishing, Cham, 2017. doi: 10.1007/978-3-319-41876-6_16.
- Jacob Granley, Lucas Relic, and Michael Beyeler. Hybrid Neural Autoencoders for Stimulus Encoding in Visual and Other Sensory Neuroprostheses. *arXiv (preprint)*, pages 1–21, may 2022.
- Nicole Han, Sudhanshu Srivastava, Aiwen Xu, Devi Klein, and Michael Beyeler. Deep Learning-Based Scene Simplification for Bionic Vision. In *Augmented Humans Conference 2021*, pages 45–54, New York, NY, USA, feb 2021. ACM. ISBN 9781450384285. doi: 10.1145/3458709.3458982.
- Jaap de Ruyter van Steveninck, Tom van Gestel, Paula Koenders, Guus van der Ham, Floris Vereecken, Marcel van Gerven, Yağmur Güçlütürk, and Richard van Wezel. Real-world indoor mobility with simulated prosthetic vision: The benefits and feasibility of contour-based scene simplification at different phosphene resolutions. *Journal of Vision*, 22(2):1–14, 2022. ISSN 15347362. doi: 10.1167/JOV.22.2.1.
- Michael Beyeler, Geoffrey Boynton, Ione Fine, and Ariel Rokem. pulse2percept: A Python-based simulation framework for bionic vision. In *Proceedings of the 16th Python in Science Conference*, number Scipy, pages 81–88. SciPy, 2017. doi: 10.25080/shinma-7f4c6e7-00c.
- Kichul Cha, Kenneth W. Horch, and Richard A. Normann. Mobility performance with a pixelized vision system. *Vision Research*, 32(7):1367–1372, jul 1992. ISSN 00426989. doi: 10.1016/0042-6989(92)90229-C.
- Burcu Küçükoğlu, Bodo Rueckauer, Nasir Ahmad, Jaap de Ruyter van Steveninck, Umüt Güçlü, and Marcel van Gerven. Optimization of Neuroprosthetic Vision via End-to-End Deep Reinforcement Learning. *International Journal of Neural Systems*, 32(11):1–15, nov 2022. ISSN 0129-0657. doi: 10.1142/S0129065722500526.
- G. S. Brindley and W S Lewin. The sensations produced by electrical stimulation of the visual cortex. *The Journal of Physiology*, 196(2):479–493, may 1968. ISSN 00223751. doi: 10.1113/jphysiol.1968.sp008519.
- W. H. Dobbelle and M. G. Mladejovsky. 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, dec 1974. ISSN 00223751. doi: 10.1113/jphysiol.1974.sp010766.
- Soroush Niketeghad, Abirami Muralidharan, Uday Patel, Jessy D. Dorn, Laura Bonelli, Robert J Greenberg, and Nader Pouratian. Phosphene perceptions and safety of chronic visual cortex stimulation in a blind subject. *Journal of Neurosurgery*, 132(6):2000–2007, jun 2020. ISSN 0022-3085. doi: 10.3171/2019.3.JNS182774.
- Michael S. Beauchamp, Denise Oswald, Ping Sun, Brett L. Foster, John F. Magnotti, Soroush Niketeghad, Nader Pouratian, William H. Bosking, and Daniel Yoshor. Dynamic Stimulation of Visual Cortex Produces Form Vision in Sighted and Blind Humans. *Cell*, 181(4):774–783.e5, may 2020. ISSN 00928674. doi: 10.1016/j.cell.2020.04.033.
- M. Bak, J. P. Girvin, F. T. Hambrecht, C. V. Kufta, G. E. Loebe, and E. M. Schmidt. Visual sensations produced by intracortical microstimulation of the human occipital cortex. *Medical & Biological Engineering & Computing*, 28(3):257–259, may 1990. ISSN 0140-0118. doi: 10.1007/BF02442682.
- E. M. Schmidt, M. J. Bak, F. T. Hambrecht, C. V. Kufta, D. K. O'Rourke, and P. Vallabhanath. Feasibility of a visual prosthesis for the blind based on intracortical micro stimulation of the visual cortex. *Brain*, 119(2):507–522, 1996. ISSN 0006-8950. doi: 10.1093/brain/119.2.507.
- Jonathan Winawer and Josef Parvizi. Linking Electrical Stimulation of Human Primary Visual Cortex, Size of Affected Cortical Area, Neuronal Responses, and Subjective Experience. *Neuron*, 92(6):1213–1219, 2016. ISSN 10974199. doi: 10.1016/j.neuron.2016.11.008.
- William H. Bosking, Michael S. Beauchamp, and Daniel Yoshor. Electrical Stimulation of Visual Cortex: Relevance for the Development of Visual Cortical Prosthetics. *Annual Review of Vision Science*, 3(1):141–166, sep 2017. ISSN 2374-4642. doi: 10.1146/annurev-vision-111815-114525.
- Peter H. Schiller, Warren M. Slocum, Michelle C. Kwak, Geoffrey L. Kendall, and Edward J. Tehovnik. New methods devised specify the size and color of the spots monkeys see when striate cortex (area V1) is electrically stimulated. *Proceedings of the National Academy of Sciences of the United States of America*, 108(43):17809–17814, 2011. ISSN 00278424. doi: 10.1073/pnas.1108337108.
- Mark H Histed, Vincent Bonin, and R Clay Reid. Direct Activation of Sparse, Distributed Populations of Cortical Neurons by Electrical Microstimulation. *Neuron*, 63(4):508–522, aug 2009. ISSN 08966273. doi: 10.1016/j.neuron.2009.07.016.
- Edward J. Tehovnik and Warren M. Slocum. Phosphene induction by microstimulation of macaque V1. *Brain Research Reviews*, 53(2):337–343, feb 2007. ISSN 01650173. doi: 10.1016/j.brainresrev.2006.11.001.
- L.A. Geddes. Accuracy Limitations of Chronaxie Values. *IEEE Transactions on Biomedical Engineering*, 51(1):176–181, jan 2004. ISSN 0018-9294. doi: 10.1109/TBME.2003.820340.
- I.C. Bruce, M.W. White, L.S. Irlich, S.J. O'Leary, Scott Dynes, Eric Javel, and G.M. Clark. A stochastic model of the electrically stimulated auditory nerve: single-pulse response. *IEEE Transactions on Biomedical Engineering*, 46(6):617–629, jun 1999. ISSN 00189294. doi: 10.1109/10.764938.
- Sungshin Kim, Thierry Callier, and Sliman J. Bensmaia. A computational model that predicts behavioral sensitivity to intracortical microstimulation. *Journal of Neural Engineering*, 14(1): 016012, feb 2017. ISSN 1741-2560. doi: 10.1088/1741-2552/14/1/016012.
- Gislin Dagnelie, Pearse Keane, Venkata Narla, Liancheng Yang, James Weiland, and Mark Humayun. Real and virtual mobility performance in simulated prosthetic vision. *Journal of Neural Engineering*, 4(1):S92–S101, mar 2007. ISSN 1741-2560. doi: 10.1088/1741-2560/4/1/S11.
- Victor Vergnien, Marc J.M. Macé, and Christophe Jouffrais. Simplification of Visual Rendering in Simulated Prosthetic Vision Facilitates Navigation. *Artificial Organs*, 41(9):852–861, sep 2017. ISSN 0160564X. doi: 10.1111/aor.12868.
- Nishant R. Srivastava, Philip R. Troyk, and Gislin Dagnelie. Detection, eye–hand coordination and virtual mobility performance in simulated vision for a cortical visual prosthesis device. *Journal of Neural Engineering*, 6(3):035008, jun 2009. ISSN 1741-2560. doi: 10.1088/1741-2560/6/3/035008.
- Jörg Sommerhalder, Benjamin Rappaz, Raoul de Haller, Angélica Pérez Fornos, Avinoam B. Safran, and Marco Pelizzone. Simulation of artificial vision: II. Eccentric reading of full-page text and the learning of this task. *Vision Research*, 44(14):1693–1706, jun 2004. ISSN 00426989. doi: 10.1016/j.visres.2004.01.017.
- Caroline J.M. Bollen, Umüt Guclu, Richard J. A. van Wezel, Marcel A.J. van Gerven, and Yagmur Gucluturk. Simulating neuroprosthetic vision for emotion recognition. In *2019 8th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, pages 85–87. IEEE, sep 2019. ISBN 978-1-7281-3891-6. doi: 10.1109/ACIIW.2019.8925229.
- Jonathan Winawer and Josef Parvizi. Linking Electrical Stimulation of Human Primary Visual Cortex, Size of Affected Cortical Area, Neuronal Responses, and Subjective Experience. *Neuron*, 92(6):1213–1219, dec 2016. ISSN 08966273. doi: 10.1016/j.neuron.2016.11.008.
- Jacob Thomas Thorn, Enrico Migliorini, and Diego Ghezzi. Virtual reality simulation of epiretinal stimulation highlights the relevance of the visual angle in prosthetic vision. *Journal of Neural Engineering*, 17(5):056019, oct 2020. ISSN 1741-2560. doi: 10.1088/1741-2552/17/5/056019.
- Melani Sanchez-Garcia, Roberto Morollon-Ruiz, Ruben Martinez-Cantin, Jose J. Guerrero, and Eduardo Fernandez-Jover. Assessing visual acuity in visual prostheses through a virtual-reality system. *arXiv (preprint)*, pages 1–19, may 2022.
- N. Parikh, L. Itti, M. Humayun, and J. Weiland. Performance of visually guided tasks using simulated prosthetic vision and saliency-based cues. *Journal of Neural Engineering*, 10(2): 026017, apr 2013. ISSN 1741-2560. doi: 10.1088/1741-2560/10/2/026017.
- Tamas Fehervari, Masaru Matsuoka, Hirotugu Okuno, and Tetsuya Yagi. Real-Time Simulation of Phosphene Images Evoked by Electrical Stimulation of the Visual Cortex. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelli-*

- gence and Lecture Notes in Bioinformatics), volume 6443 LNCS, pages 171–178. 2010. ISBN 3642175368. doi: 10.1007/978-3-642-17537-4_22.
42. Wai Ho Li. Wearable computer vision systems for a cortical visual prosthesis. *Proceedings of the IEEE International Conference on Computer Vision*, pages 428–435, 2013. doi: 10.1109/ICCV.2013.63.
 43. Jacob Granley, Alexander Riedel, and Michael Beyeler. Adapting Brain-Like Neural Networks for Modeling Cortical Visual Prostheses. *arXiv (preprint)*, sep 2022.
 44. Gislin Dagnelie. Psychophysical Evaluation for Visual Prosthesis. *Annual Review of Biomedical Engineering*, 10(1):339–368, aug 2008. ISSN 1523-9829. doi: 10.1146/annurev-bioeng.10.061807.160529.
 45. J. Boyle, A. Maeder, and W. Boles. Static image simulation of electronic visual prostheses. In *The Seventh Australian and New Zealand Intelligent Information Systems Conference, 2001*, number November, pages 85–88. IEEE, 2001. ISBN 1-74052-061-0. doi: 10.1109/ANZIS.2001.974055.
 46. Jason A. Dowling, Anthony Maeder, and Wageeh Boles. Mobility enhancement and assessment for a visual prosthesis. In Amir A. Amini and Armando Manduca, editors, *The Seventh Australian and New Zealand Intelligent Information Systems Conference, 2001*, volume 5369, page 780. IEEE, apr 2004. ISBN 1-74052-061-0. doi: 10.1117/12.534866.
 47. Fei Guo, Yuan Yang, and Yong Gao. Optimization of Visual Information Presentation for Visual Prosthesis. *International Journal of Biomedical Imaging*, 2018:1–12, 2018. ISSN 1687-4188. doi: 10.1155/2018/3198342.
 48. Alex Rasla and Michael Beyeler. The Relative Importance of Depth Cues and Semantic Edges for Indoor Mobility Using Simulated Prosthetic Vision in Immersive Virtual Reality. *28th ACM Symposium on Virtual Reality Software and Technology*, pages 1–11, nov 2022. doi: 10.1145/3562939.3565620.
 49. Antonio Lozano, Cristina Soto-Sánchez, Javier Garrigós, J. Javier Martínez, J. Manuel Fernández, and Eduardo Fernández. A 3D Convolutional Neural Network to Model Retinal Ganglion Cell's Responses to Light Patterns in Mice. *International Journal of Neural Systems*, 28(10):1850043, dec 2018. ISSN 0129-0657. doi: 10.1142/S0129065718500430.
 50. R.V. Shannon. A model of safe levels for electrical stimulation. *IEEE Transactions on Biomedical Engineering*, 39(4):424–426, apr 1992. ISSN 00189294. doi: 10.1109/10.126616.
 51. Douglas B McCreery, William F Agnew, and Leo A Bullara. The Effects of Prolonged Intracortical Microstimulation on the Excitability of Pyramidal Tract Neurons in the Cat. *Annals of Biomedical Engineering*, 30(1):107–119, jan 2002. ISSN 0090-6964. doi: 10.1114/1.1430748.
 52. Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Advances in Neural Information Processing Systems*, 32(NeurIPS), dec 2019. ISSN 10495258.
 53. J. R. Polimeni, M. Balasubramanian, and E. L. Schwartz. Multi-area visuotopic map complexes in macaque striate and extra-striate cortex. *Vision Research*, 46(20):3336–3359, 2006. ISSN 00426989. doi: 10.1016/j.visres.2006.03.006.
 54. Eric L. Schwartz. Cortical mapping and perceptual invariance: A reply to cavanagh. *Vision Research*, 23(8):831–835, jan 1983. ISSN 00426989. doi: 10.1016/0042-6989(83)90206-7.
 55. Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J van der Walt, Matthew Brett, Joshua Wilson, K Jarrod Millman, Nikolay Mayorov, Andrew R J Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, Ilhan Polat, Yu Feng, Eric W Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E A Quintero, Charles R Harris, Anne M Archibald, Antônio H Ribeiro, Fabian Pedregosa, Paul van Mulbregt, Aditya Vijaykumar, Alessandro Pietro Bardelli, Alex Rothberg, Andreas Hilboll, Andreas Kloeckner, Anthony Scopatz, Antony Lee, Ariel Rokem, C. Nathan Woods, Chad Fulton, Charles Masson, Christian Häggström, Clark Fitzgerald, David A. Nicholson, David R. Hagen, Dmitrii V. Pasechnik, Emanuele Olivetti, Eric Martin, Eric Wieser, Fabrice Silva, Felix Lenders, Florian Wilhelm, G. Young, Gavin A. Price, Gert-Ludwig Ingold, Gregory E. Allen, Gregory R. Lee, Hervé Audren, Irvin Probst, Jörg P. Dietrich, Jacob Silter, James T Webber, Janko Slavič, Joel Nothman, Johannes Buchner, Johannes Kulick, Johannes L. Schönberger, José Vinícius de Miranda Cardoso, Joscha Reimer, Joseph Harrington, Juan Luis Cano Rodríguez, Juan Nunez-Iglesias, Justin Kuczynski, Kevin Tritz, Martin Thoma, Matthew Newville, Matthias Kümmerer, Maximilian Bolingbroke, Michael Tartre, Mikhail Pak, Nathaniel J. Smith, Nikolai Nowaczyk, Nikolay Shebanov, Oleksandr Pavlyk, Per A. Brodtkorb, Perry Lee, Robert T. McGibbon, Roman Feldbauer, Sam Lewis, Sam Tygier, Scott Sievert, Sebastiano Vigna, Stefan Peterson, Surhud More, Tadeusz Pudlik, Takuya Oshima, Thomas J. Pingel, Thomas P. Robitaille, Thomas Spura, Thouis R. Jones, Tim Cera, Tim Leslie, Tiziano Zito, Tom Krauss, Utkarsh Upadhyay, Yaroslav O. Halchenko, and Yoshiki Vázquez-Baeza. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, mar 2020. ISSN 1548-7091. doi: 10.1038/s41592-019-0686-2.
 56. Jonathan C. Horton. The Representation of the Visual Field in Human Striate Cortex. *Archives of Ophthalmology*, 109(6):816, jun 1991. ISSN 0003-9950. doi: 10.1001/archophth.1991.01080060080030.
 57. E J Tehovnik, A S Tolia, F Sultan, W M Slocum, and N K Logothetis. Direct and Indirect Activation of Cortical Neurons by Electrical Microstimulation. *Journal of Neurophysiology*, 96(2):512–521, aug 2006. ISSN 0022-3077. doi: 10.1152/jn.00126.2006.
 58. Nitish Srivastava, Elman Mansimov, and Ruslan Salakhutdinov. Unsupervised Learning of Video Representations using LSTMs. *32nd International Conference on Machine Learning, ICML 2015*, 1:843–852, feb 2015.
 59. Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic Understanding of Scenes Through the ADE20K Dataset. *International Journal of Computer Vision*, 127(3):302–321, mar 2019. ISSN 0920-5691. doi: 10.1007/s11263-018-1140-0.
 60. Nadia Paraskevoudi and John S Pezaris. Full gaze contingency provides better reading performance than head steering alone in a simulation of prosthetic vision. *Scientific Reports*, 11(1):11121, may 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-86996-4.
 61. P Christiaan Klink, Xing Chen, Wim Vanduffel, and Pieter R Roelfsema. Population receptive fields in nonhuman primates from whole-brain fMRI and large-scale neurophysiology in visual cortex. *eLife*, 10:e67304, nov 2021. ISSN 2050-084X. doi: 10.7554/eLife.67304.
 62. Adam Messinger, Nikoloz Sirmplataze, Katja Heuer, Kep Kee Loh, Rogier B Mars, Julien Sein, Ting Xu, Daniel Glen, Benjamin Jung, Jakob Seidlitz, Paul Taylor, Roberto Toro, Eduardo A. Garza-Villarreal, Caleb Sponheim, Xindi Wang, R. Austin Benn, Bastien Cagna, Rakshit Dadarwal, Henry C. Evrard, Pamela Garcia-Saldivar, Steven Giavasis, Renée Hartig, Claude Lepage, Cirong Liu, Piotr Majka, Hugo Merchant, Michael P. Milham, Marcello G.P. Rosa, Jordy Tasserie, Lynn Uhrig, Daniel S. Margulies, and P. Christiaan Klink. A collaborative resource platform for non-human primate neuroimaging. *NeuroImage*, 226: 117519, feb 2021. ISSN 10538119. doi: 10.1016/j.neuroimage.2020.117519.
 63. Noah C. Benson, Omar H. Butt, Ritobrato Datta, Petya D. Radoeva, David H. Brainard, and Geoffrey K. Aguirre. The Retinotopic Organization of Striate Cortex Is Well Predicted by Surface Topology. *Current Biology*, 22(21):2081–2085, nov 2012. ISSN 09609822. doi: 10.1016/j.cub.2012.09.014.
 64. Noah C Benson, Omar H Butt, David H Brainard, and Geoffrey K Aguirre. Correction of Distortion in Flattened Representations of the Cortical Surface Allows Prediction of V1-V3 Functional Organization from Anatomy. *PLoS Computational Biology*, 10(3):e1003538, mar 2014. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1003538.
 65. Rainer Goebel, Rick van Hoof, Salil Bhat, Michael Luhrs, and Mario Senden. Reading Imagined Letter Shapes from the Mind's Eye Using Real-time 7 Tesla fMRI. In *2022 10th International Winter Conference on Brain-Computer Interface (BCI)*, pages 1–3. IEEE, IEEE, feb 2022. ISBN 978-1-6654-1337-4. doi: 10.1109/BCI53720.2022.9735031.
 66. Andrew S Bock, Paola Binda, Noah C Benson, Holly Bridge, Kate E Watkins, and Ione Fine. Resting-State Retinotopic Organization in the Absence of Retinal Input and Visual Experience. *Journal of Neuroscience*, 35(36):12366–12382, sep 2015. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.4715-14.2015.
 67. Rick van Hoof. *The brain as image processor and generator*. PhD thesis, Maastricht University, 2022.
 68. Gene Y. Fridman, Hugh T. Blair, Aaron P. Blaisdell, and Jack W. Judy. Perceived intensity of somatosensory cortical electrical stimulation. *Experimental Brain Research*, 203(3):499–515, jun 2010. ISSN 0014-4819. doi: 10.1007/s00221-010-2254-y.
 69. Sungshin Kim, Thierry Callier, and Sliman J Bensmaia. A computational model that predicts behavioral sensitivity to intracortical microstimulation. *Journal of Neural Engineering*, 14(1): 016012, feb 2017. ISSN 1741-2560. doi: 10.1088/1741-2552/14/1/016012.
 70. Bram van Vugt, Bruno Dagnino, Devavrat Vartak, Houman Safsafi, Stefano Panzeri, Stanislas Dehaene, and Pieter R Roelfsema. The threshold for conscious report: Signal loss and response bias in visual and frontal cortex. *Science*, 360(6388):537–542, may 2018. ISSN 0036-8075. doi: 10.1126/science.aar7186.
 71. John R. Bartlett, Robert W. Doty, sr., Barry B. Lee, Nubio Negrão, and William H. Overman, jr. Deleterious Effects of Prolonged Electrical Excitation of Striate Cortex in Macaques. *Brain, Behavior and Evolution*, 14(1-2):46–66, 1977. ISSN 0006-8977. doi: 10.1159/000125575.
 72. Edward J. Tehovnik, Warren M. Slocum, Stelios M. Smirnakis, and Andreas S. Tolia. Microstimulation of visual cortex to restore vision. In *Progress in Brain Research Volume 175*, pages 347–375. 2009. doi: 10.1016/S0096-6123(09)17524-6.
 73. Susana Moleirinho, Andrew J Whalen, Shelley I Fried, and John S Pezaris. The impact of synchronous versus asynchronous electrical stimulation in artificial vision. *Journal of Neural Engineering*, 18(5):051001, oct 2021. ISSN 1741-2560. doi: 10.1088/1741-2552/abc6f1.
 74. Tristan Fauvel and Matthew Chalk. Human-in-the-loop optimization of visual prosthetic stimulation. *Journal of Neural Engineering*, 19(3):036038, jun 2022. ISSN 1741-2560. doi: 10.1088/1741-2552/ac7615.
 75. Michael Beyeler and Melani Sanchez-Garcia. Towards a Smart Bionic Eye: AI-powered artificial vision for the treatment of incurable blindness. *Journal of Neural Engineering*, 19(6):063001, dec 2022. ISSN 1741-2560. doi: 10.1088/1741-2552/ac69d.
 76. Eduardo Fernández and Richard A Normann. CORTIVIS Approach for an Intracortical Visual Prostheses. In *Artificial Vision*, pages 191–201. Springer International Publishing, Cham, 2017. doi: 10.1007/978-3-319-41876-6_15.
 77. Denise Oswald, William Bosking, Ping Sun, Sameer A. Sheth, Soroush Niketeghad, Michelle Armenta Salas, Uday Patel, Robert Greenberg, Jessy Dorn, Nader Pouratian, Michael Beauchamp, and Daniel Yoshor. Multi-electrode stimulation evokes consistent spatial patterns of phosphenes and improves phosphene mapping in blind subjects. *Brain Stimulation*, 14(5):1356–1372, sep 2021. ISSN 1935861X. doi: 10.1016/j.brs.2021.08.024.
 78. Dona K. Murphey and John H.R. Maunsell. Behavioral Detection of Electrical Microstimulation in Different Cortical Visual Areas. *Current Biology*, 17(10):862–867, may 2007. ISSN 09609822. doi: 10.1016/j.cub.2007.03.066.
 79. Seyed-Reza Afraz, Roozbeh Kiani, and Hossein Esteky. Microstimulation of inferotemporal cortex influences face categorization. *Nature*, 442(7103):692–695, aug 2006. ISSN 0028-0836. doi: 10.1038/nature04982.
 80. Dona K Murphey, John H R Maunsell, Michael S Beauchamp, and Daniel Yoshor. Perceiving electrical stimulation of identified human visual areas. *Proceedings of the National Academy of Sciences*, 106(13):5389–5393, mar 2009. ISSN 0027-8424. doi: 10.1073/pnas.0804998106.
 81. H W Lee, S B Hong, D W Seo, W S Tae, and S C Hong. Mapping of functional organization in human visual cortex: Electrical cortical stimulation. *Neurology*, 54(4):849–854, feb 2000. ISSN 0028-3878. doi: 10.1212/WNL.54.4.849.
 82. Gerwin Schalk, Christoph Kapeller, Christoph Guger, Hiroshi Ogawa, Satoru Hiroshima, Rosa Lafer-Sousa, Zeynep M Saygin, Kyouosuke Kamada, and Nancy Kanwisher. Facephens and rainbows: Causal evidence for functional and anatomical specificity of face and color processing in the human brain. *Proceedings of the National Academy of Sciences*, 114(46):12285–12290, nov 2017. ISSN 0027-8424. doi: 10.1073/pnas.1713447114.