PAPER • OPEN ACCESS

Neural activity shaping utilizing a partitioned target pattern

To cite this article: Martin J Spencer *et al* 2021 *J. Neural Eng.* **18** 046025

View the [article online](#) for updates and enhancements.

You may also like

- [EEG-fMRI Bayesian framework for neural activity estimation: a simulation study](#)
Pierpaolo Croce, Alessio Basti, Laura Marzetti et al.
- [A biomimetic electrical stimulation strategy to induce asynchronous stochastic neural activity](#)
Emanuele Formento, Edoardo D'Anna, Sandra Gribi et al.
- [Development of modularity in the neural activity of childrens brains](#)
Man Chen and Michael W Deem

EEG/ECOG AMPLIFIERS
& ELECTRODES
ELECTRICAL/CORTICAL
STIMULATORS
REAL-TIME PROCESSING

SHOP NOW

PAPER

Neural activity shaping utilizing a partitioned target pattern

Martin J Spencer^{1,*} , Tatiana Kameneva^{1,2} , David B Grayden¹, Anthony N Burkitt^{1,4} and Hamish Meffin^{1,3,4} ¹ Department of Biomedical Engineering, The University of Melbourne, Parkville, Victoria, Australia² Faculty of Science, Engineering and Technology, Swinburne University of Technology, Hawthorn, Victoria, Australia³ National Vision Research Institute, Australian College of Optometry, Carlton, Victoria, Australia⁴ Denotes joint senior authorship.

* Author to whom any correspondence should be addressed.

E-mail: martin.spencer@unimelb.edu.au**Keywords:** retinal prosthesis, bionic eye, current steering, neural activity, stimulation strategy

RECEIVED

11 October 2020

REVISED

18 February 2021

ACCEPTED FOR PUBLICATION

8 March 2021

PUBLISHED

31 March 2021

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Abstract

Electrical stimulation of neural tissue is used in both clinical and experimental devices to evoke a desired spatiotemporal pattern of neural activity. These devices induce a local field that drives neural activation, referred to as an activating function or generator signal. In visual prostheses, the spread of generator signal from each electrode within the neural tissue results in a spread of visual perception, referred to as a phosphene. *Objective.* In cases where neighbouring phosphenes overlap, it is desirable to use current steering or neural activity shaping strategies to manipulate the generator signal between the electrodes to provide greater control over the total pattern of neural activity. Applying opposite generator signal polarities in neighbouring regions of the retina forces the generator signal to pass through zero at an intermediate point, thus inducing low neural activity that may be perceived as a high-contrast line. This approach provides a form of high contrast visual perception, but it requires partitioning of the target pattern into those regions that use positive or negative generator signals. This discrete optimization is an NP-hard problem that is subject to being trapped in detrimental local minima. *Approach.* This investigation proposes a new partitioning method using image segmentation to determine the most beneficial positive and negative generator signal regions. Utilizing a database of 1000 natural images, the method is compared to alternative approaches based upon the mean squared error of the outcome. *Main results.* Under nominal conditions and with a set computation limit, partitioning provided improvement for 32% of these images. This percentage increased to 89% when utilizing image pre-processing to emphasize perceptual features of the images. The percentage of images that were dealt with most effectively with image segmentation increased as lower computation limits were imposed on the algorithms. *Significance.* These results provide a new method to increase the resolution of neural stimulating arrays and thus improve the experience of visual prosthesis users.

1. Introduction

Electrode arrays are used in clinical and experimental settings to induce responses in neural tissue [1–3]. These devices aim to induce spatiotemporal patterns of neural activity by activating each electrode in the array with an appropriate current pulse [4, 5]. Electrode arrays form a key component of many neural implants that aim to treat dysfunction or lost function in the nervous system by stimulating target neural populations.

In the context of retinal implants, an implant user's visual perception is closely aligned to the pattern of neural activity across their retina or primary visual cortex [6, 7]. An important limitation of electrical stimulation is that each stimulated electrode induces a region of visual perception, referred to as a phosphene [7, 8]. This is the result of the flow of electrical current in the neural tissue [9, 10] as well as characteristics of the neural structures and connections [6]. Under conventional stimulation strategies, this current spread imposes a limitation on the

maximum resolution of the induced pattern of neural activation.

Figure 1(a) details the forward **linear-nonlinear (LNL) model** that calculates the neural activity across the retina or visual cortex for a given pattern of electrode settings [11–13]. The LNL model is based on experimental results that demonstrate that the generator signal can be estimated from recordings of neural responses to electrical stimulation [12, 13]. These results show that the LNL model is an accurate way to predict neural activity directly evoked by an arbitrary pattern of electrode amplitudes that include both cathodic-first (positive) and anodic-first (negative) stimuli. In the retina, this neural activity can be interpreted to represent the retinal ganglion cell firing rate, retinal ganglion cell firing probability, local field potential, or multi-unit activity. In the LNL model, the generator signal at each pixel is the linear sum of the contributions from each electrode, $\vec{s}$; this is the linear part of the LNL model. The two-dimensional neural activity pattern is divided into a grid of locations, which we will call pixels, that can be represented as a vector of neural activations, $\vec{r}$. At each pixel, the neural activity is modelled as a static nonlinear function of the generator signal at that pixel, $\vec{g}$, here taken to be the absolute value (i.e. rectification); this is the non-linear part of the LNL model. The generator signal may be interpreted biophysically as the trans-membrane current at the site of spike initiation, and is related to Rattay's activating function [14, 15]. It is expected that each pixel represents the activity of a large number of neurons.

Conventional stimulation strategies apply sequential electrode stimulation patterns that are

simply spatially down-sampled versions of the desired neural activity patterns [16, 17]. Under this strategy, illustrated in figure 1(b), each electrode is activated individually to induce neural activity with the aim to match the local target neural activity at the location of the electrode. To avoid the blurring of neighboring phosphenes, current from corresponding electrodes must not substantially overlap. This constraint places an upper bound on the number of perceptually useful electrodes in a given area [18].

Multipolar stimulation [19–21] and focused multipolar stimulation [19, 22–24] attempt to overcome the limitation of current spread by using simultaneous combinations of cathodic-first biphasic pulses (positive) and anodic-first biphasic pulses (negative) on electrodes with spatially overlapping current spread. These strategies seek to control the current spread to create higher contrast at a finer spatial scale in the overall neural activity pattern.

Rather than targeting control of the current in neural tissue, the neural activity shaping strategy [11] instead calculates electrode settings that aim to shape neural activity explicitly. This neural activity shaping strategy can be implemented with subsets of electrodes to induce a local pattern of neural activation or with all electrodes to induce an entire field of neural activation. The strategy effectively inverts the LNL model to find a pattern of stimulation across electrodes, $\vec{s}$, that provides an optimal match of predicted neural activity, $\vec{r}$, to a target pattern of neural activity, $\vec{r}^*$. The target pattern of neural activity is derived from the image (figure 1(c), bottom row, dashed lines) and is converted to a target generator signal by

dealing with the rectifying nonlinearity in two alternative ways.

In the **unpartitioned** version of the neural activity shaping strategy, the generator signal is chosen to take only positive values (figure 1(c) middle row, dashed lines). Note that both positive and negative electrode pulses are used (corresponding to cathodic-first and anodic-first pulses), but the negative pulses (figure 1(c), top row) are utilized *only* to attenuate positive generator signal spread in a push-pull arrangement. This can provide much greater control over neural activity than the conventional phosphene-based strategies (figure 1(b)) but can still fail to represent narrow regions of low neural activity that exist between regions of high neural activity. In images, these would be perceived as dark lines. These narrow features may be particularly important in images pre-processed to emphasize the edges of objects.

However, it has been shown experimentally that a net negative generator signal can also induce neural activity [12, 13]. This is because neural activation occurs for both cathodic-first and anodic-first biphasic pulses. Therefore, neural activity is modelled as the magnitude generator signal (to a first approximation, figure 1(a)). Figure 1(d) shows an example of using a partition vector, $\vec{\varphi}$, to partition the target neural activity into appropriate regions of positive and negative generator signal (middle row, dotted line). This is the **partitioned** version of the neural activity shaping strategy. This provides a much higher contrast distinction between the resulting two regions of high neural activity (bottom row, solid line), which emphasizes the perception of dark lines in the target neural activity. The efficient implementation of this additional partition step is computationally non-trivial and is the topic of this study.

The previously published partitioning approach uses a Hopfield network with simulated annealing [11] but suffers from limitations. There are three potential limitations of partitioning the target pattern:

- (a) The search for a beneficial segmentation of the target into regions of positive and negative generator signal is computationally infeasible for real-time processing as required by a retinal implant. For example, with a very simple target of 6×6 pixels of retinal activity, there are $N_r = 36$ locations of target neural activity. In principle, it is necessary to decide if each pixel will be achieved with either positive or negative generator signal. This leads to a discrete search space of $2^{N_r-1} = 2^{35} \approx 10^{10}$ possible combinations for each target pattern, an impossible task on a portable processor.
- (b) An exhaustive search can be avoided by using a Hopfield network to perform a guided search

similar to gradient descent. However, the search space using this method typically includes many local minima that can lead to solutions with higher errors than the solution provided by the unpartitioned strategy.

- (c) Even in cases where the global minimum solution is found, this solution can include trade-offs that misrepresent the true pattern, leading to artefacts that may be perceptually distracting.

The second and third of these potential shortcomings are illustrated in figure 2 using a 6×6 electrode array with overlapping phosphenes and a 40×40 -pixel target. Given a simple target of neural activity (figure 2, top row), the unpartitioned neural activity shaping strategy assumes that the generator signal comprises only positive values (figure 2(a), middle row), which leads to the solution depicted in figure 2(a) (bottom row). The solution that most accurately recreates the perceptual features of the target, shown in figure 2(b), partitions the neural activity pattern optimally into regions of positive and negative generator signal that, in this case, fall on the left and right of the dark line, respectively. The same target is partitioned into hypothetical regions that demonstrate the potential detrimental local minimum typical of the results of a single cycle of a Hopfield network (figure 2(c)). This counter-intuitive solution is chosen to illustrate the unusual solutions that result from algorithms that can be thought of as ‘relaxing’ or ‘quenching’ from high variation (or high energy) solutions to lower variation solutions. The outcome is that they include artefacts similar to the phase transitions between domains in the Ising model of crystal or magnetic lattices. A different target is used to illustrate a potential detrimental global minimum trade-off solution that fails to represent the break in the dark line (figure 2(d)).

This paper presents the results of a method to avoid all three of the potential shortcomings. The proposed new method uses conventional image segmentation to partition the neural activity pattern into regions of high and low activity. Each segmented region of the resulting pattern is assigned a uniform generator signal polarity. By undertaking the initial image segmentation, the search space is reduced from 2^{N_r-1} possible partitions to $2^{N_{\text{reg}}-1}$ possible partitions, where N_r is the number of target neural activity locations and N_{reg} is the number of regions in the partitioned target. If $N_{\text{reg}} \ll N_r$, then this represents a substantial reduction in computation. The Hopfield network is then applied as the final step to find the best partition. This image segmentation approach to partitioning also explicitly biases the final solution away from artefacts that arise due to local minima (figure 2(c)) and away from artefacts that arise due to trade-off solutions (figure 2(d)).

Figure 2. Potential limitations of the partitioned strategies. A 6×6 electrode array is used with overlapping phosphenes and a 40×40 -pixel target. (a) The solution provided by an unpartitioned strategy. (b) The solution that provides the most perceptual value. (c) A hypothetical local minimum solution using the partitioned strategy. (d) Another global optimum solution that includes a trade-off that misrepresents the perceptual features in the target pattern. First row: target neural activity pattern with white indicating high neural activity and black indicating no neural activity. Second row: hypothetical partition into positive and negative regions (red and blue, respectively). Third row: simulated outcome in neural activity.

This partitioning approach is additionally tested in combination with a set of standard approaches to image pre-processing.

2. Methods

The following descriptions of the LNL forward model, conventional approach, inverse LNL model, partition vector, and Hopfield network are summaries of methods applied in a previous publication [11]. The section describing simulated annealing is an improved version of the previous simulated annealing approach because it uses an optimized nonlinear annealing trajectory. The repeated Hopfield network is a new approach that has been introduced because, unlike annealing, it allows for parallelization of its implementation. The section on the Hopfield network with image segmentation is the principal new approach introduced in the present study, which is compared with other previously used methods.

2.1. LNL forward model

A LNL model is used for phenomenological prediction of neural activation by electrical stimulation [11–13] (figure 1(a)). The model uses scaled units, where (S) represents spikes, (T)—time, (I)—current, and (L)—length. The **linear** part of the model used a dimensionless transfer matrix, $\mathbf{W}$, to calculate a **generator signal**, $\vec{g}(I)$, at each pixel in the retina that results from a set of electrode currents, $\vec{s}(I)$,

$$\vec{g} = \mathbf{W}\vec{s}. \quad (1)$$

The vector, $\vec{s}$, has the same number of elements as there are electrodes in the implanted electrode array, N_e , and the vector, $\vec{g}$, has the same number of pixels

that we wish to model in the retina, N_r . (As an aside, if the same electrodes were used to record neural response to electrical stimulation so as to estimate the model [13] then $N_r = N_e$, but we make no such assumption here.) The columns of the matrix, $\mathbf{W}$, which is of size $N_r \times N_e$, quantify the weighting of the contribution of each electrode to the generator signal associated with the neurons in each pixel in $\vec{g}$. Each electrode contributes a spread of neural activity that results in the perception of phosphenes. For simplicity, this spread is modelled as Gaussian with uniform standard deviation, $\sigma(L)$. Other choices are possible, including estimating $\mathbf{W}$ by recording responses, this is an area for future work.

The spike rate at each pixel in the retina, $\vec{r}(ST^{-1})$, is assumed to be the magnitude of the generator signal at that pixel (figure 1(a)); this is the **nonlinear** part of the model,

$$\vec{r} = |\vec{g}|. \quad (2)$$

This proportionality assumption simplifies the following derivation while retaining the important qualitative properties of the experimentally measured two-sided nonlinearities. Specifically, these nonlinearities are observed to be symmetrical, with a minimum of neural activity at a generator signal value of zero [13].

2.2. Conventional approach

The basic function of a retinal implant is to activate electrodes to induce perception of a given image. The conventional approach to selecting electrode stimulation strengths, given a particular target neural activity pattern, $\vec{r}^*$, is to first downsample the target pattern. Each electrode is then stimulated in proportion to

the local intensity of the downsampled target neural activity. This approach is independent of the shape of the individual phosphenes. In this study, an enhanced version of this conventional approach is implemented by making the selection of each electrode's current amplitude dependent upon the shape of its associated neural activity profile. In this enhanced conventional approach, each electrode current is proportional to the projection of the target pattern of neural activity onto phosphene/ERF associated with that electrode,

$$\vec{s} = k \cdot \vec{r}^*{}^T \mathbf{W}, \quad (3)$$

where $\vec{r}^*$ is the given target pattern of neural activation and k is a normalization constant. This process mimics the conventional process but improves upon it by using information about the neural activation pattern induced by each electrode. This acts as a conservative benchmark against which we can compare the proposed methods.

2.3. Inverse LNL model

An improved approach is to use the inverse of the above LNL model in order to calculate the electrode activation pattern, $\vec{s}$, required to induce a given target pattern of neural activation, $\vec{r}^*$. The above LNL model (equations (1) and (2)) can be combined to give

$$\vec{r}^* = |\mathbf{W}\vec{s}|, \quad (4)$$

which is equivalent to

$$\mathbf{R}^* \vec{\varphi} = \mathbf{W}\vec{s}, \quad (5)$$

where

$$\vec{\varphi} = \text{sgn}(\mathbf{W}\vec{s}), \quad (6)$$

and $\mathbf{R}^*$ is a diagonal matrix whose diagonal elements are the vector, $\vec{r}^*$, and $\vec{\varphi}$ is here referred to as the **partitioning vector**, whose entries are ± 1 . We then multiply both sides of equation (5) by the pseudoinverse of $\mathbf{W}$ to obtain the desired set of electrode currents,

$$\vec{s} = \mathbf{W}^+ \mathbf{R}^* \vec{\varphi}, \quad (7)$$

where the '+' superscript symbol represents the pseudoinverse of the transfer matrix, $\mathbf{W}$, obtained via singular value decomposition (SVD) implemented by function 'svd' of MATLAB (version 2019a, MathWorks, MA, USA).

2.4. The partition vector

The sign of the generator signal for each pixel, given by the contents of the partition vector, $\vec{\varphi}$, are free to be chosen. The unpartitioned neural activity strategy simply sets every element of the partition vector to +1 (figure 1(c)). The problem investigated in this paper is to find the optimum values for each of the elements of $\vec{\varphi}$ in equation (7) (figure 1(d)). There are 2^{N_r-1}

possible (non-degenerate) combinations of elements of $\vec{\varphi}$. The optimum combinations are those that give the most visually useful representation of the target neural activation, $\vec{r}^*$. We generally assume that this is achieved by minimizing the mean squared error ε ($S^2 T^{-2}$) between the model predicted pattern and the actual target pattern of retinal activity,

$$\varepsilon = \|\vec{r}^* - \vec{r}\|^2. \quad (8)$$

As has been shown previously [11], this leads to the objective function,

$$\Theta = \min_{\vec{\varphi}_i \in \{-1,1\}} \vec{\varphi}^T [\mathbf{R}^* (\mathbf{I} - \mathbf{P}\mathbf{P}^T) \mathbf{R}^*] \vec{\varphi}, \quad (9)$$

where $\mathbf{I}$ is the identity matrix and $\mathbf{P}$ are the left eigenvectors, from the SVD of $\mathbf{W}$. This fits the form of a binary optimization problem, $\min_{\vec{v}_i \in \{-1,1\}} \vec{v}^T \mathbf{M} \vec{v}$, which is NP-hard for a general matrix, $\mathbf{M}$, which is of size $N_r \times N_r$.

2.5. Hopfield network

A potential (i.e. locally optimal) solution to the binary optimization problem in equation (9) is provided by a Hopfield network of N_r nodes using the elements of the symmetric matrix $\mathbf{M} = \mathbf{R}^* (\mathbf{I} - \mathbf{P}\mathbf{P}^T) \mathbf{R}^*$ as the strengths of connections between the nodes and $\vec{\varphi}$ as the states of the nodes [25]. Since the diagonal elements of the matrix, $\mathbf{M}$, are the coefficients to the terms, $\varphi_i^2 = 1$, these do not contribute to the minimization of ε and can be set to 0. This avoids self-connections in the Hopfield network, which does not function to minimize the objective function with self-connected nodes.

A random set of initial values for $\vec{\varphi}$ are used and then updated according to the iterated Hopfield network rule, $\vec{\varphi}_{i+1} \leftarrow \text{sgn}(\mathbf{M} * \vec{\varphi}_i)$, until the vector, $\vec{\varphi}_i$, is static and new iterations do not lead to any change. As is required for Hopfield optimization, the multiplication, $\mathbf{M} * \vec{\varphi}_i$, is not performed in the normal, parallel way, but instead sequentially, in which each row of $\mathbf{M}$ is multiplied by the vector, $\vec{\varphi}_k$, and the appropriate value is then updated within the vector, $\vec{\varphi}_i$, before the next row is multiplied. This process is referred to as an 'iteration' of the Hopfield network. Each iteration of this rule leads to a reduction in the error, ε , and is repeated until it reaches a (local) minimum. In the present study, the entire process of repeated iterations of the Hopfield network is referred to as a 'cycle' of the Hopfield network. The Hopfield network underlies the three approaches used in this paper. However, in practice, this single-cycle Hopfield network typically results in a detrimentally high error, ε , for the partitioning (see Results and figure 2(c)). This is due to being caught in local minima; consequently, it is not used on its own.

2.6. Approach 1: repeated Hopfield network

A single cycle of the Hopfield network frequently leads to local minima in the minimization. These local

minima are often associated with distracting artefacts in the visual field (e.g. see figure 2(c)). One approximate solution to this is to simply perform single-cycles multiple times with new random conditions, with each new cycle beginning with different random initial conditions. The process completes when a chosen computational limit, C_{lim} , has been reached and the Hopfield cycle with the lowest ε is chosen as the best result. This approach is referred to as the repeated Hopfield network strategy and has the practical advantage of being implementable in parallel because each cycle of the Hopfield network is not dependent on other cycles. In principle, parallelization provides a dramatic increase in speed when implemented using an appropriate architecture (e.g. a field programmable gate array, FPGA), though this is not done here.

2.7. Approach 2: Hopfield network with simulated annealing

The repeated Hopfield network approach with random initial conditions does not build on the results of previous cycles. Simulated annealing uses the outcome of previous cycles as a starting point for new cycles of Hopfield network iterations. Specifically, each of the cycles commences with an initial guess, $\vec{\varphi}$, that is associated with the lowest error solution of all the previous Hopfield cycles and combined with some randomization of the values of $\vec{\varphi}$. This randomization is implemented by reversing the sign of a fraction of the elements of the partition vector, $\vec{\varphi}$, with probability, F_{rev} . This process begins with maximum levels of randomization and then reduces to minimum levels of randomization, allowing the solution to ‘relax’ to its lowest level in the standard annealing fashion.

The probability, F_{rev} , is updated throughout the optimization according to the following annealing trajectory that interpolates between the maximum and minimum possible values of F_{rev} ,

$$F_{\text{rev}}(n) = 0.5 - \left(0.5 - \frac{1}{N_r}\right) (C/C_{\text{lim}})^x. \quad (10)$$

The maximum randomization is associated with a sign flip probability, $F_{\text{rev}} = 0.5$, because the error, ε , is the same with either $\vec{\varphi}$ or $-\vec{\varphi}$, and the minimum amount of randomization is a single element reversal, i.e. $F_{\text{rev}} = 1/N_r$. In equation (10), C is the number of computations completed and x determines the departure from linearity of the annealing trajectory. If $x = 1$, the trajectory is linear. If $x < 1$, the value of F_{rev} initially decreases more quickly than a linear trajectory, leading to more cycles with less randomization. If $x > 1$, the value of F_{rev} initially decreases more slowly, leading to more cycles with higher randomization. The optimum value of the parameter, x , was found to be approximately 2 and this is the value used in the investigation. The annealing cycles continue until the computation limit, C_{lim} , has been reached, and the cycle with the lowest ε is chosen as the solution.

2.8. Approach 3: Hopfield network with image segmentation

Instead of using a random initial condition for each Hopfield network cycle, we introduce an alternative approach that uses non-random initial conditions. We base the choice for the initial condition on the observation that, in practice, the results of using approaches 1 and 2 tend to partition the target activity into contiguous regions of high neural activity that are achieved with the same generator signal sign. This qualitative empirical observation can be described as two principles:

Principle 1: two adjacent retinal locations of high neural activation that are not separated by lines of near-zero activation should be stimulated with generator signals of the same sign. In this way, the generator signal in the retina on the path between those two locations will not pass through zero; passing through zero would introduce undesired artifactual lines of near-zero neural activation (figure 2(c)).

Principle 2: two adjacent retinal locations of neural activation separated by lines of low neural activation should be achieved with generator signals of opposite sign. This forces the generator signal on the path between these two locations to pass through zero, thus creating the desired region of low neural activation (figure 2(b)).

It is apparent that these two principles describe conventional image thresholding followed by image segmentation. These are implemented for the target pattern (figure 3(a)) by initially binarizing it into high and low neural activity regions by thresholding at 33% of the maximum brightness level in that target pattern (figure 3(b)). Note that this threshold parameter is later explored for its effect on the results. The high neural activity pixels are grouped according to their direct adjacency to other high neural activity pixels using the function ‘bwconncomp’ in MATLAB (figure 3(c)). Regions with a low number of pixels, fewer than either 16 pixels or 0.5% of the total number of pixels, are removed (figure 3(d)). Low neural activity pixels are initially set to be +1 in the partition vector and high neural activity pixels are initially set either +1 or -1 using an exhaustive search procedure described below. (Note an entirely equivalent approach would be to reverse the sign so that low neural activity pixels are set to be -1 and high neural activity pixels are initially set either -1 or +1).

The advantage of the segmentation approach is that it is not necessary to search every possible combination of *pixel* sign, but instead test each possible combination of *region* sign. This reduces the exhaustive search space from 2^{N_r-1} to $2^{N_{\text{reg}}-1}$, where N_{reg} is the number of regions in the target retinal pattern, $\vec{r}^*$, separated by connected lines of low neural activity.

In general, this represents a dramatic reduction in the search space because there are typically fewer regions in the target activity pattern than elements in the target activity pattern vector (i.e. $N_{\text{reg}} \ll N_r$). In the example illustrated in figure 3, the search is reduced by a factor of $\sim 10^{10}$. Each combination is used as the initial condition of a single-cycle Hopfield network.

The principle of the segmentation search is to decide which high-activity regions should be achieved with opposite partition polarities. The combinations of regions are tested in a particular order (figure 3(e)) to ensure that the earliest combinations tested are those most likely to have the largest influence on ε and thus ensure these combinations are tested within the imposed computation limit. First, all high neural activity regions are set to partition vector values of -1 and the value of ε calculated, then the largest region is flipped to be $+1$ and the value of ε is recalculated. Then, the largest region is returned to -1 and the second largest segment is set to $+1$. This process continues until the second smallest region has been tested. There is no need to vary the signs in the partition vector associated with the smallest region as these degenerate combinations are redundant. Table 1 illustrates this process for a three-region pattern. If there are many combinations, it is possible that not all possibilities will be tested because, when the computation limit, C_{lim} , is reached, the process terminates. After termination, the partition vector associated with the lowest ε is chosen as the solution.

2.9. Model details and measuring computational performance

Unless otherwise stated, the following model considers a 7×7 electrode array with a generator signal spread with a standard deviation of 1.5 times the distance between electrodes ($\sigma = 0.21$). This was value was chosen to ensure that the generator signals from each electrode overlapped. Lower and higher generator signal spreads of 1 and two times the distance

Table 1. Illustration of a partition sign search order for a neural activity pattern with three separate bright regions discovered by the image segmentation process. The numbers in the table indicate the values used in the partition vector for all elements that are associated with that region.

Partition sign	Region 1 (Most pixels)	Region 2	Region 3 (Fewest pixels)
Trial 1	-1	-1	-1
Trial 2	+1	-1	-1
Trial 3	-1	+1	-1
Trial 4	+1	+1	-1

between the electrodes ($\sigma = 0.14$, and $\sigma = 0.29$) were also tested.

Single neural activity pattern analysis is completed with a 41×41 -pixel target neural activity pattern. Results across a set of images are completed with 1000 images from the CIFAR 100 database [26]. This database is used because it is low resolution, allowing for fast computation speed, and allows comparison of the partitioning algorithms across many images of different objects and environments.

Each partitioning algorithm is run for a set amount of central processing unit (CPU) time as calculated by the function ‘cputime’ of MATLAB. This is equivalent to 1.04×10^{10} CPU clock cycles. The maximum CPU time, C_{lim} , is set to 1.56×10^{11} clock cycles, because this is found to deliver sufficient number of Hopfield network cycles to see the impact of the search process, while also being low enough to allow for testing of 1000 neural activity patterns in a feasible timeframe. This computational limit is implemented as a soft boundary; the current Hopfield cycle is allowed to continue until it reaches its own termination condition.

Given the dimensions of the matrix $\mathbf{M}$ and vector $\vec{\varphi}_i$ each iteration of the Hopfield network process requires N_r^2 multiplication operations. This makes the problem of computing *each iteration* $O(n^2)$. This

applies to all three methods, the repeated Hopfield network, Hopfield network with simulated annealing, and Hopfield network with image segmentation. However, given that we impose a computational limit on the algorithm termination, increasing the number of pixels reduces the number of iterations possible within the time available for each algorithm. In general, this is expected to make segmentation more beneficial, given that it prioritizes salient solutions.

2.10. Image pre-processing

Three standard image pre-processing approaches were applied to the 1000 image CIFAR database to test their suitability when combined with image partitioning.

- (a) Each pixel below 33% of the maximum pixel intensity was set to 0 intensity. This emphasizes dark regions of the image increasing perceptual prominence.
- (b) Sobel edge detection was applied to the image using MATLAB's 'edge' function. This edge detected version was then superimposed on the original image. This increases the perceptual prominence of borders in the image.
- (c) All pixels below 33% intensity of the maximum intensity were set to 0 intensity and all pixels above this were set to maximum intensity. After edge detection was applied, this technique leads to the appearance of connected lines of the kind that the partitioning approach can create in the neural activation pattern. These lines were superimposed on the original image. This increased the perceptual prominence of borders in the image in a way that is particularly suited to the partitioning approach to stimulation.

In approaches B and C, the edges were emphasized as regions of low intensity so as to be suited to the lines of low-intensity created by the partitioning approach.

3. Results

The aim of each method is to find the optimal electrode stimulation currents by searching for the best partition of the target neural activity pattern into positive and negative generator signal values before applying the inverse LNL model. The goal of this investigation was to compare the methods based on the results they provide, as measured by the ε of the best solution found when the chosen computational limit, C_{lim} , been reached.

3.1. Hopfield network

In practice, there are many possible initial partitions for a given neural activity pattern, specifically 2^{N_r} .

In order to visualize the results of a Hopfield network search process, a 'toy' scenario with 3×3 target pixels and 2×2 electrodes is considered for illustrative purposes. The electrode array is modelled to have current spread wider than the electrode separation. A random target vector, $\vec{r}^*$, of length, $N_r = 9$, is used, giving a partition vector, $\vec{\varphi}$, of length 9. This random target image can be best achieved using a particular partition of the generator signal pattern into negative and positive regions. Given that $\varphi_i \in \{-1, 1\}$, there are $2^9 = 512$ combinations of values of the vector. Each of these is labelled with a unique number from 1 to 512. A single attempt to find this optimum pattern using the Hopfield network is completed and each step in this search illustrated (figure 4(a)). The Hopfield network begins with a particular guess (labelled node 181), then proceeds iteratively to make small changes to its initial guess each of which reduce the error associated with the solution. The search proceeds until it cannot find any more improvements (node 224). This solution may or may not be the optimum solution. The Hopfield network based on equation (9) is then run for every vector combination as initial condition and the trajectory of each run are mapped (figures 4(b) and (c)).

It can be seen that there are six groups, each made up of vertices that flow from nodes of high ε to nodes of lower ε , finishing at a local minimum. Thus, the groups can be considered to be basins of attraction, each associated with one of six local minima, of which two are global minima. The six groups are made up of three sets of mirrored pairs. This mirroring effect arises due to $\vec{\varphi}$ and $-\vec{\varphi}$ representing identical degenerate solutions to the inversion problem with inverted partition values.

The plots in figure 4(b) show the Hopfield network results numerically, with the trajectories of ε through each node shown as a function of the number of iterations of the Hopfield network for different initial conditions. The different plots in figure 4(b) correspond to the three distinct basins of attraction. Note that the first basin of attraction has the lowest ε of the three but can also take the longest to converge depending on the initial condition.

Figure 4(c) shows detailed results associated with the set of iterations associated with an example single cycle of the Hopfield network. Although a single cycle of the Hopfield network does discover nodes of lower error compared to the initial guess (or sometimes equal error), it does not reliably discover the best solution from a random initial condition but can easily become trapped in a local minimum. In cases where an exhaustive search is not possible and the Hopfield network is applied just once, this will often lead to poor outcomes. For this reason, optimization with only a single cycle of a Hopfield network was rejected.

3.2. Approach 1: repeated Hopfield network

It is possible to perform a repeated but non-exhaustive Hopfield network search (section 2.6). The Hopfield network minimization is repeated multiple times, with each referred to as a cycle. This is tested with a target neural activity pattern of 41×41 pixels and a 7×7 electrode array. Each begins with a random initial condition. Cycles are undertaken until the CPU clock cycles limit, $C_{lim} = 1.56 \times 10^{11}$, is reached (section 2.9) and then the current Hopfield cycle is allowed to continue to its own termination, and then the lowest ϵ outcome is picked as the solution.

This approach is applied using a neural activity target with three obvious regions (figure 5(a), top insert labelled 'Target'). The result of applying an unpartitioned LNL inverse model (figure 1(c)) can be seen to appear as a blurred version of the target (figure 5(a), bottom insert labelled 'unpartitioned'). In the time available, three cycles of the Hopfield network are completed. The results of each iteration during the three cycles are shown in the main plot, while their final neural activity patterns are shown in the inset at right of figure 5(a). The time gaps between the starts of the run and the starts of

the first Hopfield network cycles represent the set-up times of the search process. The result with the lowest ϵ is labelled '3'. The unusual regions in the resulting pattern are highly artifactual, which is a typical sign of a Hopfield network being trapped in detrimental local minima. These artefacts are similar to the Ising model phase transitions between magnetic or crystal domains as explained in the introduction and illustrated in figure 2(c). Avoiding these outcomes efficiently is one of the aims of this publication.

3.3. Approach 2: Hopfield network with simulated annealing

The repeated Hopfield network (figure 5(a)) approach effectively throws away information from the results of previous cycles. An alternative approach, described in section 2.7, is to use the lowest ϵ solution of previous cycles as the basis of a new initial guess for a new cycle of the Hopfield network (figure 5(b)). The initial guess is additionally perturbed by randomly flipping the sign of some values in the partition vector with the contribution of the randomization reducing during later cycles (equation (10)). This approach of reducing randomization under repeated trials is

Figure 5. Results for the three approaches for a 41×41 pixels activity patterns and a 7×7 electrode array. In each plot, ε is shown as a function of each iteration of a Hopfield network. The final outcome of each cycle is indicated by an open circular marker. The resulting neural activity patterns are shown as insets to the right of each plot along with their value of ε , where the best outcome is the highest numbered result. The horizontal dashed lines show the value of ε that results from using an unpartitioned approach. The vertical dotted lines show the CPU clock cycles cutoff, a soft limit of $C_{lim} 1.56 \times 10^{11}$. The time gaps between the starts of the run and the starts of the first Hopfield network cycles represent the set-up times of the search process. (a) The repeated Hopfield network approach. The two insets show the target pattern and the result of using the unpartitioned neural activity shaping approach. (b) The Hopfield network with simulated annealing approach. (c) The Hopfield network with image segmentation approach. The inset images show the image segmentations associated with each of the tested solutions (1–4). The positive and negative regions are shown in blue and red, respectively.

a form of simulated annealing. The same maximum number of CPU clock cycles, $C_{lim} = 1.56 \times 10^{11}$, was used. In this particular implementation with this target neural activity pattern, the simulated annealing approach gave $\varepsilon = 0.065$, an improved result over the repeated Hopfield network approach ($\varepsilon = 0.068$) within the time allowed (figure 5(b)). This can be confirmed by comparing the qualitative neural activity patterns resulting from the repeated Hopfield network (figure 5(a), ‘3’) with the Hopfield network with simulated annealing (figure 5(b), ‘4’). Figure 5(b) also illustrates that the amount of improvement over successive annealing runs reduces markedly for later runs.

3.4. Approach 3: Hopfield network with image segmentation

The Hopfield network with image segmentation approach is applied to the same three-region target pattern, as described in section 2.8. It can be seen in figure 5(c) that this results in a greatly reduced number of computational operations and, in this particular case, lower $\varepsilon = 0.026$. The solution, shown in the inset (figure 5(c), ‘4’), is also the obvious optimal outcome for this particular target.

It can be seen by examining the delays before the first iteration of the first cycle of the Hopfield network (figures 5(a)–(c)) that the additional image segmentation step did not significantly impact upon the results of the image segmentation approach. It can also be seen, by comparing the number of computations within each cycle of the Hopfield network, that each cycle of the image segmentation approach (figure 5(c)) is much faster than those associated

with the repeated Hopfield network (figure 5(a)) and the Hopfield networks with simulated annealing (figure 5(b)). This is because these networks are provided with an initial guess driven by image segmentation that is already close to a local minimum of the Hopfield network.

3.5. Comparison of the methods using a large image dataset

The results shown in figure 5 deal with only a single simple target image. The same process is completed for 1000 images taken from the test set of the CIFAR database [26]. A 7×7 electrode array and target neural activity pattern of 32×32 pixels are chosen for investigation. The CPU clock cycles limit was kept at the same value, $C_{lim} = 1.56 \times 10^{11}$. The minimal ε for each of the 1000 images is compared for the repeated Hopfield network and the Hopfield network with image segmentation (figure 6). The results are normalized to the ε associated with the unpartitioned approach (figures 6(a) and (b), black line, largely obscured in figure 6(a)).

The results across the image database are shown in figure 6. Figure 6(a) shows results for each algorithm ordered by ε of the image segmentation approach. Since algorithm choice for retinal implants is not performed on an image-by-image basis, but rather across a representative sample of salient images, the results are also shown in figure 6(b) ordered by ε for each algorithm independently. This allows for an algorithm-by-algorithm comparison. When the results are shown in these ways, it becomes obvious that image segmentation provided the best overall results but was not the best in all cases.

Of the 1000 neural activity patterns, all partitioning approaches gave the same unpartitioned result in 682 cases. Of the 318 neural activity patterns that were improved through partitioning, 169 were most improved using the image segmentation approach, 86 using the Hopfield network with simulated annealing approach, and 63 using the repeated Hopfield network approach.

The comparison of image segmentation with simulated annealing, the next best solution, is shown by calculating the ratio between ε for each target neural activity pattern resulting from simulated annealing vs. image segmentation. In figure 6(c), the images are ordered by the magnitude of the resulting ratio. A ratio greater than 1 indicates that, for a particular image, image segmentation provided a better outcome than the repeated Hopfield network. It can be seen that for those images that have lower ε for image segmentation, the benefit is often substantial with one image with a ratio greater than 2, while for the smaller number of images that had a lower ε for simulated annealing, the benefit was more limited with ε ratio up to 1.36.

The data presented in figure 6 provide quantitative evidence that, based on mean squared error, ε , image segmentation is a superior choice of algorithm. However, given that ultimately it is image perception that is important for retinal implant users rather than a numerical measure of pixel-by-pixel spike rate differences, it is relevant to look at qualitative results. Figure 7(a) shows a selection of three favourable examples of image segmentation with low values of ε relative to the unpartitioned approach. It can be seen that the particular target neural activity patterns for which image segmentation is particularly relevant are those target images with narrow connected lines of low intensity. Median outcomes and poor outcomes for the image segmentation approach are also shown (figures 7(b) and (c)).

3.6. Image segmentation parameter exploration

The results above relating to the image segmentation approach (figures 5–7) are for a particular set of parameters, namely: a threshold in image segmentation of 33%, electrode number of 49, and generator signal spread of $\sigma = 0.21$ (where the length of the electrode array is set to be 1). The value of these parameters may conceivably influence the results, and so the effect of varying them is examined by comparing the improvements in outcomes in the relative ε across the 1000 images. Figure 8(a) shows that there is an optimal range for the choice of target image threshold level between 25% and 40% of the maximum brightness level in the image. These selections produced the greatest number of images that were improved by the partitioned approach, while the choice of 10% or 50% reduced the numbers of improved images. Increasing the number of electrodes improves the outcomes for both the partitioned and unpartitioned strategies but reduces the benefit of partitioning relative to the unpartitioned strategy (figure 8(b)). This is because the low-activity lines induced by partitioning are more influential in improving low-resolution images. The extent of phosphene overlap, as defined by the spread of generator signal, does not influence ε relative to the unpartitioned solution under the same conditions (figure 8(c)).

3.7. Image pre-processing

Emphasizing features in images provides more opportunity for implant users to benefit from the low information provided by visual prostheses. Pre-processing is a standard approach in visual prostheses. By applying a set of standard image pre-processing strategies, the partitioning approach was tested for its suitability (figure 9).

In the context of pre-processed images, the value of the partitioning approach to stimulation is

dramatically increased (figure 9(d)). This outcome appears to be due to the fundamental property of the partitioning approach in representing isolated low-intensity regions of a target image (figures 9(a)–(c)).

Using pre-processing approach C, different computational limits were explored to discover the influence on the selection of the partitioning algorithm. It was found that image segmentation was the most effective strategy choice with lower computation limits (figure 10).

4. Discussion

This study compares methods for partitioning target neural activity patterns into regions that should optimally be achieved with either positive or negative generator signals. In the context of retinal implants, the aim of partitioning is to improve the image perceived by retinal implant users by introducing high contrast lines of low neural activity and, therefore, perception of low intensity lines. This goal was quantified by using the mean squared error, ε , between

the target pattern of neural activity and the expected pattern of neural activity as calculated using a LNL model.

The results in this paper indicate that image segmentation combined with a Hopfield network provides a large reduction in the time required to reach low ε solutions compared to other methods, including a repeated Hopfield network or a Hopfield network with simulated annealing. It was shown that, given a particular image and a fixed available computations limit, the image segmentation approach achieved lower ε solutions than the alternative approaches in many cases (figure 6(a)). Also, given a set of images, image segmentation was shown to comprehensively provide a better outcome in terms of the number of images that were improved (figure 6(b)). This result applies both to the fraction of images in a set that were improved with target partitioning, and the magnitude of the reduction in error achieved for the most improved images.

Visual prostheses do not typically use natural images directly as their target images. Instead they used pre-processed images that emphasize perceptual features such as the depth, intensity, or object edges [27–32]. The partitioning approach was found to be particularly effective when applied to pre-processed images that are ubiquitous in the field of visual prostheses. Given these results, it is expected that partitioning may be particularly appropriate for images pre-processed to emphasize depth by using darker or lighter image intensity combined with edge detection pre-processing to represent the distance to objects [33, 34].

The partitioning approach detailed in this study improves upon the methods detailed in a previous publication [11] that assumes that the target generator signal values, $\bar{g}^*$, are always positive. In the present study, by allowing for negative values of generator signal, the method provides a greater number of possible solutions.

Previous studies focus on the goal of creating virtual phosphenes [19–24]. The results of the neural activity shaping strategy presented in this study and a previous study [11] supersede the need for virtual phosphenes, or conventional phosphenes. Instead the method manipulates a pattern of neural activity as a whole. However, the approach can achieve a virtual phosphene outcome, if that is required, by using a phosphene as a target pattern of neural activity.

The mathematical problem of partitioning dealt with in this paper is more broadly known as the unconstrained binary quadratic programming problem. This problem arises in the context of a range of physics and optimization problems [35] and is NP-hard [36]. There is no known way to efficiently calculate an exact solution to these problems and the common approach is to apply

heuristic methods of the kind developed in this publication [35].

The problem scales with $O(n^2)$. This makes it important to consider the issue of reduced time availability (figure 10). We found that as the CPU time is reduced the image segmentation method becomes the preferred solution. It should be noted that the reduction in CPU cycles explored in figure 10, represents a reduction in MATLAB computation time from 15 s per image to 500 ms per image. In practice, the entire process must be concluded in the order of 30 ms, something that we expect can be achieved with optimized code running on an FPGA.

Although this study used all electrodes to simultaneously manipulate the entire pattern of neural activity, the method could be implemented semi-sequentially. This would be achieved by dividing the target image and electrode array into sub-sets which can be calculated and applied sequentially. This is a topic for future work.

In this study, the partitioning approach was combined with the neural activity shaping algorithm. It may alternatively be possible to use the partitioning algorithm in combination with the simpler conventional stimulation strategy. Under this approach, the image segmentation would be undertaken to decide the polarity of the target generator signal. The conventional approach, applied simultaneously, could then be used as usual. Experimental validation of the results of these approaches is an area for further work. The neural activity shaping strategy, either under the partitioned or the unpartitioned approach, is sensitive to the numerical measurement of the phosphene shapes contained in the inversion matrix $\mathbf{W}$ [11]. It is possible that this would not be less so when combined with the conventional down-sampled strategy because it does not require knowledge of the numerical values of the phosphenes. This is also an area for further work.

It is important to note that the three algorithms compared in this study were each computed as a serial process, with each additional cycle of the Hopfield network run after the previous one had finished. While this is a necessity for simulated annealing, where later cycles of the Hopfield network depend on the results of earlier cycles, it is not a necessity for the repeated Hopfield network or the image segmentation approach. These two approaches could be computed in parallel, leading to very large reductions in real time of computation, proportional to the number of threads available in the hardware used. This possibility does not change the results of this study; however, it does provide an opportunity for future work to investigate the further reduction in time available from parallelization.

Previous experimental results examining the neural response in retinal cells have found that most two-electrode interactions are linear [37]. These results suggest that the neural activity shaping strategy is

highly suited for use in retinal implants. In these studies, a small subset of the interactions were found to be highly nonlinear; an example is the neuron shown in figure 5(a) [37]. This cell appears to be entirely dominated by the level of one electrode or another, without any summation between the two. The neural activity shaping strategy, either unpartitioned or partitioned, is not capable of modelling the effects of these kind of cells. However, these cell types were a small minority of the total cell population investigated in that study. Note that the present study relies on experimental results that justify the choice of the LNL model [12, 13].

Given that the partitioned strategy has the capacity to induce complex patterns, it might be thought that the results are only applicable to arrays with large numbers of electrodes. However, this is not the case—the partitioning method is also applicable with low numbers of electrodes (figure 8(b)). This is because, with few electrode numbers, the alternative strategies result in very low resolution and, for certain target neural activity patterns with low activity lines, the resulting pattern could be dramatically improved by a partitioning approach. It may also be thought that the method only applies to devices with a low number of electrodes and not to high density electrode arrays. However, the benefits of the method depend on the existence of overlapping phosphenes and not the number of electrodes. Higher density electrodes are, by definition, more likely to have overlapping phosphenes, making partitioning more relevant to higher density electrodes than low density electrodes.

In this study, there appeared to be little difference between the repeated Hopfield network approach and the Hopfield network with simulated annealing approach when compared across a population of images. As mentioned above, this result would change in favour of the Repeated Hopfield network approach were it to be implemented in its parallel form. It is also possible that simulated annealing is better for target neural activity patterns where the global minima is influential across the whole search space. Speculatively, this situation may be associated with complex images with many separable regions in the target neural activity pattern.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgments

This project was supported by an Australian Research Council Linkage Project Grant (LP170101162). The authors thank Brian Gordon and Lindsay Powles from Bionic Vision Technologies Ltd, Matt Petoe from the Bionics Institute, and Kevin Meng from The

University of Melbourne's Department of Biomedical Engineering for their feedback on the partitioning approach to neural activity shaping.

ORCID iDs

Martin J Spencer <https://orcid.org/0000-0002-2001-9054>

Tatiana Kameneva <https://orcid.org/0000-0003-0081-8569>

Hamish Meffin <https://orcid.org/0000-0003-4307-6841>

References

- [1] Chuang A T, Margo C E and Greenberg P B 2014 Retinal implants: a systematic review *Br. J. Ophthalmol.* **98** 852–6
- [2] Tehovnik E J 1996 Electrical stimulation of neural tissue to evoke behavioral responses *J. Neurosci. Methods* **65** 1–17
- [3] Cogan S F 2008 Neural stimulation and recording electrodes *Annu. Rev. Biomed. Eng.* **10** 275–309
- [4] Margalit E et al 2002 Retinal prosthesis for the blind *Surv. Ophthalmol.* **47** 335–56
- [5] Weiland J D, Liu W and Humayun M S 2005 Retinal prosthesis *Annu. Rev. Biomed. Eng.* **7** 361–401
- [6] Beyeler M, Nanduri D, Weiland J D, Rokem A, Boynton G M and Fine I 2019 A model of ganglion axon pathways accounts for percepts elicited by retinal implants *Sci. Rep.* **9** 9199
- [7] Sinclair N C, Shivdasani M N, Perera T, Gillespie L N, McDermott H J, Ayton L N and Blamey P J 2016 The appearance of phosphene elicited using a suprachoroidal retinal prosthesis *Invest. Ophthalmol. Vis. Sci.* **57** 4948–61
- [8] Luo Y H-L, Zhong J J, Clemons M and da Cruz L 2016 Long-term repeatability and reproducibility of phosphene characteristics in chronically implanted Argus II retinal prosthesis subjects *Am. J. Ophthalmol.* **170** 100–9
- [9] Bonham B H and Litvak L M 2008 Current focusing and steering: modeling, physiology, and psychophysics *Hear. Res.* **242** 141–53
- [10] Dumm G, Fallon J B, Williams C E and Shivdasani M N 2014 Virtual electrodes by current steering in retinal prostheses *Invest. Ophthalmol. Vis. Sci.* **55** 8077–85
- [11] Spencer M J, Kameneva T, Grayden D B, Meffin H and Burkitt A N 2019 Global activity shaping strategies for a retinal implant *J. Neural Eng.* **16** 026008
- [12] Halupka K J, Shivdasani M N, Cloberty S L, Grayden D B, Wong Y T, Burkitt A N and Meffin H 2017 Prediction of cortical responses to simultaneous electrical stimulation of the retina *J. Neural Eng.* **14** 016006
- [13] Maturana M I, Apollo N V, Hadjicicolaou A E, Garrett D J, Cloberty S L, Kameneva T, Grayden D B, Ibbotson M R and Meffin H 2016 A simple and accurate model to predict responses to multi-electrode stimulation in the retina *PLoS Comput. Biol.* **12** e1004849
- [14] Rattay F 1999 The basic mechanism for the electrical stimulation of the nervous system *Neuroscience* **89** 335–46
- [15] Esler T B, Maturana M I, Kerr R R, Grayden D B, Burkitt A N and Meffin H 2018 Biophysical basis of the linear electrical receptive fields of retinal ganglion cells *J. Neural Eng.* **15** 055001
- [16] Chen S C, Suaning G J, Morley J W and Lovell N H 2009 Simulating prosthetic vision: I. Visual models of phosphene *Vision Res.* **49** 1493–506
- [17] Chen S C, Suaning G J, Morley J W and Lovell N H 2009 Simulating prosthetic vision: II. Measuring functional capacity *Vision Res.* **49** 2329–43
- [18] Wilke R G H, Moghadam G K, Lovell N H, Suaning G J and Dokos S 2011 Electric crosstalk impairs spatial resolution of multi-electrode arrays in retinal implants *J. Neural Eng.* **8** 046016
- [19] Eiber C D, Dokos S, Lovell N H and Suaning G J 2017 Multipolar field shaping in a suprachoroidal visual prosthesis *IEEE Trans. Neural Syst. Rehabil. Eng.* **25** 2480–7
- [20] Schmid E W, Fink W and Wilke R 2013. Simultaneous vs. sequential and unipolar vs. multipolar stimulation in retinal prostheses 2013 6th Int. IEEE/EMBS Conf. on Neural Engineering (NER) pp 190–3
- [21] Spencer T C, Fallon J B, Abbott C J, Allen P J, Brandli A, Luu C D, Epp S B and Shivdasani M N 2018 Electrical field shaping techniques in a feline model of retinal degeneration 2018 40th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society (EMBC) pp 1222–5
- [22] Dmochowski J P, Datta A, Bikson M, Su Y and Parra L C 2011 Optimized multi-electrode stimulation increases focality and intensity at target *J. Neural Eng.* **8** 046011
- [23] George S S, Shivdasani M N, Wise A K, Shepherd R K and Fallon J B 2015 Electrophysiological channel interactions using focused multipolar stimulation for cochlear implants *J. Neural Eng.* **12** 066005
- [24] Spencer T C, Fallon J B, Thien P C and Shivdasani M N 2016 Spatial restriction of neural activation using focused multipolar stimulation with a retinal prosthesis *Invest. Ophthalmol. Vis. Sci.* **57** 3181–91
- [25] Hopfield J J and Tank D W 1985 'Neural' computation of decisions in optimization problems *Biol. Cybern.* **52** 141–52
- [26] Krizhevsky A and Hinton G 2012 Learning Multiple Layers of Features from Tiny Images Technical Report (Toronto: University of Toronto)
- [27] Parikh N, Itti L and Weiland J 2010 Saliency-based image processing for retinal prostheses *J. Neural Eng.* **7** 016006
- [28] Asher A, Segal W A, Baccus S A, Yaroslavsky L P and Palanker D V 2007 Image processing for a high-resolution optoelectronic retinal prosthesis *IEEE Trans. Biomed. Eng.* **54** 993–1004
- [29] Boyle J, Maeder A and Boles W 2003 Scene specific imaging for bionic vision implants *Proc. 3rd Int. Symp. on Image and Signal Processing and Analysis, 2003. ISPA 2003* vol 1 pp 423–7
- [30] Buffoni L-X, Coulombe J and Sawan M 2005 Image processing strategies dedicated to visual cortical stimulators: a survey *Artif. Organs* **29** 658–64
- [31] Li H, Han T, Wang J, Lu Z, Cao X, Chen Y, Li L, Zhou C and Chai X 2017 A real-time image optimization strategy based on global saliency detection for artificial retinal prostheses *Inf. Sci.* **415–416** 1–18
- [32] Li H, Su X, Wang J, Kan H, Han T, Zeng Y and Chai X 2018 Image processing strategies based on saliency segmentation for object recognition under simulated prosthetic vision *Artif. Intell. Med.* **84** 64–78
- [33] McCarthy C, Barnes N and Lieby P 2011 Ground surface segmentation for navigation with a low resolution visual prosthesis 2011 Annual Int. Conf. IEEE Engineering in Medicine and Biology Society pp 4457–60
- [34] Jung J-H, Aloni D, Yitzhaky Y and Peli E 2015 Active confocal imaging for visual prostheses *Vision Res.* **111** 182–96
- [35] Kochenberger G, Hao J-K, Glover F, Lewis M, Lü Z, Wang H and Wang Y 2014 The unconstrained binary quadratic programming problem: a survey *J. Comb. Optim.* **28** 58–81
- [36] Pardalos P M and Jha S 1992 Complexity of uniqueness and local search in quadratic 0–1 programming *Oper. Res. Lett.* **11** 119–23
- [37] Jepson L H, Hottowy P, Mathieson K, Gunning D E, Dąbrowski W, Litke A M and Chichilnisky E J 2014 Spatially patterned electrical stimulation to enhance resolution of retinal prostheses *J. Neurosci.* **34** 4871–81

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Spencer, MJ;Kameneva, T;Grayden, DB;Burkitt, AN;Meffin, H

Title:

Neural activity shaping utilizing a partitioned target pattern

Date:

2021-08-01

Citation:

Spencer, M. J., Kameneva, T., Grayden, D. B., Burkitt, A. N. & Meffin, H. (2021). Neural activity shaping utilizing a partitioned target pattern. JOURNAL OF NEURAL ENGINEERING, 18 (4), <https://doi.org/10.1088/1741-2552/abecc4>.

Persistent Link:

<http://hdl.handle.net/11343/290390>

License:

[CC BY](#)