



Universiteit
Leiden
The Netherlands

Optimizing stimulus energy for cochlear implants with a machine learning model of the auditory nerve

Nobel, J. de; Kononova, A.V.; Briaire, J.J.; Frijns, J.H.M.; Back, T.H.W.

Citation

Nobel, J. de, Kononova, A. V., Briaire, J. J., Frijns, J. H. M., & Back, T. H. W. (2023). Optimizing stimulus energy for cochlear implants with a machine learning model of the auditory nerve. *Hearing Research*, 432. doi:10.1016/j.heares.2023.108741

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3607996>

Note: To cite this publication please use the final published version (if applicable).



Optimizing stimulus energy for cochlear implants with a machine learning model of the auditory nerve

Jacob de Nobel^{a,*}, Anna V. Kononova^a, Jeroen J. Briaire^b, Johan H.M. Frijns^{b,c}, Thomas H.W. Bäck^a

^a Leiden Institute of Advanced Computer Science, Niels Bohrweg 1, Leiden, Netherlands

^b Department of Otorhinolaryngology, Leiden University Medical Center, Albinusdreef 2, Leiden, Netherlands

^c Leiden Institute for Brain and Cognition, Wassenaarseweg 52, Leiden, Netherlands

ARTICLE INFO

Article history:

Received 8 November 2022

Revised 9 February 2023

Accepted 12 March 2023

Available online 14 March 2023

Keywords:

Cochlear implants

Optimization

Auditory nerve

Neural model

Machine learning

Evolutionary algorithms

ABSTRACT

Performing simulations with a realistic biophysical auditory nerve fiber model can be very time-consuming, due to the complex nature of the calculations involved. Here, a surrogate (approximate) model of such an auditory nerve fiber model was developed using machine learning methods, to perform simulations more efficiently. Several machine learning models were compared, of which a Convolutional Neural Network showed the best performance. In fact, the Convolutional Neural Network was able to emulate the behavior of the auditory nerve fiber model with extremely high similarity ($R^2 > 0.99$), tested under a wide range of experimental conditions, whilst reducing the simulation time by five orders of magnitude. In addition, a method for randomly generating charge-balanced waveforms using hyperplane projection is introduced. In the second part of this paper, the Convolutional Neural Network surrogate model was used by an Evolutionary Algorithm to optimize the shape of the stimulus waveform in terms of energy efficiency. The resulting waveforms resemble a positive Gaussian-like peak, preceded by an elongated negative phase. When comparing the energy of the waveforms generated by the Evolutionary Algorithm with the commonly used square wave, energy decreases of 8%–45% were observed for different pulse durations. These results were validated with the original auditory nerve fiber model, which demonstrates that the proposed surrogate model can be used as its accurate and efficient replacement.

© 2023 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Cochlear Implants (CI) are neuroprostheses that use direct electrical stimulation of the auditory nerve to restore a sense of sound and speech understanding to people with profound sensorineural hearing loss. A condition, which is estimated to affect roughly 5% of the global population (Olusanya et al., 2014). These devices consist of an external part containing a battery, a microphone, and a sound processor, which transforms sounds into coded signals that are sent via a wireless transmitter to an implanted component. The implanted electronics include a receiver and a stimulator that converts the coded signals to a pattern of electrical pulses that stimulate the auditory nerve via a multi-channel electrode array in the cochlea (inner ear).

The external part of the device has to be worn on the side of the head, behind the ear, and the social stigma (Rapport et al.,

2022) that is associated with wearing a hearing aid can be a potential barrier for use. In fact, the size and visibility of hearing aids have been identified as one of the main factors limiting adoption (David and Werner, 2015). For the CI, a large part of the external unit is taken up by the battery, and thus decreasing the implant's power consumption would lead to smaller batteries and, consequently, smaller devices.

In this work, we focus on neural stimulation of the auditory nerve, which has been reported to account for up to 90% of the total power consumed by the device (Yip et al., 2017). Specifically, the shape of the waveforms used to stimulate the auditory nerve fiber (ANF) is analyzed, whilst optimizing energy efficiency. In practice, the CI is often configured to use square biphasic waveforms, even though many studies show the potential benefit of using alternate shapes. For example, the introduction of an inter-phase gap (Carlyon et al., 2005) or the use of asymmetric pulses (Macherey et al., 2006) has been shown to lower the activation threshold. Moreover, it has been shown that the human auditory system is more sensitive to the anodic phase of bipha-

* Corresponding author.

E-mail address: nobeljpde1@liacs.leidenuniv.nl (J. de Nobel).

sic stimuli (Macherey et al., 2008; Undurraga et al., 2010). In some studies (Jezernik and Morari, 2005; Sahin and Tie, 2007), the reduced threshold energy of using exponentially decaying/increasing waveforms have been reported. Wongsarnpigoon et al. (2010) investigates square, exponential, and linear ramp waveforms for monophasic stimulation, but finds that none of these waveforms are simultaneously energy, power and charge optimal. In a later study, Wongsarnpigoon and Grill (2010) used an Evolutionary Algorithm¹ to optimize the shape of the stimulus waveform for energy and propose a truncated Gaussian pulse followed by a rectangular charge balancing phase. This methodology was used specifically for an ANF in Yip et al. (2017), where it was shown that an exponentially decaying cathodic phase followed by a rectangular anodic phase provides the lowest energy at threshold level.

In order to investigate the effect of electrical stimulation on the auditory nerve fibers in the cochlea, simulation with computer models can be used (Frijns et al., 2000; 1995; Hanekom and Hanekom, 2016). This can be an attractive alternative for performing experiments with CI, as *in vivo* experiments require human subjects, and a considerable effort to conduct. Early work in computational modeling with multi-compartment models yielded the activating function (Rattay, 1986; Warman et al., 1992), which is proportional to the second-order spatial derivative of the extracellular potential along the axon and can be used to approximate the influence of electrical stimulation on neural excitation. Contemporary models (Kalkman et al., 2022; Rattay et al., 2001; Smit et al., 2010) combine 3D modeling techniques with non-linear multi-node ANF models, and conceptually split up the simulation task into two processes: I) the calculation of the electrical field potentials in the geometry of the cochlea, and II) the simulation of the response of a nerve fiber to an externally applied potential field (Briaire and Frijns, 2000b). The increased complexity of these later models comes at a substantial computational cost, as they require the solution to a large set of differential equations to accurately simulate fiber behavior.

Surrogate modeling (Grefenstette and Fitzpatrick, 1985; Jin, 2011) is a technique used in engineering and optimization that uses an approximate model of a process for which the outcome is not easily measured or computed. This has the potential to massively speed up calculation time, as the evaluation of a surrogate is often orders of magnitude faster than the original process. In this work, a surrogate modeling approach is used to create a statistical approximate for the second of the two simulation processes of the model presented in Kalkman et al. (2022), the ANF model (Briaire and Frijns, 2005). With the surrogate, it becomes possible to do a much larger number of simulations than was previously possible with the original model, which allows researchers to perform experiments on a much larger scale.

This paper features two contributions. Firstly, five different machine learning models are compared in detail when tasked with constructing a surrogate for the active nerve fiber model used in Kalkman et al. (2022). This results in a Convolutional Neural Network (CNN) to be favored, which predicts the neural response with very high accuracy. Secondly, we draw inspiration from Wongsarnpigoon and Grill (2010) and Yip et al. (2017) and use the surrogate to optimize the stimulus waveform for energy with an Evolutionary Algorithm (EA). The surrogate effectively enables this type of analysis, as the large number of evaluations required to do the optimization would have been infeasible to do with the original model. A new method for generating charge-balanced waveforms is proposed, using hyperplane projection. This allows for the generation of arbitrary charge-balanced waveforms, whilst in previous works, the shapes that could be generated were




Fig. 1. Three-dimensional cut-through of a cochlea with a laterally implanted electrode array, as used by the volume conduction model. Nerve fibers are shown in yellow, and the electrode contacts can be identified in grey/black. Figure taken from Kalkman et al. (2022).

limited. This experiment suggests a potential alternative to the square biphasic pulses used in practice and indicates substantial energy savings.

2. Methods

2.1. Model of electrical hearing

This research builds on the work of Kalkman et al. (2022), where an updated version of a computational model of the implanted human cochlea is presented. This model consists of two underlying models, which will both be described in more detail here.

2.1.1. 3D Volume conduction model

The volume conduction part of the model uses a boundary element model (BEM) of the cochlea that can model arbitrary cochlear geometries, implanted with a multi-channel electrode array (Briaire and Frijns, 2000a). An example of such a geometry can be seen in Fig. 1. These geometries are characterized by a 3D mesh of polygons, which represent the boundaries between the different tissue types in the cochlea. When current is injected into the inner ear by the electrodes, this induces a potential distribution. The calculation of this distribution, termed the volume conduction problem, consists of finding the solution to the Poisson equation:

$$\nabla^2 \psi = -\frac{I}{\sigma}, \quad (1)$$

where ψ represents the potential distribution, I the injected current, and σ the conductivity of the medium. Because the cochlea does not consist of a heterogeneous medium (i.e., different tissue types have different levels of conductivity), every polygon in the mesh is associated with the levels of conductivity σ corresponding to the tissues it separates. In the model, fibers are represented by several cylindrical segments, for which at each segment the extracellular potential is calculated. This yields a vector of approximately 115 regularly spaced measurements of the electrical potential along the fiber. This vector is then normalized over the input current, which produces a transfer function with its values in kΩ. For simplicity, we refer to this transfer function by the term potential distribution in the remainder of the paper. The potential distribution induced by electrode e for a nerve fiber f is denoted with ψ_e^f , and can be compactly expressed in $\mathbb{R}^5$ with a fifth-order polynomial (see Appendix A).

¹ Described as a Genetic Algorithm (GA) by the authors.




Fig. 2. Schematic overview of the design of a surrogate model for the ANF model.

2.1.2. Active model of the auditory nerve fiber

The second part of the computational model architecture consists of an auditory nerve fiber model (Dekker et al., 2014; Kalkman et al., 2022) with human Schwarz-Reid-Bostock (SRB) kinetics (Schwarz et al., 1995), and calculates the neural response elicited by a stimulus. It is a non-linear double cable model developed to display fundamental mammalian nerve fiber properties, such as spike conduction velocity and repetitive firing behavior (Frijns et al., 2000). The model represents a human bipolar High Spontaneous Rate (HSR) fiber, which is the most prevalent type of primary auditory nerve fiber (Frijns et al., 2000) and is modeled in accordance with known human data (Briare and Frijns, 2005).

The ANF model $\mathbf{f}$ takes an arbitrary stimulus waveform $\phi_e(t)$, which is a discrete univariate time series of current I over a given experimental time window. Based on the potential distribution ψ_e^f provided by the 3D volume conduction model and a stimulus waveform $\phi_e(t)$, the model calculates the activation threshold $\hat{I}_e^f$:

$$\mathbf{f}(\mathbf{x}) = \hat{I}_e^f \in \mathbb{R}, \quad (2)$$

where $\mathbf{x} := (\psi_e^f, \phi_e(t))$. $\hat{I}_e^f$ represents the relative amount of current that must be injected by an electrode e using a waveform $\phi_e(t)$ for a fiber f to spike. In other words, if the waveform $\phi_e(t)$ is scaled by $\hat{I}_e^f$, an action potential can be observed. The computation of $\hat{I}_e^f$ requires the solution to a system of differential equations. This is solved using a classic backward Euler method, which takes a considerable amount of time to compute, and grows linearly with the length of $\phi_e(t)$.

2.2. Surrogate model of the auditory nerve fiber

The surrogate model is intended as a substitution for the ANF model described in the previous section, in order to reduce its substantial computation time. The goal is to construct a statistical approximate $\hat{\mathbf{f}}$ of the nerve fiber model $\mathbf{f}$, through the use of machine learning. By presenting the surrogate with a large number of training examples generated with the ANF model, the surrogate can be trained to function as its replacement. A schematic of this process is shown in Fig. 2.

2.2.1. Regression

Designing a surrogate for the ANF model can be postulated as a regression task, modeling the relationship between the independent model parameters $\mathcal{X}$ and the dependent model outcomes $\mathbf{y}$. Formally, the data set for training the surrogate model is defined as a set of input-target pairs $\mathcal{D} : -\{(\mathbf{x}_i \in \mathbb{R}^{85}, y_i \in \mathbb{R})_{i=1, \dots, n}\}$, where each input $\mathbf{x}_i$ consists of two entities: the potential distribution ψ_e^f computed by the volume conduction model and the stimulus waveform $\phi_e(t)$. The target value y_i equals the activation threshold $\hat{I}_e^f$.

A regression model is defined as a function $\hat{\mathbf{f}}_\theta(\mathcal{X}) \rightarrow \hat{\mathbf{y}}$ parameterized by θ , which outputs a predicted value $\hat{\mathbf{y}}$, representing activation thresholds ($\hat{I}_e^f$), for a given data set $\mathcal{X}$ of ANF model simulation parameters $(\psi_e^f, \phi_e(t))$. By providing the model with a set of training examples, we search for the regression model parameter

set θ , which causes it to predict a $\hat{\mathbf{y}}$ with a minimal error compared to the true value $\mathbf{y}$. This is evaluated using the coefficient of determination (R^2), which is used to indicate the “goodness of fit” for a regression model. Defined on $[-\infty, 1]$, R^2 denotes the proportion of the dependent variable ($\mathbf{y}$) which is correctly inferred from the independent variables ($\mathcal{X}$):

$$R^2(\mathbf{y}, \hat{\mathbf{y}}) = 1 - \frac{\sum_{i=1}^n (\mathbf{y}_i - \hat{\mathbf{y}}_i)^2}{\sum_{i=1}^n (\mathbf{y}_i - \bar{\mathbf{y}})^2} \quad (3)$$

Here, $\bar{\mathbf{y}}$ denotes the average value of $\mathbf{y}$. Note that a constant model which predicts the expected value for $\mathbf{y}$ yields an R^2 of 0, and a perfect model receives a value of 1. R^2 is calculated for the training, and test data sets separately, see also Section 3.2.

2.2.2. Regression models

As was mentioned in the introduction, five machine learning models are considered for the surrogate (regression) model:

1. **Polynomial Elastic Net (PEN)** is a polynomial regression model, that combines the $\mathcal{L}_1$ and $\mathcal{L}_2$ regularization factors of the lasso and ridge regression methods. Regularization reduces model complexity, effectively performing an in-place feature selection on the input data by penalizing the regression model for having large weights. This produces a sparse model, which can be especially useful when many of the input features are correlated (Zou and Hastie, 2005).
2. **Random Forests (RF)** (Breiman, 2001) are a type of ensemble learning that uses multiple decision trees, combining many weak classifiers to form a strong classifier through a process called bagging. Bagging involves randomly selecting subsets of data to train decision trees on, and predictions are made through majority voting over all the trees.
3. **Gradient Tree Boosting (GB)** is another type of ensemble learning which uses an ensemble of decision trees, similar to RF. In contrast to RF, the decision trees are created additively, where each tree is built to improve upon the already existing ones. This is achieved via gradient descent of a differentiable loss function specified by the user (Chen and Guestrin, 2016).
4. **Multi-Layer Perceptron (MLP)** is a fully connected feed-forward artificial neural network (Goodfellow et al., 2016). It consists of at least three layers of perceptrons or nodes, an input layer, several hidden layers, and an output layer. Each node in the network maps the weighted inputs to a scaled output: $z(\mathbf{w} \cdot \mathbf{x}_i) \mapsto \mathbb{R}$, which is then passed to the next layer of the network. The weights $\mathbf{w}$ of each node are obtained via the backpropagation of error terms in a form of gradient descent. In regression, the single-node output layer yields a continuous value.
5. **Convolutional Neural Network (CNN)** is a deep neural network, which sets itself apart from regular neural networks by the fact that they use a convolution instead of a matrix multiplication when processing data (Goodfellow et al., 2016). This allows CNN's to be highly effective at processing data with an inherent structure, such as images and audio signals, as they are invariant to the shifts to the input. The structure of the network is similar to an MLP, except that it has at least one convolutional layer as one of its hidden layers.

The selection of the models was based on the following considerations: Random Forests was included due to its favorable out-of-the-box performance and low requirement for parameter tweaking, which is a good indicator of the problem complexity. Gradient Boosting was added as it often demonstrates better performance than Random Forests. Deep Learning models were included because of their proven effectiveness in various machine learning tasks, and their non-linear nature fits the non-linear character of the ANF model. The NN-based models provide a universal function

Table 1

High-level overview of the data sets used in this work; **A**, **B** and **C**.

	A	B	C
electrode type	LW, PM	LW	LW
fiber health	H, ST, CD	H, ST, CD	H
# cochleae	5	1	1
# electrodes	16	1	1
# fibers	3 200	3 200	1
# pulses	1	1 296	48 000
# samples	1 466 189	12 441 600	48 000

approximation and can demonstrate better capabilities in out-of-sample predictions compared to the tree-based models that struggle with extrapolation. The CNN-based architecture was included due to its suitability for handling sequential data. The PEN model was also added as a simple indicator of the problem's complexity. Other popular models such as Support Vector Machines and k-nearest-neighbors were not included due to the large size of the data set, making them inefficient for this particular problem.

For PEN, RF, and MLP, the implementations as provided by the Python library SCIKIT-LEARN (Pedregosa et al., 2011) (version 0.24.2) were used. For the gradient boosting algorithm, the implementation from the library XGBOOST (Chen and Guestrin, 2016) (version 1.5.2) was used, and for the CNN, the KERAS (Chollet et al., 2015) API for TENSORFLOW (Abadi et al., 2015) (version 2.6.2) was used. All experiments were conducted using Python 3.7.5.

2.3. Simulated data

Data was collected by performing simulations with the ANF model for various experimental conditions. The research data includes potential distributions ψ_e^f for five different cochlear geometries (Kalkman et al., 2022), implanted with both lateral wall (LW) and perimodiolar (PM) electrode arrays, which are model equivalents of a HiFocus1J cochlear implant with 16 electrode contacts. Samples are collected for up to 3 200 nerve fibers (placed uniformly along three cochlear turns) with three levels of neural degeneration: i) healthy fibers (H), ii) "short-terminal" (ST) fibers, and iii) fibers with completely degenerated (CD) dendrites (Kalkman et al., 2015).

Three data sets are considered, for which an overview is given in Table 1. Each data set has different components, which allows for an investigation of the parameters of the ANF model from contrasting aspects. Data set **A** has the highest variability in potential distributions ψ_e^f , and includes data for several cochlear geometries and electrodes, but only considers a single stimulus waveform $\phi_e(t)$. A much larger number of stimulus waveforms are included in data set **B** for a single geometry and electrode contact. In contrast, data set **C** only varies the stimulus waveform and keeps all other parameters constant. Note that the number of samples for data set **A** is not a perfect multiple of the number of experimental conditions used to generate it. This is because simulations with parameters that were not able to yield a finite value for $\tilde{I}_e^f$ were excluded, as the maximum value for $\tilde{I}_e^f$ was limited to 100. Additionally, all data was scaled to zero mean and unit variance prior to training the surrogate model.

2.3.1. Stimulus

In practice, a stimulus waveform $\phi_e(t)$ is a univariate real-valued time series with a fixed sampling rate, representing the active current I in ampere (A) of an electrode e over a sampling window. Here, a time window of 160 μ s is considered with a sampling rate of 2 μ s. This allows for the representation of a stimulus waveform as a vector in $\mathbb{R}^{80}$, with each component bounded on $[-0.01, 0.01]$.

Table 2

An overview of the stimulus waveforms considered in the experiments for data set **B**, and the methods used to generate their variants. The proportion of samples in the data for a given base is also shown.

Name	Example	Extension method	Proportion
Biphasic	Fig. 3(a)	$\mathcal{GI}, \mathcal{E}, \mathcal{I}$	0.247
Asymmetric biphasic	Fig. 3(b)	$\mathcal{GI}, \mathcal{E}, \mathcal{I}$	0.517
Triangle	Fig. 3(b)	$\mathcal{GI}, \mathcal{E}, \mathcal{I}$	0.097
Linear ramp	Fig. 3(b)	$\mathcal{E}, \mathcal{I}$	0.052
Mirrored Linear ramp	Fig. 3(c)	$\mathcal{E}, \mathcal{I}$	0.046
Sinusoid	Fig. 3(c)	$\mathcal{I}$	0.009
Gaussian	Fig. 3(c)	$\mathcal{I}$	0.032

In standard operating conditions, a cochlear implant is often configured to produce square biphasic pulses (Macherey and Cazals, 2016; Macherey et al., 2006; Wongsarnpigoon and Grill, 2010; Yip et al., 2017), which can be adjusted parametrically in terms of the phase-duration, amplitude and polarity. An example of a cathodic-first biphasic pulse, with a phase duration of 50 μ s, a delay of 20 μ s and an inter-phase gap of 10 μ s is shown in Fig. 3(a). Nevertheless, the cochlear implant can be configured to transmit arbitrary stimulus waveforms, given that they are charge-balanced, i.e.:

$$\sum_{t=0}^{\infty} \phi_e(t) = 0 \quad (4)$$

This is to ensure that the implant does not cause tissue damage due to an unbalanced charge, which can introduce gases, corrosion products, and pH changes as byproducts that can be detrimental to the user (Huang et al., 1999).

In this work, all simulations are conducted using monopolar stimulation (Kalkman et al., 2022). Data set **A** (see Table 1) contains data for a single cathodic-first biphasic pulse, with a phase duration of 18 μ s, no inter-phase gap, and no offset, with simulations for all the cochlear geometries and electrode configurations.

For data set **B**, data was generated for pulses found commonly in the literature (an overview of these can be found in Fig. 3 and Table 2). For each of these pulses, variations were generated by scaling them to a different time duration, in 2 μ s intervals, up to a length of 160 μ s, using three different methods:

- **Gap-Insertion** ($\mathcal{GI}$) takes the base pulse and inserts an inter-phase gap at all zero-crossings of the pulse.
- **Extension** ($\mathcal{E}$) grows the pulse by copying every element, i.e:

$$(x_1, \dots, x_n) \mapsto (x_1, x_1, \dots, x_n, x_n).$$

- **Interpolation** ($\mathcal{I}$) computes the linear interpolation of the pulse at $2n$ points and replaces the original with the interpolation, i.e:

$$(x_1, \dots, x_n) \mapsto (x_1^a, x_1^b, \dots, x_n^a, x_n^b),$$

where x_i^a and x_i^b denote the interpolated points around the point x_i .

Starting with a base pulse x of length n , the $\mathcal{E}$ and $\mathcal{I}$ operators recursively increase the duration of the pulse up to the sampling window (160 μ s). Using these three operators, 648 different stimulus waveforms were generated. Each of these waveforms was then also considered with reversed polarity by multiplying it by -1, yielding a total of 1 296 pulses/waveforms. Simulations with the ANF model for these pulses (i.e. data set **B**) were only conducted for the HC3A geometry, and one electrode (electrode 9, i.e., in the middle of the array) due to the considerable amount of time in running these simulations (several weeks on a 64 machine compute cluster).




Fig. 3. Examples of the predefined waveforms used in this paper.




Fig. 4. Visualization of the charge-balancing transformation (Eq. (6)), in three dimensions. The non-violating hyperplane H is shown in blue, the original samples $\mathbf{x}$ in red, and the transformed samples $\mathbf{x}'$ in green. The dashed line indicates the applied transformation on each sample. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2.3.2. Generating charge-balanced waveforms

As previously mentioned, the stimulus waveform is represented by a real-valued vector in $[-0.01, 0.01]^{80}$. In order to satisfy the charge-balancing constraint (Eq. (4)), a randomly generated vector $\mathbf{x}$ can be projected to the non-violating hyperplane:

$$H := \{\mathbf{x} : \langle \mathbf{x}, \mathbf{1} \rangle = b\}, \quad (5)$$

where $b = 0$. The inner product between $\mathbf{x}$ and $\mathbf{1}$ is denoted by $\langle \mathbf{x}, \mathbf{1} \rangle$, and $\mathbf{1}$ represents a vector with all ones of the same size as $\mathbf{x}$. To ensure that every projected vector $\mathbf{x}'$ becomes charge-balanced and can be used as a stimulus waveform $\phi_e(t)$, each sampled vector $\mathbf{x}$ is projected onto H as follows:

$$\mathbf{x}' = \mathbf{x} - \frac{\langle \mathbf{x}, \mathbf{1} \rangle - b}{\|\mathbf{1}\|^2} \mathbf{1} \quad (6)$$

Figure 4 provides a visual intuition of this transformation for three-dimensional samples. From the figure, it can be seen that for every point in the three-dimensional space, an orthogonal transformation onto H can be computed, which satisfies Eq. (4).

2.3.3. Latin hypercube sampling

Since the distribution of stimulus waveforms used in data sets **A** and **B** have a relatively low variance, a Latin Hypercube Sampling (LHS) of $\phi_e(t)$ is used to get a data set with a broader range of

training examples for the surrogate model. This is labeled data set **C** (see Table 1).

Latin Hypercube sampling is a technique for generating quasi-random samples from a continuous multidimensional distribution, which ensures the drawn samples are distributed evenly in the sample space (Mckay et al., 1979). It involves dividing the cumulative density function (CDF) for each of the dimensions into N evenly spaced partitions, which are used to construct a multidimensional grid or hypercube. Random samples are then taken in such a way that only one sample is drawn for each axis-aligned hyperplane. For example, in a two-dimensional grid, this means that for every row and every column, there is only one sample (Tang, 1993). This produces a set of robustly stratified random samples, which, in contrast to traditional random sampling, provides the assumption that the samples reflect the true variability of the underlying distribution.

As it is computationally infeasible to simulate a large number of LH samples using the ANF model for all geometries and nerve fibers, this is only done for a single nerve fiber and electrode combination. In this way, only the stimulus waveform $\phi_e(t)$ is varied, while the potential distribution ψ_e^f remains fixed. Since the $\phi_e(t)$ and ψ_e^f are independent model parameters, it is expected that observations made from this data set will generalize to all fibers. The potential distribution induced by electrode 9 on fiber 1213 is used, which was chosen because this fiber has the lowest average threshold I_{th}^f for electrode 9, which is also included in data set **B**.

Data set **C** contains 8 000 quasi-random samples for stimulus waveforms $\phi_e(t)$ of length 80. In order to include shorter waveforms in the sample, an additional 2 000 samples are drawn for each length in $\{2, 6, \dots, 74, 78\}$. Each of these samples is then padded with zeros on the right in order to produce a vector in $\mathbb{R}^{80}$. This yields a total of 48 000 vectors in the sample for data set **C**. Note that these vectors cannot directly be used as stimulus waveforms, since because they are generated in a pseudo-random fashion, and do not satisfy constraint (Eq. (4)). However, by applying the transformation given by Eq. (6), each of the 48 000 random samples in data set **C** is charge-balanced.

3. Model selection

To establish which of the investigated machine learning models is best suited for surrogate construction, two different experiments are conducted, discussed in more detail in the following sections. The first is a hyperparameter tuning experiment on a small subset of the data set **A**, to find good (hyper) parameter sets for each of the investigated models. Then, using these parameter sets, model performance is evaluated through a large-scale cross-validation experiment on all the data sets (**A**, **B**, **C**). All experiments in this section are performed on a Linux server running CentOS 7, with 256 CPUs and 500 GB of RAM.

Table 3

Best values of the parameters of machine learning models investigated, found by the grid search.

Model	Parameter	Best value
Polynomial Elastic Net	DEGREE	3
	α	0.1
Multi-Layer Perceptron	$\mathcal{L}1$ -RATIO	0.1
	ACTIVATION	TANH
Convolutional Neural Network	HIDDEN_LAYER_SIZES	(500, 250, 100)
	N_FILTERS	512
Gradient Tree Boosting	KERNEL_SIZE	16
	POOL_SIZE	4
Random Forests	DROPOUT_RATE	0.2
	N_NODES_DENSE	100
Random Forests	γ	0.8
	LEARNING_RATE	0.1
Random Forests	MAX_DEPTH	14
	REG_{α}	0.1
Random Forests	REG_{λ}	12.8
	TWEEDIE_VARIANCE_POWER	1.2
Random Forests	N_ESTIMATORS	500
	MAX_DEPTH	30
Random Forests	MIN_SAMPLES_SPLIT	2
	MIN_SAMPLES_LEAF	1

3.1. Hyperparameter tuning

In order to determine the best settings for each of the tested machine learning models, hyperparameter tuning by means of a grid search is used. A grid search involves an iterative procedure that performs an exhaustive search over all combinations of parameter values specified by the user and returns the parameters which yield the highest performance. In order to get a robust estimate of the performance associated with a given parameter set, 3-fold cross-validation is used. This requires a large number of fitting operations for each of the models and would be very costly if evaluated on one of the complete data sets, due to their relatively large size. Thus, a down-sampled grid-search data set $\mathcal{D}_g$ is constructed from 1% of the samples in data set **A**, taken uniformly at random, which includes 14 661 samples. 33% of this data is held back for testing ($\mathcal{D}_g^{test}$) and perform the grid search on the remaining 67% ($\mathcal{D}_g^{train}$). Then for each model, the performance of the best parameter set found by the grid search is validated by testing on $\mathcal{D}_g^{test}$, using a model trained on $\mathcal{D}_g^{train}$.

Table 3 summarizes the list of searched parameters and their best values. For more detail on the parameters, we refer the interested reader to the SCIKIT-LEARN, TENSORFLOW and XGBOOST documentation (Abadi et al., 2015; Chen and Guestrin, 2016; Chollet et al., 2015; Pedregosa et al., 2011).

In Fig. 5, the results of the grid search experiments are shown. The figure shows two kinds of results. Firstly, it shows the distribution of the 3-fold cross-validation R^2 scores on $\mathcal{D}_g^{train}$ graphed as box plots for parameter sets of each considered model. Secondly, the R^2 score attained by each model when using the best parameter set on $\mathcal{D}_g^{test}$ is shown in the figure as circles. The precise value (R^2) for each dot is shown in the legend. The figure clarifies that the best parameter sets for all models show similar performance on this data set ($\mathcal{D}_g$), except for the PEN model, which seems to be less suited for this particular regression problem. Since PEN is the least complex model, this is to be expected. The highest score on the test set (0.904) was attained by the MLP model.

3.2. Validation

In this section, the hyperparameters sets which showed the highest performance (see Table 3) in the grid-search experiment are evaluated. Each model, with its best hyperparameter setting, is evaluated on the full **A**, **B** and **C** data sets (see Table 1) in this

Table 4

Results for 10-fold cross-validation for all data sets and machine learning models. Average over cross-validation folds for prediction quality metrics R^2 , MAE and MAPE are shown, in addition to the average fitting time in seconds. Best values attained for each experiment set are shown in boldface. Results of PEN on data set **B** are dropped due to the prohibitively large size of the data set.

	Metric	PEN	MLP	CNN	GB	RF
A	R^2	0.599	0.979	0.942	0.940	0.980
	MAE	0.950	0.177	0.312	0.304	0.137
	MAPE	0.277	0.045	0.079	0.075	0.034
	time (s)	783.85	24419.65	4830.11	43.23	202.56
	R^2	X	0.998	0.987	0.995	0.983
B	MAE	X	0.196	0.551	0.199	0.353
	MAPE	X	0.035	0.116	0.020	0.036
	time (s)	X	227128.67	232049.12	905.37	8819.14
	R^2	0.197	0.997	0.999	0.998	0.998
	MAE	4.870	0.188	0.147	0.272	0.310
C	MAPE	2.112	0.064	0.061	0.119	0.142
	time (s)	6518.32	302.36	956.77	37.53	137.46

experiment. This shows the true performance of each model on each data set. Moreover, since each of the three data sets is generated with different experimental conditions, performance recorded on $\mathcal{D}_g$ might not generalize to the other data sets. In order to assess how the results for each model will generalize to unseen data, cross-validation (Stone, 1974) is used, which is a re-sampling method that uses different partitions of the data to iteratively train and test a machine learning model. The goal of cross-validation is to evaluate the performance of a model on data that was not used to train it, providing insight into how it generalizes to an independent data set.

10-fold cross-validation With 10-fold cross-validation, each data set is partitioned into 10 equally sized groups or so-called folds, selected uniformly at random. Each group is sequentially held back for testing, while the machine learning model is trained on the union of the other groups. Subsequently, R^2 is computed for each fold by evaluating the model on the withheld test set. In addition to R^2 , model performance is evaluated using the Mean Absolute Error (MAE) and the Mean Absolute Percentage Error (MAPE) (see Appendix B). Note that while R^2 is to be maximized, both MAE and MAPE are to be minimized.

For each of these metrics, averages over the 10 cross-validation folds are shown in Table 4. The average training time is also given. From these results one can see that for data set **A**, the RF model outperforms all the other models. The R^2 of the RF model is close to the R^2 for the MLP model, but it has a lower MAE, which shows that its average error is of a lower magnitude. For data set **B**, the MLP shows the highest performance, and for data set **C** the CNN has the highest value for all recorded metrics. The performance of the PEN is rather poor compared to the other models. Additionally, it was not able to perform a fitting operation on the complete data set **B**, since the polynomial expansion of the input data causes the machine to go out of memory. Note also that the fitting times for both neural network models (MLP and CNN) are considerably higher than for the other models but show very good performance at the cost of an increased learning effort. Notably, most models (except for PEN & MLP) display the highest R^2 on data set **C**, which shows that the variance introduced by the LHS is beneficial in training a machine learning model. Overall, it can be observed that the models (except for PEN) show very high R^2 scores on all data sets, which indicates that they can emulate the behavior of the ANF model almost perfectly.

Leave-P-Groups-Out cross-validation The fact that data set **B** contains data for 1 296 unique stimulus waveform shapes, allows for an additional experiment to test a model's quality when predicting unseen waveforms. In the previous experiment, this was not explicitly enforced; thus, the machine learning models could poten-




Fig. 5. Grid search results. Each boxplot shows the distribution of the mean value of R^2 for every searched parameter set over a 3-fold cross-validation on the training data set $\mathcal{D}_g^{\text{train}}$. In addition, the validation score attained on $\mathcal{D}_g^{\text{test}}$ for the best-found parameter set for each of the models is shown as a circle, the precise values of these dots are given in the legend. Note that the y-axis is clipped at 0.

Table 5

Results for the Leave-P-Groups-Out with data set **B**, and the cross-prediction (predicting data set **B** with a model trained exclusively on data set **C**) experiments. Prediction quality metrics R^2 , MAE and MAPE are shown, in addition to the fitting time in seconds. Best values attained for each experiment set are shown in boldface. Note that the top row shows averages over a 10-fold cross-validation, while the bottom row only shows a single experiment.

	Metric	PEN	MLP	CNN	GB	RF
B _{grouped}	R^2	X	0.665	0.963	0.920	0.824
	MAE	X	0.752	0.805	0.926	1.433
	MAPE	X	0.186	0.138	0.093	0.152
	time (s)	X	178285.10	205791.83	626.32	7632.34
C → B _{f1213}	R^2	-0.991	0.954	0.974	0.692	0.366
	MAE	1.724	0.176	0.114	0.497	0.656
	MAPE	2.272	0.215	0.147	0.521	0.721
	time (s)	179.01	383.65	3181.94	43.99	148.52

tially still learn from similar data from the same waveform shape. This can be enforced by splitting the data into groups, where each group contains the complete data set for one stimulus waveform, and then performing a *Leave-P-Groups-Out* cross-validation experiment. In a similar fashion as with the previous 10-fold cross-validation experiment, 10 cross-validation folds are created, but now the folds are generated by using random subsets of groups, and not of individual samples. This ensures that every validation set contains only data from groups that were not used to train the model. As the grouping is based on waveform shape, this experiment can be used to answer the question: *If a model is generated using a set of known stimulus waveforms, how well can the model be expected to generalize to new waveforms?*

The top row of Table 5 shows the results of this experiment. From these results, a substantial decrease in model quality can be observed for all models, which is to be expected, due to the increased complexity of the experiment. However, with the RF model, this is more pronounced, which is probably due to the fact that RF models are known to fall short in extrapolation when predicting samples out of the domain of the training data (Zhang et al., 2017). Interestingly, all three other models are showing the best value for one of the metrics, with the highest R^2 (0.963) attained by the CNN model.

Cross-prediction Since data set **B** contains 1 296 samples for fiber 1213, and data set **C** exclusively contains samples for fiber 1213, a model trained on data set **C** can be used to predict the samples for fiber 1213 from data set **B**. Since there now only exists a single training (**C**) and validation set (**B**_{f1213}), only one experiment is conducted for each model, with no additional cross-

validation. This experiment is directed to the question: *Can a model trained on exclusively pseudo-random samples be used to predict data for waveforms used in clinical practice?*

Results for this experiment are shown in the bottom row of Table 5. From this, it can be seen that only the neural network (NN) based models are offering satisfactory performance in this experiment, with the CNN clearly outperforming the MLP model. Moreover, both models actually are showing better performance in this experiment than in the *Leave-P-Groups-Out* experiment. This gives rise to the notion that for training a neural network model (i.e. MLP or CNN) that generalizes to unseen data, the LHS-based data set **C** is better than data set **B**. Furthermore, the fact that the other models are not able to do well in this experiment, might mean that they are overfitting to the training data in the previous experiments, which was much more homogeneous with the test data than in this experiment.

Overall, the validation experiments presented in the previous paragraphs show that when only predicting samples from a single data set source, all models except for the PEN model show satisfactory performance. The GB model is the quickest to train but can be of lower quality than some of the other models tested. When testing for generalizability and evaluating how the models perform on data from other data set sources the CNN model was the most promising. *Therefore, for the experiments in the remainder of the paper, the CNN model will be used.*

Running time In Fig. 6 the running (prediction) time of the surrogate model is shown in comparison to the running time of the original ANF model. For the surrogate, the CNN model is shown, since this model showed the best performance in the previous ex-




Fig. 6. $\log_{10}$ scaled running time in seconds of both the ANF model f and the CNN Surrogate $\hat{f}$ vs. the duration of the pulse (a) and the number of evaluated samples (b). Samples of running time are generated using bootstrapping, and the average values are shown. The bands represent the standard deviation.

periments. From this figure, it can be observed that the surrogate model realizes a speedup of 5 orders of magnitude w.r.t the running time of the ANF model. In addition, it can be seen that while the running time increases linearly with the duration of the pulse for f it remains constant for $\hat{f}$.

4. Waveform optimization

The experiments in this section are based on the methodology presented in Wongsarnpigoon and Grill (2010), where a Genetic Algorithm (GA) is used to find energy-optimal stimulus waveforms. As an analytical solution for an energy-optimal stimulus waveform is not easily derived due to the complex non-linear nature of the nerve fiber model, heuristic optimization with a real-valued GA can be used.

A Genetic Algorithm (Holland, 1975) is a nature-inspired meta-heuristic, which seeks to evolve a population of candidate solutions to an optimization problem through an iterative procedure. By combining the properties of good solutions with small random changes, the GA can move towards better solutions over time. It ranks solutions based on an objective or fitness function $\mathcal{F}(\mathbf{s}) \mapsto \mathbb{R}$, which computes the quality of a candidate solution $\mathbf{s}$ relative to optimization problem the GA is solving. Historically, the GA is a method for solving discrete optimization problems, and the modifications made in Wongsarnpigoon and Grill (2010) for solving a continuous optimization problem, cause it to be closer to an Evolution Strategy (Beyer and Schwefel, 2002) (ES) than a GA. Therefore, in the remainder of the paper, the GA used by Wongsarnpigoon and Grill (2010) will be referred to by the overarching term Evolutionary Algorithm (EA) for correctness.

To save computational time, a surrogate model, previously introduced in this paper, is used to calculate the nerve fiber response instead of the computationally expensive ANF model. The used surrogate is a CNN model, retrained using all available data from $\mathbf{C} \cup \mathbf{B}_{f1213}$. With this experiment, we aim to demonstrate the applicability of the surrogate model as a stand-in replacement of the ANF model and ask the question: *Can the shape of the stimulus waveform used by the original ANF model be accurately optimized while performing the expensive simulations with the surrogate model?*

4.1. Experimental setup

In this work, candidate solutions are stimulus waveforms ($\mathbf{s} = \phi_e(t)$), which are vectors in $\mathbb{R}^d$, where each component represents the amplitude of the waveform at a given time step ($dt = 2 \mu\text{s}$).




Fig. 7. Schematic overview of the waveform optimization process with an Evolutionary Algorithm. The EA generates a new candidate solution $\phi_e(t)'$. Energy and the threshold for $\phi_e(t)'$ (calculated by the surrogate ANF model $\hat{f}$), are then combined into $\mathcal{F}(\phi_e(t)')$, which is fed back into the EA.

The fitness of a candidate solution (i.e., how good a candidate solution is) is based on the signal energy of the pulse $E(\phi_e(t))$, defined by the sum of the squared magnitude of the waveform signal:

$$E(\phi_e(t)) = dt \sum_{t=0}^{\infty} |\phi_e(t)|^2, \quad (7)$$

where dt is the size of the time step. Fitness is defined as $E(\phi_e(t))$ plus a penalty term P :

$$\mathcal{F}(\phi_e(t)) = E(\phi_e(t)) + P. \quad (8)$$

For each wave $\phi_e(t)$, the threshold value $\hat{f}_e$ is calculated with the surrogate model $\hat{f}$. If $\phi_e(t)$ is stimulating below threshold (i.e. $\hat{f}_e > 1$) the penalty term P is assigned to 1e5, which is several orders of magnitude greater than $E(\phi_e(t))$, and 0 otherwise:

$$P = \begin{cases} 1e5, & \hat{f}_e > 1 \\ 0, & \text{otherwise.} \end{cases}$$

A schematic overview of this process is shown in Fig. 7.

Experiments are conducted for the pulse durations: $\{20, 40, \dots, 140, 160\} \mu\text{s}$, which are encoded in the EA by real-valued vectors of dimensionality: $\{10, 20, \dots, 70, 80\}$. In order to satisfy the constraint in Eq. (4), the following strategies are investigated:

1. **Projection:** Candidate solutions are projected onto a non-violating hyper-plane though Eq. (6). This constricts the search space of the optimization algorithm to $d - 1$ but allows for non-biphasic pulses to be generated.
2. **Mirroring:** Only the (first) positive phase of the stimulus is optimized, and the second phase consists of the mirror opposite with respect to zero of the first phase to balance the charge.




Fig. 8. The optimal EA waveforms, grouped by pulse duration and constraint handling strategy. Each curve represents the best solution given by the EA for a trial and is scaled by $\hat{I}_i^0$ computed by the ANF model. For each subgraph, 10 trials are shown. Figure (a) shows the waveforms generated for the projected constraint handling strategy and Figures (b–d) resp. show the waveforms generated by the balanced_pre, balanced_post and mirrored strategies.

3. **Balancing:** Only the positive phase of the stimulus is optimized, and a rectangular pulse is used to balance the charge of the first phase. For the balancing rectangular wave, both preceding and ensuing pulses (referred to later here as “balancing_pre” and “balancing_post”) are investigated.

Note that for both balancing and mirroring, the optimization problem becomes considerably simpler since instead of d points, only $\frac{d}{2}$ points have to be optimized by the optimization algorithm.

The EA was implemented to the specifications given in Wongsarnpigoon and Grill (2010), with two-point crossover and where mutation of each candidate solution occurs via scaling with a random factor: $\mathbf{s}_i^t = \mathbf{s}_i^{t-1} \cdot \mathcal{N}(1, 0.025)$. For the projection constraint handling strategy, a scaling mutation operator would constrict pulses to the same polarity, which is why for this strategy, the mutation occurs not via scaling, but by shifting with a random factor drawn from a normal distribution, i.e.: $\mathbf{s}_i^t = \mathbf{s}_i^{t-1} + \mathcal{N}(0, 0.025)$. A population size of 50 was used, where




Fig. 9. (a) pulse duration vs. the value of $\mathcal{F}$ of the best solutions found by the EA computed using the surrogate, grouped by constraint handling strategy. (b) pulse duration vs. energy, scaled with $\hat{f}_e^f$ computed by the ANF model. (c) pulse duration vs. energy values, and the EA-generated waveforms (aggregated over all constraint handling strategies) are compared against the energy values of the predefined pulse shapes. Each line with dots shows the mean value over 10 trials, with the shaded area showing the standard deviation.

for each generation the top 10 candidate solutions were passed on to the next generation (elitism). For every combination of the aforementioned constraint handling strategies and each dimensionality, the EA was run with 10 independent trials for 10 000 generations. For every trial, the EA thus conducts a total of 400 010 evaluations of the surrogate.

4.2. Results

Looking at the shape of the pulses generated by the EA (see Fig. 8), some interesting observations can be made. It can be observed that, for pulse durations exceeding 40 microseconds, the positive phase of the generated pulses resembles a Gaussian curve, with an elongated balancing phase. This is consistent with the curves found in the study by Wongsarnpigoon and Grill (2010). The exception to this is the mirrored strategy, where due to the enforced symmetry both phases are resembling Gaussians. The shorter pulses have peaks that are notably sharper and are exponential in shape for some cases (e.g. 40 μs balanced_pre). The results are quite stable across trials, and most of the trial runs are all converging to similar shapes. This indicates that (local) optima are being found by the EA, and the optimization problem can be consistently solved. The least stable results are from the projection constraint handling strategy, but often only the peak offset is varying. As was mentioned before, the optimization task when using the projection constraint handling strategy is particularly more complex, due to the greater number of variables optimized. With this in mind, it might be that strategy needs a longer time to converge, and hence the less stable results. This increased complexity, however, comes with the added benefit that pulses generated by this strategy are not restricted to biphasic pulses. Interestingly however, the shapes this strategy generates, are still of that general shape, with only a few exceptions of tri-phasic pulses (e.g. for 120 μs). Moreover, the strategy seems to converge into shapes that are almost mono-phasic, where for longer pulse durations, the elongation is only used to balance the negative charge more evenly over a longer period. By definition, such shapes are also being generated by the balancing strategy, which shows that even though the balancing strategies are restricted to a biphasic shape, the shapes they converge to are close to the converged shapes of an unrestricted strategy. This means that the increased flexibility that is introduced by the projection strategy is not required for this specific problem, since the shapes that it converges to can also be

generated with a simpler strategy that uses balancing. Nonetheless, when comparing the pulses of the balancing pre and balancing post strategies, it can be seen that the projection strategy is actually able to ultimately find shapes that are similar to that of the better (lower energy) of these two strategies (e.g. 20 μs).

4.2.1. Validation

In order to validate the predictions made by the surrogate $\hat{\mathbf{f}}$ during the optimization process, the threshold value is also computed with the original ANF model. This is used to compute a validated energy value for each of the waveforms that resulted from the optimization, by scaling each of these waveforms by the threshold computed by the original ANF model $\mathbf{f}(\phi_e(t))$:

$$E^*(\phi_e(t)) = E(\mathbf{f}(\phi_e(t))\phi_e(t)) \quad (9)$$

In the results presented in the sequel, we refer to this validated energy value when discussing energy, and objective function value $\hat{\mathbf{f}}$ otherwise.

4.2.2. Energy efficiency

The objective function (left) and energy (right) values for the final solutions provided by the EA for each constraint handling strategy are shown in Fig. 9. When comparing Fig. 9(a) with the Fig. 9(b), a discrepancy between the actual energy value and the predicted value by the surrogate can be observed, and on average, the surrogate predicts a 6.61% lower threshold. This seems to be especially prevalent in the waveforms generated using the mirrored strategy, for which the surrogate tends to predict a lower threshold than the ANF model.

Figure 9 shows that, in general, the energy of the solutions decreases with pulse duration. The waveforms generated by balancing_pre and projection strategies are yielding the most energy-efficient pulses, while the mirroring strategy is less efficient. Balancing with a rectangular phase after the optimized phase is also resulting in pulses with lower energy than balancing before the optimized phase. Figure 9 (c) shows the union of all the constraint handling strategies as a single line in comparison with the energy values for the predefined pulses from Section 2.3.1. From this, it can be observed that for all pulse durations, the optimized waveforms are of lower energy, and the EA provides better solutions than the handcrafted pulses (data set B). This effect becomes even more pronounced with a longer pulse duration.




Fig. 10. The figure shows dimension vs. the energy value of the best solutions found by the EA in a series of 10 runs grouped by constraint handling strategy, relative to the energy of a square biphasic pulse. (a) shows the comparison with an anodic-first square wave, (b) shows the comparison with a cathodic-first square wave. Energy values are computed using the ANF model.

Figure 10 shows a comparison between the energy of the pulses generated by the EA and the energy of the typically used square wave of that same pulse duration, for both anodic-first and cathodic-first pulses. For each strategy, the solution with minimal energy over the 10 trials is shown. The figure shows a similar trend as can be seen in Fig. 9, in that the energy efficiency relative to a square wave tends to increase with pulse duration, even though the trend is not as pronounced. Interestingly, it seems that for shorter pulse lengths, anodic square pulses seem more efficient, while for longer pulses this trend reverses, which seems to be contradicting previous works (Kalkman et al., 2022; Macherey et al., 2008; Undurraga et al., 2010). For shorter pulse lengths ($< 20 \mu\text{s}$), the anodic-first square wave is already quite energy efficient, and only the balanced_post strategy can yield a substantially more energy-efficient pulse. Overall, it can be observed that the waveforms generated by the EA are of considerably lower energy for all pulse durations ($> 20 \mu\text{s}$). They on average have 31% lower energy than a cathodic-first and 41% lower energy than an anodic-first square wave of that same pulse duration.

5. Conclusion

In this work, a surrogate to the ANF model from Kalkman et al. (2022) was presented and used to conduct an optimization experiment. The surrogate model was able to predict the behavior of the ANF model with very high confidence, whilst reducing computation time by five orders of magnitude. Five different machine learning models were tested, of which the CNN model was the most promising and best able to extrapolate its prediction to unseen experimental conditions. The data was comprised of a combination of LHS samples and handcrafted simulation data, and it was found that the CNN model was able to accurately predict the handcrafted simulation data while having only trained on the (pseudo-random) LHS samples. With this study, we have aimed to show that through the use of a surrogate, a complex process such as the ANF model can be accurately approximated, allowing for large-scale experimentation. This is not restricted to this specific ANF model, however, and this methodology can be extended to other ANF models and processes. For example, the prediction of the time to an action potential or the detection of the node where the action potential starts.

The second part of the study used the CNN surrogate model with an EA to optimize stimulus waveforms for energy efficiency at threshold level, which suggests an alternate waveform to the commonly used square biphasic pulses. The waveform resembles a positive Gaussian curve offset by an elongated rectangular negative phase, which is in line with earlier research (Sahin and Tie,

2007; Wongsarnpigoon and Grill, 2010). The shape is in disagreement with the pulses generated in Yip et al. (2017), where decaying exponential waveforms were found to be the most efficient, but this might be due to the fact that the ANF model used here includes a 3D model interface, fewer restrictions on waveform shape were applied here. The shapes differ with varying pulse duration, and it was found that the energy efficiency of the waveforms generated by the EA increases with pulse duration. Compared to the commonly used square biphasic pulses, 8%–45% energy decreases were observed for pulse durations of 20–160 μs . Additionally, it was found that for shorter pulses, an anodic-first square wave is more energy efficient, while for longer pulses a cathodic-first pulse is better. The generated waveforms were also considerably more energy efficient than those from a large set of handcrafted pulses.

A flexible approach that uses projection for generating charge-balanced pulses is proposed, which can consistently produce the most energy-efficient pulses for each of the charge-balancing strategies tested. Interestingly, the shapes generated with this unrestricted strategy are close to the converged shapes of a restricted strategy, which was used in earlier research. This indicates that even while considering the full possible range of waveform shapes, the most energy-optimal shape is still a relatively uncomplicated biphasic waveform.

It must be noted that even though the modeling *in silico* suggests these potential energy savings, further work with in-patient testing is required to validate these results. This would then also consider the effects of capacity and the power dissipated at the electrode-tissue interface on the realized energy savings. Moreover, state-of-the-art analysis shows that the effect of an adaptive power supply can have a drastic impact on the energy efficiency of certain waveform shapes (Varkevisser et al., 2022), which is not taken into account here. Future work might focus on a wider range of pulse durations or the extension of this approach to other ANF models. Additionally, the benefits of using a surrogate modeling approach to full speech simulation might also be an interesting line of research.

Data availability

Data will be made available on request.

CRediT authorship contribution statement

Jacob de Nobel: Conceptualization, Methodology, Software, Investigation, Writing – original draft, Visualization. **Anna V. Kononova:** Conceptualization, Supervision, Writing – review & editing, Funding acquisition. **Jeroen J. Briaire:** Conceptualization, Supervision, Writing – review & editing, Funding acquisition.

Johan H.M. Frijns: Supervision, Writing – review & editing, Funding acquisition. **Thomas H.W. Bäck:** Supervision, Writing – review & editing, Funding acquisition.

Acknowledgments

We would like to thank Randy Kalkman for his vital support in collecting the simulation data, and Hao Wang for his advice and feedback. This work was supported by Health Holland, under the research project “TEMPORAL” in collaboration with Advanced Bionics.

Appendix A. Potential distributions

The potential distribution, represented by ψ_e^f , consists of approximately 115 real-valued measurements taken along a fiber f . These measurements reflect the potential energy induced by an electrode e and are measured at evenly spaced segments. The fact that not every fiber in the model has the same length, and thus not the same number of segments, causes the number of measurements to differ from fiber to fiber. An example of the potential distributions, induced by electrode 1 on 3 200 fibers modeled

using the HC3A (Kalkman et al., 2015) cochlear geometry is shown in Fig. A.1. As electrode numbers are ordered from base to apex, meaning this electrode is placed closest to the base of the cochlea, it can be seen that fibers located at the base (darker colors) show greater response to this electrode than fibers located more deeply in the cochlea.

Additionally, it can be observed from the figure that for each fiber the potential distribution follows the same curve, only the magnitude differs. From this observation, it follows that each of these curves can be approximated by a polynomial of degree k :

$$f_w(x) = \sum_{i=0}^k w_i x^i, \quad (\text{A.1})$$

parameterized by a set of real-valued weights w , obtained using e.g. least squares regression. By using the weights from this polynomial as the representation of ψ_e^f , the number of features for this parameter component is effectively reduced and standardized from ≈ 115 to the polynomial degree k .

In order to determine the appropriate degree d , the Root Mean Squared Error (RMSE) between the original potential distributions and their corresponding polynomial representations is computed




Fig. A.1. Electric potential distributions of 3200 Healthy fibers induced by electrode 1 taken from HC3A cochlear geometry, with PM electrode placement. The hue of a line represents the id of fiber, which is ordered from base to apex.




Fig. A.2. Average RMSE vs. the degree of the fitted polynomial with standard deviation displayed as error bars. The knee-point of the RMSE curve is shown in a dashed line.

for $d \in [1, 20]$ for data set **A**, which is shown in Fig. A.2. Note that the values shown are averages over all fiber and electrode combinations in the data. Subsequently, the knee point of the RMSE vs. degree curve is computed using the *Kneedle* (Satopaa et al., 2011) algorithm. As can be seen in the figure, this results in the value 5, which is used for the degree of polynomial representation of ψ_e^f .

Appendix B. Performance metrics

Mean Absolute Error (MAE) or $\mathcal{L}_1$ -norm, representing the expected value of the absolute error. It is defined on $[\infty, 0]$ as:

$$\text{MAE}(y, \hat{y}) = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (\text{B.1})$$

Mean Absolute Percentage Error (MAPE) represented the expected value of the absolute error of the prediction in proportion to the relative magnitude of the output. This makes the metric more sensitive to relative errors. It is defined on $[\infty, 0]$ as:

$$\text{MAPE}(y, \hat{y}) = \frac{1}{n} \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{\max(\epsilon, |y_i|)}, \quad (\text{B.2})$$

where ϵ is set to an arbitrary small value.

References

- Abadi, M., et al., 2015. TensorFlow: large-scale machine learning on heterogeneous systems. Software available from <https://www.tensorflow.org/>.
- Beyer, H.-G., Schwefel, H.-P., 2002. Evolution strategies—a comprehensive introduction. *Nat. Comput.* 1 (1), 3–52. doi:10.1023/A:1015059928466.
- Breiman, L., 2001. Random forests. *Mach. Learn.* 45 (1), 5–32. doi:10.1023/A:1010933404324.
- Briaire, J.J., Frijns, J.H.M., 2000. 3D mesh generation to solve the electrical volume conduction problem in the implanted inner ear. *Simul. Pract. Theory* 8 (1), 57–73. doi:10.1016/S0928-4869(00)00007-0.
- Briaire, J.J., Frijns, J.H.M., 2000. Field patterns in a 3D tapered spiral model of the electrically stimulated cochlea. *Hear. Res.* 148 (1), 18–30. doi:10.1016/S0378-5955(00)00104-0.
- Briaire, J.J., Frijns, J.H.M., 2005. Unraveling the electrically evoked compound action potential. *Hear. Res.* 205 (1), 143–156. doi:10.1016/j.heares.2005.03.020.
- Carlyon, R.P., van Wieringen, A., Deeks, J.M., Long, C.J., Lyzenga, J., Wouters, J., 2005. Effect of inter-phase gap on the sensitivity of cochlear implant users to electrical stimulation. *Hear. Res.* 205 (1), 210–224. doi:10.1016/j.heares.2005.03.021.
- Chen, T., Guestrin, C., 2016. XGBoost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794. doi:10.1145/2939672.2939785.
- Chollet, F., et al., 2015. Keras: deep learning for humans. <https://keras.io>.
- David, D., Werner, P., 2015. Stigma regarding hearing loss and hearing aids: a scoping review. *Stigma and Health* 1. doi:10.1037/sah0000022.
- Dekker, D.M.T., Briaire, J.J., Frijns, J.H.M., 2014. The impact of internodal segmentation in biophysical nerve fiber models. *J. Comput. Neurosci.* 37 (2), 307–315. doi:10.1007/s10827-014-0503-y.
- Frijns, J.H.M., Briaire, J.J., Schoonhoven, R., 2000. Integrated use of volume conduction and neural models to simulate the response to cochlear implants. *Simul. Pract. Theory* 8 (1), 75–97. doi:10.1016/S0928-4869(00)00008-2.
- Frijns, J.H.M., de Snoo, S., Schoonhoven, R., 1995. Potential distributions and neural excitation patterns in a rotationally symmetric model of the electrically stimulated cochlea. *Hear. Res.* 87 (1), 170–186. doi:10.1016/0378-5955(95)00090-Q.
- Goodfellow, I., Bengio, Y., Courville, A., 2016. *Deep Learning*. MIT press doi:10.5555/3086952.
- Grefenstette, J.J., Fitzpatrick, J.M., 1985. Genetic search with approximate function evaluation. In: *Proceedings of the 1st International Conference on Genetic Algorithms*, pp. 112–120. doi:10.5555/645511.657078.
- Hanekom, T., Hanekom, J.J., 2016. Three-dimensional models of cochlear implants: a review of their development and how they could support management and maintenance of cochlear implant performance. *Netw. Comput. Neural Syst.* 27 (2–3), 67–106. doi:10.3109/0954898X.2016.1171411.
- Holland, J.H., 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI doi:10.5555/531075.
- Huang, C.Q., Shepherd, R.K., Center, P., Seligman, P.M., Tabor, B., 1999. Electrical stimulation of the auditory nerve: direct current measurement in vivo. *IEEE Trans. Biomed. Eng.* 46 (4), 461–469. doi:10.1109/10.752943.
- Jezernik, S., Morari, M., 2005. Energy-optimal electrical excitation of nerve fibers. *IEEE Trans. Biomed. Eng.* 52 (4), 740–743. doi:10.1109/TBME.2005.844050.
- Jin, Y., 2011. Surrogate-assisted evolutionary computation: recent advances and future challenges. *Swarm Evol. Comput.* 1 (2), 61–70. doi:10.1016/j.swevo.2011.05.001.
- Kalkman, R.K., Briaire, J.J., Dekker, D.M., Frijns, J.H.M., 2022. The relation between polarity sensitivity and neural degeneration in a computational model of cochlear implant stimulation. *Hear. Res.* 415, 108413. doi:10.1016/j.heares.2021.108413.
- Kalkman, R.K., Briaire, J.J., Frijns, J.H.M., 2015. Current focussing in cochlear implants: an analysis of neural recruitment in a computational model. *Hear. Res.* 322, 89–98. doi:10.1016/j.heares.2014.12.004.
- Macherey, O., Carlyon, R.P., Van Wieringen, A., Deeks, J.M., Wouters, J., 2008. Higher sensitivity of human auditory nerve fibers to positive electrical currents. *J. Assoc. Res. Otolaryngol.* 9 (2), 241–251. doi:10.1007/s10162-008-0112-4.
- Macherey, O., Cazals, Y., 2016. Effects of pulse shape and polarity on sensitivity to cochlear implant stimulation: a chronic study in guinea pigs. *Physiol. Psychoacoust. Cognit. Normal Impaired Hear.* 133–142. doi:10.1007/978-3-319-25474-6_15.
- Macherey, O., Van Wieringen, A., Carlyon, R.P., Deeks, J.M., Wouters, J., 2006. Asymmetric pulses in cochlear implants: effects of pulse shape, polarity, and rate. *J. Assoc. Res. Otolaryngol.* 7 (3), 253–266. doi:10.1007/s10162-006-0040-0.
- Mckay, M., Beckman, R., Conover, W., 1979. A comparison of three methods for selecting vales of input variables in the analysis of output from a computer code. *Technometrics* 21, 239–245. doi:10.1080/00401706.1979.10489755.
- Olusanya, B.O., Neumann, K.J., Saunders, J.E., 2014. The global burden of disabling hearing impairment: a call to action. *Bull. World Health Organ.* 92, 367–373. doi:10.2471/BLT.13.128728.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cour-napeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Rapport, F., Lo, C.Y., Elks, B., Warren, C., Clay-Williams, R., 2022. Cochlear implant aesthetics and its impact on stigma, social interaction and quality of life: a mixed-methods study protocol. *BMJ Open* 12 (3). doi:10.1136/bmjopen-2021-058406.
- Rattay, F., 1986. Analysis of models for external stimulation of axons. *IEEE Trans. Biomed. Eng.* 33 (10), 974–977. doi:10.1109/TBME.1986.325670.
- Rattay, F., Lutter, P., Felix, H., 2001. A model of the electrically excited human cochlear neuron: I. Contribution of neural substructures to the generation and propagation of spikes. *Hear. Res.* 153 (1–2), 43–63. doi:10.1016/S0378-5955(00)00256-2.
- Sahin, M., Tie, Y., 2007. Non-rectangular waveforms for neural stimulation with practical electrodes. *J. Neural Eng.* 4 (3), 227. doi:10.1088/1741-2560/4/3/008.
- Satopaa, V., Albrecht, J., Irwin, D., Raghavan, B., 2011. Finding a “kneedle” in a haystack: detecting knee points in system behavior. In: *31st International Conference on Distributed Computing Systems Workshops*, pp. 166–171. doi:10.1109/ICDCSW.2011.20.
- Schwarz, J.R., Reid, G., Bostock, H., 1995. Action potentials and membrane currents in the human node of Ranvier. *Pflügers Arch.* 430 (2), 283–292. doi:10.1007/BF00374660.
- Smit, J.E., Hanekom, T., Van Wieringen, A., Wouters, J., Hanekom, J.J., 2010. Threshold predictions of different pulse shapes using a human auditory nerve fibre model containing persistent sodium and slow potassium currents. *Hear. Res.* 269 (1–2), 12–22. doi:10.1016/j.heares.2010.08.004.
- Stone, M., 1974. Cross-validatory choice and assessment of statistical predictions. *J. R. Stat. Soc. Ser. B (Methodological)* 36 (2), 111–133. doi:10.1111/j.2517-6161.1974.tb00994.x.
- Tang, B., 1993. Orthogonal array-based latin hypercubes. *J. Am. Stat. Assoc.* 88 (424), 1392–1397. doi:10.2307/2291282.
- Undurraga, J.A., Van Wieringen, A., Carlyon, R.P., Macherey, O., Wouters, J., 2010. Polarity effects on neural responses of the electrically stimulated auditory nerve at different cochlear sites. *Hear. Res.* 269 (1–2), 146–161. doi:10.1016/j.heares.2010.06.017.
- Varkevisser, F., Costa, T.L., Serdijn, W.A., 2022. Energy efficiency of pulse shaping in electrical stimulation: the interdependence of biophysical effects and circuit design losses. *Biomed. Phys. Eng. Express* 8 (6), 065009. doi:10.1088/2057-1976/ac8c47.
- Warman, E.N., Grill, W.M., Durand, D., 1992. Modeling the effects of electric fields on nerve fibers: determination of excitation thresholds. *IEEE Trans. Biomed. Eng.* 39 (12), 1244–1254. doi:10.1109/10.184700.
- Wongsarnpigoon, A., Grill, W.M., 2010. Energy-efficient waveform shapes for neural stimulation revealed with genetic algorithm. *J. Neural Eng.* 7, 046009. doi:10.1088/1741-2560/7/4/046009.
- Wongsarnpigoon, A., Woock, J.P., Grill, W.M., 2010. Efficiency analysis of waveform shape for electrical excitation of nerve fibers. *IEEE Trans. Neural Syst. Rehabil. Eng.* 18 (3), 319–328. doi:10.1109/tnsre.2010.2047610.
- Yip, M., Bowers, P., Noel, V., Chandrakasan, A., Stankovic, K.M., 2017. Energy-efficient waveform for electrical stimulation of the cochlear nerve. *Sci. Rep.* 7. doi:10.1038/s41598-017-13671-y.
- Zhang, H., Nettleton, D., Zhu, Z., 2017. Regression-enhanced random forests. *JSM Proc. Sect. Stat. Learn. Data Sci.* 636–647. doi:10.48550/ARXIV.1904.10416.
- Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B (statistical methodology)* 67 (2), 301–320. doi:10.1111/j.1467-9868.2005.00503.x.