Published in final edited form as:

J Neural Eng. 2017 February ; 14(1): 016014. doi:10.1088/1741-2552/aa52d1.

Particle Swarm Optimization for Programming Deep Brain Stimulation Arrays

Edgar Peña^{1,†}, Simeng Zhang^{1,†}, Steve Deyo¹, YiZi Xiao¹, and Matthew D. Johnson^{1,2,*}

¹Department of Biomedical Engineering, University of Minnesota

²Institute for Translational Neuroscience, University of Minnesota

Abstract

Objective—Deep brain stimulation (DBS) therapy relies on both precise neurosurgical targeting and systematic optimization of stimulation settings to achieve beneficial clinical outcomes. One recent advance to improve targeting is the development of DBS arrays (DBSAs) with electrodes segmented both along and around the DBS lead. However, increasing the number of independent electrodes creates the logistical challenge of optimizing stimulation parameters efficiently.

Approach—Solving such complex problems with multiple solutions and objectives is well known to occur in biology, in which complex collective behaviors emerge out of swarms of individual organisms engaged in learning through social interactions. Here, we developed a particle swarm optimization (PSO) algorithm to program DBSAs using a swarm of individual particles representing electrode configurations and stimulation amplitudes. Using a finite element model of motor thalamic DBS, we demonstrate how the PSO algorithm can efficiently optimize a multi-objective function that maximizes predictions of axonal activation in regions of interest (ROI, cerebellar-receiving area of motor thalamus), minimizes predictions of axonal activation in regions of avoidance (ROA, somatosensory thalamus), and minimizes power consumption.

Main Results—The algorithm solved the multi-objective problem by producing a *Pareto front*. ROI and ROA activation predictions were consistent across swarms (<1% median discrepancy in axon activation). The algorithm was able to accommodate for (1) lead displacement (1 mm) with relatively small ROI (9.2%) and ROA (1%) activation changes, irrespective of shift direction; (2) reduction in maximum per-electrode current (by 50% and 80%) with ROI activation decreasing by 5.6% and 16%, respectively; and (3) disabling electrodes (n=3 and 12) with ROI activation reduction by 1.8% and 14%, respectively. Additionally, comparison between PSO predictions and multi-compartment axon model simulations showed discrepancies of <1% between approaches.

Significance—The PSO algorithm provides a computationally efficient way to program DBS systems especially those with higher electrode counts.

* Address correspondence to: Matthew D. Johnson, Department of Biomedical Engineering, University of Minnesota, 7-105 NHH, 312 Church Street SE, Minneapolis, MN, 55455, phone: (612)-626-6492, fax: (612)-626-6583, john5101@umn.edu.

† Contributed equally to the work.

1. Introduction

Deep brain stimulation (DBS) therapy has shown tremendous promise and growth over the past decades as an invasive neurosurgical technique for treating numerous brain disorders [1]. The clinical success of DBS relies on both accurate implantation of one or more leads of electrodes into deep brain target(s), and identification of stimulation parameters that alleviate symptoms without inducing adverse side effects. One recent advance in the field of DBS that can address both factors is the development of DBS leads with electrodes distributed both along and around the shank of the lead [2]–[6]. With this increase in number and distribution of electrode sites, these so-called DBS arrays (DBSAs) expand the programming options for steering, shifting, and sculpting volumes of neural activation [3], [7]. Such functionality may be especially important when DBS leads are positioned in a brain region with a non-uniform target morphology [8], or when DBS leads are implanted in close proximity to nuclei or fiber pathways that, when stimulated, evoke adverse side effects [3], [5], [6].

However, increasing the number of independent electrodes creates the logistical challenge of identifying (or programming) the stimulation settings that optimize therapy for a patient. DBS leads with four cylindrical electrode contacts are known to require hours of meticulous and tedious trial-and-error programming in some cases [9], and the increase to eight [4] or thirty-two electrodes [2] is likely to create an intractable problem to optimize electrode configurations and stimulation parameters within a clinical setting. One solution to decrease the high-dimensionality of this programming problem is to construct computational neuron models of DBS that are fit to patient imaging data [10]–[13] and then apply efficient algorithms to identify stimulation parameters on a subject-specific basis [14], [15].

One previously developed algorithmic approach is based upon machine learning and training a classifier on thousands of computational neuron model simulations [14] that are based on biophysically realistic axon models [16]. These axon models are distributed across a range of orientations relative to a DBS lead, and finite element models are used to solve for the tissue voltage across a range of electrode configurations and stimulation settings [17], [18]. Geometrical features that describe the resulting volume or tracts of tissue activated from these neuron model simulations are then used to train a machine learning algorithm. The trained algorithm can then predict stimulation settings based on a new target volume of tissue that one wishes to activate. While comprehensive in formulation, the machine learning approach requires a large number of upfront simulations to train the classifier.

A more computationally efficient approach for automating the DBS programming process is to optimize a less complex model of neural activation without actually simulating multi-compartment neuron models. In such models of electrical stimulation, one typically estimates the stimulus-induced currents applied to each neuronal membrane compartment through an activating function [19], driving function [20], or weighted driving function [21] that is calculated from extracellular voltages obtained from solving finite element models. For example, Xiao and colleagues showed that one can maximize the sum of the activating function values within a region of interest using convex optimization to automate programming of DBS arrays [15]. Similarly, genetic algorithms have been developed for

programming stimulation settings through peripheral nerve cuff electrodes [22]. The challenge, however, with this overall approach is knowing how well the less complex model simulations actually compare with the more detailed multi-compartment neuron models that incorporate more biophysically realistic elements and parameters.

In this study, we formulate the problem of predicting neural activation within regions of interest and regions of avoidance as a non-convex, discrete-valued objective function with local minima. To solve such a problem, we developed a particle swarm optimization (PSO) methodology, which works by iterative exploration of the electrode configuration and stimulation amplitude parameter space. The PSO approach has been successfully applied in a number of optimization problems [23], which range from permutations [24], inversion of ocean color observations [25], training multi-layer neural networks [26], predicting tremor onset [27], and tracking human motion without markers [28]. In addition to implementing the PSO approach to solve the non-convex, threshold-based problem we have formulated, we also extend the optimization problem to a multi-objective one that optimizes for three separate clinically relevant objectives: (1) maximize activation of the therapeutic target volume, (2) minimize activation of side effect volumes, and (3) minimize overall power consumption.

2. Methods

2.1 Finite Element Modeling (FEM)

A three-dimensional finite element model (COMSOL Multiphysics v5.2) was developed for a DBS array [29], which consisted of 32 elliptical electrodes (0.53 mm major axis, 0.3 mm curved minor axis, 0.1 mm thick) arranged in eight rows and four columns along the cylindrical lead (0.5 mm diameter). Conductance values for lead insulation ($\sigma = 1 \times 10^{-12}$ S/m) and electrodes ($\sigma = 1 \times 10^6$ S/m) were set according to a previous model from our group [15], approximating the conductance of silicon carbide and polyimide insulation and conductive platinum electrodes, respectively. Though more complex tissue conductance models of DBS have been developed [30]–[33], for the purposes of demonstrating the PSO algorithm, we assigned simple isotropic conductance values to the encapsulation layer (0.1 mm thick; $\sigma = 0.18$ S/m) [34] and to the bulk tissue (100 mm diameter; $\sigma = 0.3$ S/m) [35]. Quadratic tetrahedral mesh elements were generated by Delaunay triangulation with variable resolution mesh refinements set such that further refinement of the mesh yielded less than 5% changes in the activating function measure. The resulting mesh consisted of 4,104,421 domain elements, 204,990 boundary elements, and 12,708 edge elements [15], [36].

While there are myriad number of electrode configurations and stimulation amplitudes that could be modeled, we leveraged the principle of superposition and solved the Poisson equation in COMSOL to calculate the tissue voltage map individually for each of the 32 electrodes. In this case, for each simulation, a single electrode was set as the boundary current source with a cathodic current of 1 mA (current density 8.63×10^3 A/m²) and the outer surface of the bulk tissue was set to ground. These single boundary current source simulations resulted in tissue voltage maps, which were then scaled and superimposed to generate finite element modeling solutions for more complex electrode configurations with independent current sources.

2.2 Predictions of Thalamocortical Axon Model Activation

2.2.1 Motor and Somatosensory Thalamus Reconstructions—For this study, we modeled stimulation targeting the rhesus macaque cerebellar-receiving area of motor thalamus (VPLo: ventral posterolateral pars oralis) [15], which is the homologue of the ventral intermediate nucleus (Vim) in humans and primary DBS target for treating essential tremor. Of note, clinical DBS implants targeting this nucleus can be difficult to program because of low-threshold side effects such as paresthesia, which are thought to result from the spread of stimulation into the somatosensory thalamus (VPLc: ventral posterolateral pars caudalis) [6], [37]. Surface reconstructions of the region of interest (ROI; in this case, VPLo thalamus) and the region of avoidance (ROA; VPLc thalamus) were generated from the rhesus macaque brain [15], [38]. A DBSA was placed within the VPLo volume at 77° above the horizontal plane and at 10° from the sagittal plane in an anterior to posterior trajectory (Fig. 1A). VPLo and VPLc volumes were populated with 4,549 and 5,937 thalamocortical axon models (0.2 mm internodal spacing, with each axon arranged in a grid 0.72 mm away from adjacent neighbors) with simplified linear trajectories running from ventromedial to dorsolateral at 45° from the axial intercommissural plane (Fig. 1B). The linear trajectory and 45° angle were approximated from a previous investigation of efferent thalamic fibers in non-human primates [39]. Axon models that overlapped the DBSA were removed from subsequent analysis.

2.2.2 Modeling Myelinated Axon Activation—One way to estimate axonal activation resulting from extracellular stimulation is to solve the inhomogeneous cable model equation of a myelinated axon [40]–[42]:

$$\lambda^2 \frac{\partial^2 V_m}{\partial x^2} - \tau_m \frac{\partial V_m}{\partial t} - V_m = -\lambda^2 \frac{\partial^2 V_e}{\partial x^2} \quad (1)$$

where λ is the axon space constant that is dependent on the axon dimensions and geometry, V_m is the membrane voltage, τ is the time constant of the axonal membrane, and V_e is the interpolated extracellular potential calculated from the finite element model solution. For simplicity, the source term can be approximated with the difference approximation:

$$\frac{\partial^2 V_e}{\partial x^2} \approx \frac{V_{e,n-1} - 2V_{e,n} + V_{e,n+1}}{\Delta x^2} \quad (2)$$

where $V_{e,n}$ corresponds to the extracellular voltage at node n , and x is the distance between adjacent nodes of Ranvier. This source term or “activating function” has been shown to provide a reasonable approximation of the non-faradaic transmembrane currents that result from the initial onset of an extracellular stimulation pulse [19], [40].

Here, we use a modified activating function (MAF) to predict axonal activation. We compute the second spatial difference using extracellular potentials from non-adjacent nodes of Ranvier:

$$MAF = V_{e,n-2} - 2V_{e,n} + V_{e,n+2} \quad (3)$$

Notably, this approach is fundamentally the same as the activating function, except that it yields a smoother version of the activating function. Because this modified spatial difference is a linear function of the extracellular potentials, the superposition principle can be used to efficiently predict MAF values for an arbitrary electrode configuration [15]. For each axon, we constructed an N-by-32 matrix (denoted as the “C matrix”) containing the MAF value for all N nodes of Ranvier when stimulating through each electrode individually. This enabled computing the MAF value at every node for arbitrary electrode configurations by multiplying the C matrix with the 32-by-1 vector of currents going through each electrode of the DBSA.

2.2.3 Defining MAF Thresholds for Axonal Activation—An axon was considered activated if, for a given electrode configuration, the MAF value exceeded a predefined threshold (MAFT) at one or more of its nodes. Because the choice of MAFT is dependent on the specific axonal geometry, we tuned our MAFT empirically to maximize its predictive accuracy for our thalamic fiber geometry. We simulated the axonal fiber geometries as multi-compartment myelinated axon models in NEURON (v7.4). Axon models consisted of 2 μm diameter fibers with compartments representing nodes of Ranvier, myelin attachment segments, paranode main segments, and internode segments connected through an axial resistance [16]. Axonal membrane compartments were each driven using the extracellular mechanism in NEURON (e_extracellular) [6], [8], [43]. We applied a waveform with a 90 μs cathode-leading phase, 400 μs interphase delay, 3 ms charge-balanced anodic phase, and 135 Hz pulse rate [44]. The specific extracellular potential localized to each axonal node from current-controlled DBS inputs (0 to -1 mA per electrode) was estimated with the FEM and superposition. After running 30 primer simulations with a range of electrode configurations, we computed a best fitting MAFT value by minimizing the mean squared error between MAFT predictions and the 30 primer NEURON simulations. The obtained MAFT value (0.0023) was used for all subsequent PSO runs.

2.2.4 Estimating Axonal Activation—The MAFT value was then used to construct functions that predicted the number of axons activated in ROI, $R(x)$, and the number of axons activated in ROA, $S(x)$. These threshold-based functions were described as follows:

$$R(\vec{x}) = \sum_{i=1}^U H(\max(C_{ROI,i} \vec{x}) - \alpha) \quad (3)$$

$$S(\vec{x}) = \sum_{j=1}^V H(\max(C_{ROA,j} \vec{x}) - \alpha) \quad (4)$$

where x is a vector of size 32 corresponding to the current through each electrode; $C_{ROI,i}$ and $C_{ROA,j}$ are the C matrices for ROI axon i and ROA axon j , respectively; $H(\bullet)$ is the

Heaviside function; and α is the MAFT value. Using a MAFT value of 0.0023 (obtained as described in Section 2.2.3), the discrepancy between MAFT-based predictions and NEURON model predictions was less than 1% of the total axons in either ROI or ROA (Fig. 2). Given this low discrepancy, the MAFT value and the functions $R(x)$ and $S(x)$ were used in all subsequent PSO simulations.

2.3 Optimization Problem

For the simplified thalamocortical axon geometry modeled, the optimization algorithm's objectives were to (1) maximize the number of activated axons in the ROI, (2) minimize the number of activated axons in the ROA, and (3) minimize the power dissipated by the stimulator. These three distinct objectives reflect the desired clinical outcome, in which robust therapy is delivered with little to no side effects and with low power consumption. Additionally, we constrained the stimulus waveform through each electrode contact to have a cathode-leading phase (between 0 to -0.5 mA), which was within the current density safety limits based on the electrode size. Conceptually, this optimization problem can be denoted as follows:

$$\text{maximize } R(\vec{x}) \quad (5)$$

$$\text{minimize } S(\vec{x}) \quad (6)$$

$$\text{minimize } P(\vec{x}) = \sum_{k=1}^{32} (x_k^2) \quad (7)$$

$$\text{subject to } -0.5 \leq x_k \leq 0, k=1, \dots, 32 \quad (8)$$

where $P(x)$ is the power in mA^2 (not scaled by impedance).

To solve this multi-objective optimization problem, we employed a standard approach known as linear aggregation, which involves creating a single scalar objective function from the weighted sum of the three distinct objectives. We chose a weighting based on our ranking of importance of the objectives:

$$\text{minimize } -R(\vec{x}) + (2)S(\vec{x}) + (0.5)P(\vec{x}) \quad (9)$$

$$\text{subject to } -0.5 \leq x_k \leq 0, k=1, \dots, 32 \quad (10)$$

However, one limitation of the linear aggregation approach, from the clinical perspective, is that it assumes that there is one correct prioritization between the three objectives. This limits the user who seeks to prioritize outcomes in a different, subject-specific manner. Ideally, we would be able to solve the optimization problem for every possible combination of weightings. To efficiently obtain solutions for a range of weighting combinations, we mixed the aggregation-based method with elements of a Pareto dominance-based method. Specifically, we leveraged the fact that exploratory optimization algorithms like particle swarm optimization evaluate many intermediate solutions in search of the “best” solution. We collected these intermediate solutions, and labeled as ‘Pareto dominated’ any candidate solution for which there was at least one ‘better’ solution (i.e. better in all three objectives). By collecting these intermediate solutions and getting rid of solutions that were Pareto dominated, we were able to construct a set of solutions known as the *Pareto front* (Fig. 3). Solutions on that Pareto front are considered *Pareto optimal* because they were not Pareto dominated by any other solution. As such, each solution on the resulting Pareto front represents the best electrode configuration possible given the specified tradeoffs. The Pareto front obtained from such a mixed approach enables the user to readily select among multiple Pareto-optimal electrode configurations without having to directly solve the problem for different weighting combinations. A schematic flow of this approach is illustrated in Figure 4.

2.4 Particle Swarm Optimization

In this study, axons were considered activated when MAF exceeded a certain threshold (MAFT) at any point along the axon. Mathematically, this type of discrete function does not meet the criteria for convex optimization techniques. Therefore, to solve this problem without relaxing the parameters of the problem, we implemented a particle swarm optimization (PSO) approach in MATLAB (v2014b).

2.4.1 PSO Behavior—PSO solves optimization problems through a series of searches performed by a collection of interacting individuals [45]. Each individual (particle) exhibits three simple behaviors that enable it to search through the solution space, and to communicate its findings to neighboring individuals (swarm). The most fundamental of these behaviors is the particle’s persistent and somewhat random movement through solution space, which enables it to explore different potential solutions over time (i.e. inertia). The second behavior is the particle’s tendency to move toward a point that the swarm considers to be the best so far (i.e. the “social” or “global” best). The third behavior is the particle’s tendency to move towards the best point the particle itself has found (i.e. the “cognitive” best).

Mathematically, a particle’s position in the solution space at a given time can be described as follows:

$$x_i(t+1)=x_i(t)+v_i(t+1) \quad (11)$$

$$\text{where } x_i(0) \sim U(X_{min}, X_{max}) \quad (12)$$

where x is the position (i.e. electrode current) bounded between X_{min} and X_{max} , i is the particle number, t is iteration, and v is the “velocity” term (i.e. change in current). The three simple behaviors of the particle are included in this velocity term by three components: inertial, social, and cognitive [46]. Hence, velocity is defined as:

$$v_{ij}(t+1) = w \cdot v_{ij}(t) + c_1 \cdot r_{1ij}(t) \cdot [y_{ij}(t) - x_{ij}(t)] + c_2 \cdot r_{2ij}(t) \cdot [\hat{y}_j(t) - x_{ij}(t)] \quad (13)$$

where j is the electrode number, w is the inertial component, c_1 is the cognitive component, c_2 is the social component, y is the particle-specific best as of iteration t , and $\hat{y}$ is the global best as of iteration t . Meanwhile, r_{1ij} and r_{2ij} are additional weighting factors that allow for randomness, which promotes exploration. In our implementation, we set r_{1ij} to be randomly distributed between 0 and 1 (inclusive), while r_{2ij} was set to 1 for all i and j . As such, the random component of motion was solely due to the cognitive component. Finally, we defined w as follows:

$$w(t) = -\frac{w(0) - w(n)}{n}t + w(0) \quad (14)$$

where n is the maximum number of generations. We set w such that it decreased linearly from 0.9 to 0.4 in order to promote initial exploration followed by a tendency to converge toward the end [46] (Table 1).

2.4.2 Cognitive vs. Social Velocity Components—PSO exploration is driven by social and cognitive components, such that the choice of c_1 (cognitive) and c_2 (social) significantly influences the effectiveness of exploration in a problem-specific way. To determine appropriate c_1 and c_2 values for our thalamic geometry, we performed a parameter sweep by running the PSO algorithm five times for every possible combination of c_1 and c_2 across a range of values (0.025, 0.05, 0.1, 0.2, 0.4, 0.9, 1.8, 2.7, 3.6, 4.5, 5.4, and 6.3). We then assessed each combination based on (1) the optimality and consistency of its objective function value (i.e. mean and standard deviation of objective function value), and (2) its termination status (i.e. convergence vs. stalling vs. generation limit reached).

2.4.3 Number of Particles—The number of particles affects how effectively and efficiently the swarm can explore the search space. More particles increase the extent of exploration, but also increase computational demand. To select an appropriate number of particles for the modeled thalamic geometry, we assessed the performance of PSO with respect to 50, 100, 200, and 400 particles. Accuracy was measured using the minimum objective function value achieved, while computational demand was measured by the number of objective function calls by the program.

2.4.4 Stimulation Current Constraints—In our implementation of PSO for DBSA programming, each electrode configuration was represented as a particle in 32-dimensional space (corresponding to the 32 electrodes). Each dimension was bounded between 0 mA to -0.5 mA, corresponding to current constraints below typical current density limits (Table 1). We generated 100 particles, 68 of which were initialized to have random current through each electrode (within 0 to -0.25 mA). The remaining particles were initialized to have exactly -0.25 mA going through only a single electrode. This so-called “multi-start initialization” approach promoted exploration of the search space [45]. Constraints were enforced by checking for any out-of-bounds particles at each iteration. Particles that exceeded the current constraints at any given electrode were stripped of their velocity (i.e. $v=0$) in that electrode’s dimension, and that electrode’s max stimulus waveform current was set to either 0 mA or -0.5 mA, depending on which bound was exceeded. Notably, these current bounds were the only constraints applied to the particles.

2.4.5 Generations and Termination Criteria—Similar to genetic algorithm approaches, the PSO algorithm involves the exploration of a state space over the course of multiple “generations” in order to find optimal solutions. We established three termination criteria: (1) convergence, (2) stalling, and (3) reaching the generation limit. The criteria for convergence, stalling, and generation limits were tuned by trial and error. Specifically, we considered particles to be converged if the particles’ currents were close to each other in at least 30 out of 32 electrodes, as measured by a given electrode having a the standard deviation less than 0.004 mA across all particles. In addition, we considered particles to be stalled if the objective function score did not improve for 100 consecutive generations (Table 1). Meanwhile, we set the generation limit to be 200 generations (Table 1) since we found this enabled enough iterations for convergent runs to terminate while also enabling non-convergent runs to explore the space and improve Pareto front estimates. If any of the termination criteria were met, the PSO algorithm terminated and returned the estimated Pareto front (Fig. 4).

3. Results

3.1 PSO Variable Parameter Sweeps

Consistently low objective function values were obtained when c_1 was 3.6, 4.5, or 5.4 and when c_2 was 0.1, 0.2, or 0.4 (white outlined boxes in Fig 5A–B and black outlined box in Fig. 5C). We used low objective function values as a measure of accuracy because such values were obtained by having either (1) increased ROI activation, (2) decreased ROA activation, or (3) decreased power. As such, the lowest objective function value also corresponded to the most favorable electrode configuration possible. While the stochastic nature of our algorithm did not guarantee that it would find the lowest possible objective function value on any given run, we took a lower objective function value to mean that the algorithm had more closely approached this “true” best. Thus, from the set of c_1 and c_2 combinations that had low objective function values, we selected center values for c_1 as 4.5 and c_2 as 0.2 for all subsequent PSO runs (Table 1). With respect to number of particles, an increase from 50 to 400 caused computational demand to increase linearly, while generating

only a small improvement in accuracy (Fig 5D). We chose a balanced tradeoff by selecting 100 particles for the remaining simulations (Table 1).

3.2 Consistency of PSO Solutions

We measured the PSO algorithm's consistency by quantifying the variation in Pareto fronts across 30 independent runs (Fig. 6). In addition, we compared each estimate to a "combined" Pareto front that was constructed from all 30 Pareto front estimates. For each point Y on each of the 30 Pareto fronts, we computed the discrepancy between Y and the closest "combined" Pareto front point. As shown in Figure 6A for five sample Pareto fronts, the fronts exhibited considerable overlap, and their respective "best" electrode configurations had in common 8 out of 11–13 active electrodes (Fig. 6B). Furthermore, predicted ROI activation variation was within 2.0% across all 30 runs. Constructing a "combined" Pareto front (Fig. 6C) and comparing its distribution of values to the all 30 runs (uncombined), we observed a relatively small difference in median ROI activation (19 axons), median ROA activation (–20 axons), and median power (–4.11 mA²) (Fig. 6D).

3.3 Robustness

We evaluated robustness by measuring ROI, ROA, and Power under three conditions: (1) the existence of open electrode sites (i.e. those unusable because of very high impedances); (2) low per electrode current limit (e.g. low battery); and (3) lead shifts by 1 mm anterior, posterior, medial, and lateral relative to the DBSA trajectory. The PSO was able to accommodate for disabling 3 or 12 active electrodes, with ROI activation reduced by only 1.8% and 14%, respectively (Fig. 7A). Reducing the maximum per-electrode current by 50% and 80% reduced ROI activation by 5.6% and 16%, respectively (Fig. 7B). Lead displacement by 1 mm resulted in different yet relatively similar percent ROI activation (center: 27.8%, posterior: 20.3%, anterior: 24.2%, lateral: 20.9%, medial: 37.0%). ROA activation was not significantly different across shifts (<1% change in any direction) (Fig. 7C–D). Notably, additional simulations with larger lead shifts at 3 mm yielded more divergent values of percent ROI activation (center: 27.8%, posterior: 10.4%, anterior: 0%, lateral: 10.3%, medial: 41.4%). The active electrodes chosen by the algorithm were similar between the 1 mm and 3 mm shifts.

3.4 Efficiency

Efficiency was measured as the average runtime for constructing a Pareto front and obtaining a best electrode configuration. We assessed this by running the PSO algorithm five times on a PC with eight cores, 64-bit operating system, 24.0 GB RAM, and an Intel Core i7 processor at 3.40 GHz. The algorithm took 3.19 seconds per generation, resulting in an average of 10.6 minutes per run. These times reflected the duration for obtaining a solution after segmentation of the brain volume, positioning of the DBSA within the volume, incorporation of axonal tract morphologies, and identification of the MAFT for the ROI and ROA.

4. Discussion

In this study, we developed a particle swarm optimization approach to identify DBS electrode configurations and stimulation amplitudes that generate the most selective and efficient activation of a region of interest around a DBS array. This multi-objective problem incorporated a Pareto front, which offers a range of optimal electrode configurations from which a user can choose based on patient-specific programming goals. For example, a programmer may decide that activation within an ROA above 10% is undesirable for a given patient, whereas in another case, this value can be higher or lower. Thus, while the PSO is an automated algorithm, the user can select and test a set of optimized solutions along the Pareto front as needed, making the overall programming process more intuitive despite the complex geometries and nontrivial electrode configurations involved.

4.1 Particle Swarm Optimization

Particle swarm optimization algorithms employ cognitive and social components that adapt across multiple iterations, much like genetic algorithms and other evolutionary algorithms [22], [47]. In the case of the PSO algorithms, the individuals (or particles) survive throughout all iterations and continue to refine the solution to the overall problem [45]. This iterative refining process enables PSO to efficiently and effectively search the solution space. For multi-objective PSO, there are a number of approaches for selecting best solutions and subsequently updating particle positions. We used a mix between the linear aggregation-based approach and a Pareto dominance-based approach. Specifically, we used a linear aggregation objective function to guide exploration of the particles, and we used Pareto dominance criteria to construct an archive of Pareto-optimal points from all explored candidate solutions.

The linear aggregation-based approach combines all the objectives of a multi-objective problem into a single objective function with fixed weights. This is the simplest approach to multi-objective PSO [23]. However, this approach gives only one solution, and does not consider the range of the solutions that could be obtained by altering the weighting of the objectives. To address this, we coupled a Pareto dominance-based method to keep track of all the Pareto-optimal points at every intermediate iteration [48]. Notably, the Pareto front estimate was fundamentally independent of the final solution obtained from the aggregate-based approach, which yielded Pareto front estimates that were robust to choice of weighting for ROI, ROA, and power consumption. In other words, assigning a different weighting to each of the objectives (ROI, ROA, power) would yield a comparable Pareto front estimate in spite of the algorithm evolving differently across iterations.

In addition to linear aggregation-based approaches, there are also a number of more complex Pareto dominance-based methods [49] that may yield more efficient Pareto front estimates (e.g. in fewer generations or with fewer particles). These Pareto dominance-based approaches rely on selecting non-dominated “leader” solutions to guide the algorithm. Employing such methods for DBSA programming would require testing and tuning several different parameters, such as leader selection schemes, swarm diversity, particle front spread, and archive maintenance [23].

Also important to PSO algorithms in general is the topological network of connections amongst particles and neighborhoods, which govern the convergence behavior of the algorithm. Here, we fully connected all particles to one another such that all particles in the swarm were directly informed of the global best solution at each iteration. Using such a global network enables the swarm to converge more rapidly than using a local network (e.g. ring, tree, wheel, von Neumann networks), since a locally connected network is slower to propagate the information of the global best solution. The global network's faster convergence, however, may be undesirable if it leads to premature convergence on local optima [48]. As a result, using a local neighborhood approach may improve the diversity of exploration and consistency of the Pareto front estimates [50], [51].

To avoid premature convergence and facilitate exploration, we employed three strategies: (1) we implemented a multi-start initialization approach (see Section 2.4.4); (2) we chose low social velocity relative to cognitive velocity, which promotes exploration [48]; and (3) we set the convergence criteria to be very strict, as the particles are considered converged only when the standard deviation of current was less than 0.004 mA (<1% of max current) in at least 30 electrodes.

Notably, while exploration is essential to obtaining consistent Pareto front estimates, convergence is not necessary. We found that similar and comparable results were obtained from non-converging runs that reached the generation limit. This finding enabled us to specify a generation limit (see 2.4.5) such that non-converging runs continued to make small improvements before terminating. Such a strategy produced more exploration at the cost of additional generations of the PSO. Future studies may wish to try different topological network organizations to aid the PSO algorithm in optimizing its parameter space exploration and convergence.

4.2 PSO Performance

4.2.1 Consistency—Across 30 independent runs, the means and standard deviations of the best points on the Pareto fronts were 1281 ± 22 (1.7%) for ROI activation, 34 ± 5 (14.7%) for ROA activation, and $34.0 \pm 4.6 \text{ mA}^2$ (13.5%) for Power. Furthermore, visual inspection of five sample Pareto fronts indicated a significant overlap between the fronts. This level of consistency meant that one could expect the PSO algorithm to provide fairly similar outputs across runs (particularly for ROI activation) in spite of varying initial conditions.

The discrepancies between the “combined” Pareto front and the 30 individual fronts were 19 axons for median ROI activation, -20 axons for median ROA activation, and -4.11 mA^2 for median power. The signs reflect the fact that the “combined” Pareto front had higher ROI activation, lower ROA activation, and lower power, which was consistent with expectations. Considering these discrepancies in terms of the total number of axons for each region (4,549 for ROI and 5,937 for ROA) and the maximum power attained (60.04 mA^2), these discrepancies corresponded to <1% for ROI and ROA axons, and 6.8% for power. The overall small discrepancies further suggested that the PSO algorithm could come up with similar estimates of the Pareto front across runs. However, despite this overall low discrepancy, there were still points on the individual Pareto fronts that reached high

discrepancies (e.g. up to ~300 axons, or ~7% of maximum ROI activated). Future implementations involving purely dominance-based Pareto front estimation (section 4.1) may help reduce the incidence of such deviations.

4.2.2 Robustness—The PSO was also robust to translation of the lead trajectory in relation to the ROA, which was located posterior, lateral, and slightly ventral to the ROI. As a result, shifting the lead 1 mm in the posterior direction positioned the lateral and posterior portions of the lead closer to the ROA. This caused a reduction in stimulation through the posterior electrodes (21, 25, 29) and the lateral electrodes (24, 28). The ventral electrodes (21, 22, 24) also had less current because the shape of the ROA at the new lead location had a greater ventral extent. The consequence of fewer electrodes being active and the lead being closer to the ROA was lower Power consumption, lower ROI activation, and relatively unchanged ROA activation. This result aligned with our objective function weightings in which low ROA activation was more highly prioritized (Equation 9). Shifting the lead 1 mm in the lateral direction had a similar impact as the posterior shift, since it positioned the lead closer to the ROA. As a result, the posterior electrodes (site 21) and lateral electrodes (sites 24, 28) experienced a reduced current. In contrast to the posterior shift, however, the ventral electrodes were still active since the ventral end of the lead in the lateral shift still remained surrounded by ROI axons. The resulting solution had lower ROI activation, lower ROA activation, and lower Power (though Power was not as low as in the posterior shift). Shifting the lead 1 mm in the anterior direction had the opposite effect of the posterior shift. Since the lead at that location was further away from the ROA, there was less ROA activation. In addition, virtually every electrode in the ROI was sourcing current, resulting in an overall larger Power consumption. However, since this shift also caused the lead to have fewer surrounding ROI fibers, there was a reduction in ROI activation relative to the original location. Shifting the lead 1 mm in the medial direction positioned the lead farther away from the ROA in all directions, as well as positioning it deeper in the ROI. This resulted in more active ventral electrodes (sites 17–20), which increased Power (though not as much as in the anterior shift) and increased ROI activation. Notably, ROA activation was also increased, reflecting a tradeoff between the large increase in ROI activation and the relatively smaller increase in ROA activation (Equation 9).

Additionally, the PSO showed robustness to disabling ‘bad’ electrodes on the lead as well as modifying maximum current amplitudes resulting from a low battery, for instance. In the PSO runs, disabling posterior electrode 25 and anterior electrodes 27 and 31 led to an increase in current through the neighboring anterior electrode 23. Disabling the three most dorsal rows of electrodes (21–32) resulted in relatively higher current through the ventral electrodes 14, 17, 18, and 19. Overall, we observed the expected trend of reduced ROI, ROA, and Power as more electrodes were disabled. When employing more stringent current constraints, we also observed that the electrode configurations did not appreciably change in topology as current limits were reduced. As expected, lower current limits were associated with lower activation of ROI and ROA, as well as lower Power consumption. This was consistent with the idea that a higher power consumption would allow for more activation, all else being equal. Notably, if we had instead set constraints that force the algorithm to

budget a total amount (e.g. sum) of current across all electrodes, the reduced current limits may have evoked more complex changes in electrode configuration pattern.

4.3 Models for Predicting Axonal Activation

The PSO approach relies on accurate prediction of neural modulation from electrical stimulation. We used a modified version of the activating function (MAFT-predicted activation) to demonstrate the PSO approach. This MAFT-predicted activation tracked the multi-compartment NEURON models well (Fig. 2). While one could solve the nonlinear differential equations governing a multi-compartment neuron model as performed in NEURON and as we have shown here for a thalamocortical axon model, the computational time required for these calculations is high. This limitation led several groups to develop other simplified and more computationally efficient axonal excitability prediction models. These models have ranged from a discrete difference approximation of transmembrane current induced by extracellular stimulation (i.e. the so-called ‘activating function’) [19], [40], to variations in the activating function that account for redistribution of current via space constants [52], to inclusion of additional source terms at other points along an axon that can affect a given node through intracellular ohmic conductance [20], to a weighted distribution of source terms [21]. It is important to note that the PSO algorithm as described here can utilize other computational approaches for modeling axonal excitability resulting from DBS, as long as the predictive functions are specified (Equations 3 and 4).

It is also important to acknowledge that predicting axonal activation depends on a variety of factors including axon diameter, distance between nodes of Ranvier, ion channel density at the nodes and surrounding compartments, current waveforms, and geometry of the axons relative to the induced electric field [16]. In terms of latter, the MAFT value was found to vary slightly amongst axonal orientations, which means that one would need to identify an activating function threshold a priori and for each axonal pathway that one wishes to evaluate with the PSO. Along these lines, it would be possible to introduce more complex axonal geometries, such as nonlinear axonal trajectories, while keeping intact the fundamental framework of PSO for programming DBSA’s (equations 9 & 10). However, the activation prediction functions, $R(x)$ and $S(x)$, would need to be tuned to obtain reasonable predictions for the new, non-linear geometries. Future developments in activation predictions could be readily deployed into the PSO programming algorithm to facilitate robust prediction across stimulation parameters.

Because of this limitation on axonal geometry, biophysical models for the axonal geometry of interest must be constructed and run once as a primer for identifying an appropriate MAFT value. Since accurate prediction of axonal activation relied on the appropriate selection of a MAFT value, a poor choice of MAFT value would lead to suboptimal predictions. While this need for a small-scope biophysical model simulation adds an extra setup step in applying the PSO approach, the overall process remains computationally efficient. Another important limitation in the application of the PSO in this study is that we modeled only a single axonal pathway between thalamus and cortex. There are indeed several pathways that are likely activated during thalamic DBS, including the reticular nucleus to VPLo, the cerebellothalamic tract, and corticothalamic tract. Their activation as

well as the dynamic synaptic modulation resulting from the activation may underlie components of the therapy.

4.4 Practical Considerations

The results on robustness highlight an important practical point for DBS lead placement and the capabilities of current steering algorithms. While the simulations in Figure 7C–D suggest that current steering is able to compensate for off-target placement, this ability of directional DBS leads may be limited to displacements of a couple of millimeters or less [36]. At larger displacements, current steering through a single lead of electrodes may not adequately compensate for off-target lead placement, depending on the size of target. For the VPL example demonstrated in this paper, applications of the PSO to lead displacements of 3 mm yielded greatly reduced to no activation for shifts away from ROI.

For proof-of-concept and consistency with other ongoing work in our group, the PSO approach presented here was implemented for a 32-electrode DBS array with four radial electrodes per row. However, PSO programming is readily applicable to the eight-electrode directional DBS leads that have recently emerged on the market, such as the Vercise and Infinity DBS systems. While any given row of these commercial DBS leads only has three radial electrodes, this reduced number of radial electrodes does not significantly diminish steering capabilities when current is allowed to be non-uniformly distributed across the electrodes [36]. As such, the PSO approach can be readily integrated into a system that uses such DBS leads, enabling the development of a system for efficient, subject-specific DBS programming.

The PSO approach was observed to be efficient by running in a matter of minutes. However, it is important to acknowledge that running the PSO algorithm requires a series of preparatory steps that include segmenting the target structures from subject-specific anatomical images, as well as specifying the trajectory of the axonal pathways. Ultimately, the PSO approach for DBS programming is intended for use in subject-specific contexts. As such, widespread use of the algorithm will require that DBS systems become streamlined for anatomically-informed, subject-specific DBS. We anticipate that this streamlining process will occur in tandem with the current developments in directional DBS, due to the complementary nature of these two technologies. For the time being, centers that already have the technical expertise to implement prospective subject-specific models can leverage the PSO approach to enable optimized programming of directional DBS leads.

4.5 Future Applications

The PSO algorithm demonstrated here was applied to a pathway thought to be involved in DBS therapy for treating Essential Tremor [6], [53]. However, it could readily be applied to other known DBS targets. A number of DBS targets have been characterized in terms of their relevant fiber tracts and morphologies using probabilistic tractography and/or histological processing in the context of computational neuron models of DBS. For instance, in Parkinson's disease, STN-DBS therapy is thought to stem from a combination of modulating pathways in and adjacent to the STN [10], [11], [54], and a similar combination of therapeutic pathways is thought to occur with GPi-DBS [55], [56]. There is also ongoing

research to characterize targets for other indications, such as in Major Depressive Disorder, where retrospective modeling suggests that the forceps minor, uncinate fasciculus, and cingulum bundle may play a critical role in the therapeutic DBS mechanisms [13]. It is important to note that the PSO approach presented here can be readily extended to cases in which there are multiple ROI's. In this case, the PSO algorithm could provide the user with a set of electrode configurations that intuitively demonstrate the tradeoff between stimulating one target vs. another target (i.e. via the Pareto Front).

In addition, one could also consider including other stimulation parameters such as bipolar/multipolar stimulation, a range of pulse widths, and potentially different stimulation pulse train patterns. The key to incorporating additional parameters would be to have a mathematical representation of how these factors influence axonal excitability. While this study used a simplistic homogeneous and isotropic model of tissue conductance to obtain extracellular potential values at the axonal nodes of Ranvier, one may consider using more complex models of brain tissue with the PSO algorithm. For example, studies have shown that finite element models that incorporate inhomogeneous and anisotropic tissue properties (assigning different conductance values to various tissue types) improve modeling results and may better reflect the physiological properties of the brain [11], [32], [57]. Finally, while the implementation presented here was designed for programming a set of electrodes that have already been implanted, it is conceivable to leverage the efficiency of this approach for pre-surgical planning of DBS lead placement. By simulating different lead placements (e.g. location, angle of attack, etc.) and aiming for the one that best achieves the objectives, one may be able to optimize lead placement and therefore improve targeting and functional outcomes of DBS.

5. Conclusion

We have shown particle swarm optimization to be an efficient, consistent, and robust method for programming DBS arrays. The proposed implementation provides the end-user flexibility to select among alternative configurations and objective function specification along the Pareto front to accommodate subject-specific needs.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by the NIH (R01-NS081118, R01-NS094206, P50-098573) and the Michael J. Fox Foundation. We thank Drs. Noam Harel, Essa Yacoub, and Gregor Adriany at the Center for Magnetic Resonance Research (P41-EB015894, P30-076408, U54-MH091657) for help with the thalamic reconstructions used in this study. We also thank members of the Neuromodulation Research and Technology Lab, especially Ben Teplitzky, for helpful feedback on the simulations and manuscript.

References

1. Wichmann T, Delong M. Deep Brain Stimulation for Neurologic and Neuropsychiatric Disorders. *Neuron*. Oct; 2006 52(1):197–204. [PubMed: 17015236]

2. Contarino MF, et al. Directional steering: A novel approach to deep brain stimulation. *Neurology*. Sep; 2014 83(13):1163–9. [PubMed: 25150285]
3. Martens HC, et al. Spatial steering of deep brain stimulation volumes using a novel lead design. *Clin Neurophysiol*. Mar; 2011 122(3):558–66. [PubMed: 20729143]
4. Pollo C, et al. Directional deep brain stimulation: an intraoperative double-blind pilot study. *Brain*. Jul; 2014 137(Pt 7):2015–26. [PubMed: 24844728]
5. Connolly AT, Jensen AL, Baker KB, Vitek JL, Johnson MD. Classification of pallidal oscillations with increasing parkinsonian severity. *J Neurophysiol*. Jul; 2015 114(1):209–218. [PubMed: 25878156]
6. Keane M, Deyo S, Abosch A, Bajwa JA, Johnson MD. Improved spatial targeting with directionally segmented deep brain stimulation leads for treating essential tremor. *J Neural Eng*. Aug.2012 9(4):46005.
7. van Dijk KJ, et al. A novel lead design enables selective deep brain stimulation of neural populations in the subthalamic region. *J Neural Eng*. 2015; 12(4):46003.
8. Zitella LM, Mohsenian K, Pahwa M, Gloeckner C, Johnson MD. Computational modeling of pedunculopontine nucleus deep brain stimulation. *J Neural Eng*. May.2013 10(4):45005.
9. Hunka K, Suchowersky O, Wood S, Derwent L, Kiss ZHT. Nursing time to program and assess deep brain stimulators in movement disorder patients. *J Neurosci Nurs J Am Assoc Neurosci Nurses*. Aug; 2005 37(4):204–210.
10. Butson CR, Cooper SE, Henderson JM, Wolgamuth B, McIntyre CC. Probabilistic analysis of activation volumes generated during deep brain stimulation. *Neuroimage*. Feb; 2011 54(3):2096–2104. [PubMed: 20974269]
11. Chaturvedi A, Butson CR, Lempka SF, Cooper SE, McIntyre CC. Patient-specific models of deep brain stimulation: Influence of field model complexity on neural activation predictions. *Brain Stimulat*. Apr; 2010 3(2):65–77.
12. Lujan JL, Chaturvedi A, Malone DA, Rezai AR, Machado AG, McIntyre CC. Axonal pathways linked to therapeutic and nontherapeutic outcomes during psychiatric deep brain stimulation. *Hum Brain Mapp*. Apr; 2012 33(4):958–968. [PubMed: 21520343]
13. Riva-Posse P, et al. Defining Critical White Matter Pathways Mediating Successful Subcallosal Cingulate Deep Brain Stimulation for Treatment-Resistant Depression. *Biol Psychiatry*. Dec; 2014 76(12):963–969. [PubMed: 24832866]
14. Chaturvedi A, Luján JL, McIntyre CC. Artificial neural network based characterization of the volume of tissue activated during deep brain stimulation. *J Neural Eng*. Oct.2013 10(5):56023.
15. Xiao Y, Pena E, Johnson MD. Theoretical Optimization of Stimulation Strategies for a Directionally Segmented Deep Brain Stimulation Electrode Array. *IEEE Trans Biomed Eng*. Feb; 2016 63(2):359–371. [PubMed: 26208259]
16. McIntyre CC, Grill WM. Extracellular stimulation of central neurons: influence of stimulus waveform and frequency on neuronal output. *J Neurophysiol*. Oct; 2002 88(4):1592–1604. [PubMed: 12364490]
17. Wei XF, Grill WM. Current density distributions, field distributions and impedance analysis of segmented deep brain stimulation electrodes. *J Neural Eng*. Dec; 2005 2(4):139–147. [PubMed: 16317238]
18. Butson CR, McIntyre CC. Current steering to control the volume of tissue activated during deep brain stimulation. *Brain Stimul*. Jan; 2008 1(1):7–15. [PubMed: 19142235]
19. Rattay F. Analysis of models for external stimulation of axons. *IEEE Trans Biomed Eng*. Oct; 1986 33(10):974–977. [PubMed: 3770787]
20. Warman EN, Grill WM, Durand D. Modeling the effects of electric fields on nerve fibers: determination of excitation thresholds. *IEEE Trans Biomed Eng*. Dec; 1992 39(12):1244–1254. [PubMed: 1487287]
21. Peterson E, Izad O, Tyler D. Predicting myelinated axon activation using spatial characteristics of the extracellular field. *J Neural Eng*. Aug.2011 8(4):46030.
22. Brill N, Tyler D. Optimizing nerve cuff stimulation of targeted regions through use of genetic algorithms. *Conf Proc Annu Int Conf IEEE Eng Med Biol Soc IEEE Eng Med Biol Soc Annu Conf*. 2011; 2011:5811–5814.

23. Parsopoulos KE, Vrahatis MN. Recent approaches to global optimization problems through Particle Swarm Optimization. *Nat Comput.* Jun; 2002 1(2–3):235–306.
24. Huang KW, Yang CS, Tsai CW. A Two-Phase Hybrid Particle Swarm Optimization Algorithm for Solving Permutation Flow-Shop Scheduling Problem. *Int J Comput Appl.* Jun; 2012 48(1):11–18.
25. Slade W. Computational Intelligence Approaches to Ocean Color Inversion. *Electron Theses Diss.* Dec.2004
26. Santos, SM., Valença, MJS., Bastos-Filho, CJA. Comparing Particle Swarm Optimization Approaches for Training Multi-Layer Perceptron Neural Networks for Forecasting. In: Yin, H.Costa, JAF., Barreto, G., editors. *Intelligent Data Engineering and Automated Learning - IDEAL 2012.* Springer; Berlin Heidelberg: 2012. p. 344–351.
27. Wu D, et al. Prediction of Parkinson's disease tremor onset using a radial basis function neural network based on particle swarm optimization. *Int J Neural Syst.* Apr; 2010 20(2):109–116. [PubMed: 20411594]
28. Saini S, Zakaria N, Rambli DRA, Sulaiman S. Markerless Human Motion Tracking Using Hierarchical Multi-Swarm Cooperative Particle Swarm Optimization. *PLoS ONE.* May.2015 10(5):e0127833. [PubMed: 25978493]
29. Connolly AT, et al. A Novel Lead Design for Modulation and Sensing of Deep Brain Structures. *IEEE Trans Biomed Eng.* Jan; 2016 63(1):148–157. [PubMed: 26529747]
30. Yousif N, Liu X. Investigating the depth electrode–brain interface in deep brain stimulation using finite element models with graded complexity in structure and solution. *J Neurosci Methods.* Oct; 2009 184(1):142–151. [PubMed: 19596028]
31. Miocinovic S, et al. Experimental and theoretical characterization of the voltage distribution generated by deep brain stimulation. *Exp Neurol.* Mar; 2009 216(1):166–176. [PubMed: 19118551]
32. Schmidt C, van Rienen U. Modeling the field distribution in deep brain stimulation: the influence of anisotropy of brain tissue. *IEEE Trans Biomed Eng.* Jun; 2012 59(6):1583–1592. [PubMed: 22410323]
33. Howell B, McIntyre CC. Analyzing the tradeoff between electrical complexity and accuracy in patient-specific computational models of deep brain stimulation. *J Neural Eng.* Jun; 2016 13(3):36023–36039.
34. Grill WM, Mortimer JT. Electrical properties of implant encapsulation tissue. *Ann Biomed Eng.* 22(1):23–33.
35. Ranck JB. Specific impedance of rabbit cerebral cortex. *Exp Neurol.* Feb; 1963 7(2):144–152. [PubMed: 13990734]
36. Teplitzky BA, Zitella LM, Xiao Y, Johnson MD. Model-Based Comparison of Deep Brain Stimulation Array Functionality with Varying Number of Radial Electrodes and Machine Learning Feature Sets. *Front Comput Neurosci.* 2016; 10:58. [PubMed: 27375470]
37. Kuncel AM, et al. Clinical response to varying the stimulus parameters in deep brain stimulation for essential tremor. *Mov Disord.* Nov; 2006 21(11):1920–1928. [PubMed: 16972236]
38. Xiao Y, et al. Multimodal 7T Imaging of Thalamic Nuclei for Preclinical Deep Brain Stimulation Applications. *Brain Imaging Methods.* 2016:264.
39. Crouch RL. The nuclear configuration of the hypothalamus and the subthalamus of *Macacus rhesus*. *J Comp Neurol.* Jun; 1934 59(3):431–449.
40. Rattay F. Analysis of models for extracellular fiber stimulation. *IEEE Trans Biomed Eng.* Jul; 1989 36(7):676–682. [PubMed: 2744791]
41. Plonsey, R. *Bioelectric Phenomena.* 1. New York u.a: McGraw-Hill Inc., US; 1969.
42. SASMOR, L., SMITH, ED. *Handbook of Biomedical Engineering.* Academic Press; 1988. Chapter 17 - Artificial Intelligence in Medicine A2 - KLINE, JACOB; p. 417–437.
43. Teplitzky BA, Connolly AT, Bajwa JA, Johnson MD. Computational modeling of an endovascular approach to deep brain stimulation. *J Neural Eng.* Apr.2014 11(2):26011.
44. Lempka SF, Johnson MD, Miocinovic S, Vitek JL, McIntyre CC. Current-controlled deep brain stimulation reduces in vivo voltage fluctuations observed during voltage-controlled stimulation. *Clin Neurophysiol Off J Int Fed Clin Neurophysiol.* Dec; 2010 121(12):2128–2133.

45. Kennedy J, Eberhart R. Particle swarm optimization. 1995; 4:1942–1948.
46. Shi, Y., Eberhart, R. A modified particle swarm optimizer. The 1998 IEEE International Conference on Evolutionary Computation Proceedings, 1998. IEEE World Congress on Computational Intelligence; 1998; p. 69-73.
47. Holland, JH. Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence. Cambridge, Mass: MIT Press; 1992. 1st MIT Press ed
48. Engelbrecht, AP. Fundamentals of computational swarm intelligence. Hoboken, NJ: Wiley; 2005.
49. Reyes-sierra M, Coello CAC. Multi-Objective particle swarm optimizers: A survey of the state-of-the-art. Int J Comput IN TellIGENCE Res. 2006; 2(3):287–308.
50. Kennedy, J. Small worlds and mega-minds: effects of neighborhood topology on particle swarm performance. Proceedings of the 1999 Congress on Evolutionary Computation, 1999. CEC 99; 1999; 1938.
51. Kennedy, J., Mendes, R. Population structure and particle swarm performance. Proceedings of the 2002 Congress on Evolutionary Computation, 2002. CEC '02; 2002; p. 1671-1676.
52. Altman KW, Plonsey R. Analysis of excitable cell activation: relative effects of external electrical stimuli. Med Biol Eng Comput. Nov; 1990 28(6):574–580. [PubMed: 2287182]
53. Groppa S, Herzog J, Falk D, Riedel C, Deuschl G, Volkmann J. Physiological and anatomical decomposition of subthalamic neurostimulation effects in essential tremor. Brain J Neurol. Jan; 2014 137(Pt 1):109–121.
54. Miocinovic S, et al. Computational analysis of subthalamic nucleus and lenticular fasciculus activation during therapeutic deep brain stimulation. J Neurophysiol. Sep; 2006 96(3):1569–80. [PubMed: 16738214]
55. Johnson MD, McIntyre CC. Quantifying the neural elements activated and inhibited by globus pallidus deep brain stimulation. J Neurophysiol. Nov; 2008 100(5):2549–2563. [PubMed: 18768645]
56. Johnson MD, Zhang J, Ghosh D, McIntyre CC, Vitek JL. Neural targets for relieving parkinsonian rigidity and bradykinesia with pallidal deep brain stimulation. J Neurophysiol. Jul; 2012 108(2): 567–577. [PubMed: 22514292]
57. Zitella LM, et al. Subject-specific computational modeling of DBS in the PPTg area. Front Comput Neurosci. 2015; 9:93. [PubMed: 26236229]

Figure 1.

Models of thalamic DBSA stimulation. (A) Reconstructions of the cerebellar-receiving area of motor thalamus (VPLo, green) and somatosensory thalamus (VPLc, blue) from the rhesus macaque, showing DBSA placed 77° above the horizontal plane and at 10° from the sagittal plane in an anterior to posterior trajectory [15]. (B) Trajectories of thalamocortical axons (VPLo axons: red, VPLc axons: blue) extending from both thalamic nuclei.

Figure 2.

Comparison of MAFT predictions and NEURON model predictions in the ROI and ROA. (A) Spatial cross sectional view of a subset of axons illustrated overall agreement between the MAFT (o) and NEURON predictions (x). (B) Comparison of MAFT and NEURON predictions in terms of number of axons activated. Results indicated a discrepancy of -1.3 ± 38 axons (slight MAFT underprediction) for ROI, and 1.6 ± 20 axons (slight MAFT overprediction) for ROA. This constituted a discrepancy of less than 1% of all axons.

Figure 3. Typical output from a PSO run. (A) A three-dimensional Pareto front showing five unique points on the Pareto front, where each point corresponds to a different predicted ROI and ROA activation (i.e. number of axons activated) and power consumption (in mA^2 without the impedance). Among them, point number 1 achieved the lowest objective function value (Equation 9). (B) Electrode configurations for each of the five highlighted points. Depending on a user's desired prioritization between ROI, ROA, and Power, a different point along the Pareto front can be selected for stimulation. The table shows the PSO-predicted axonal activation (bold, in number of axons) in relation to the total number of axons in ROI and ROA, as well as the PSO-predicted power consumption (in mA^2 without the impedance) for each electrode configuration.

Figure 4.
Flowchart for the PSO algorithm

Figure 5.

PSO algorithm parameter sweeps across cognitive and social velocities. Plots show (A) mean objective function values across 5 runs, (B) standard deviation of objective function values across 5 runs, and (C) termination behavior. The boxes highlight a range of c_1 and c_2 values that consistently yielded low objective function values with little stalling. (D) Effects of number of particles on objective function value and computational demand.

Figure 6. Pareto fronts from multiple PSO algorithm runs. The point that minimized the objective function is highlighted on each of the five independent Pareto front (A) and the corresponding electrode configurations are provided (B). The table shows the PSO-predicted number of axons activated in relation to the total number of axons in ROI and ROA, as well as the PSO-predicted power consumption for each electrode configuration. Pareto fronts and electrode configurations exhibited consistent topographies across runs. The “combined” Pareto front (C) was obtained from all points from 30 independent Pareto fronts. (D) Comparing the distribution of ROI, ROA, and Power across all runs to the distribution across the combined Pareto front in (C), there was a relatively small difference in axons activated and power. ROI and ROA are in units of “number of axons” and power is in unit of “mA²” without the impedance.

Figure 7.

Three tests for PSO algorithm robustness. (A) When three select electrodes near the ROI were disabled, the algorithm maintained relatively stable activation of ROI. Disabling additional electrodes resulted in ineffective targeting. (B) The algorithm exhibited similar robustness when reducing the upper bound current. (C–D) When shifting the DBSA lead from original center position “C” to 1 mm posterior “P”, anterior “A”, lateral “L”, and medial “M” relative to the DBSA trajectory, the algorithm adjusted stimulation to reflect the new location of the spatial target. The tables show the PSO-predicted number of axons activated in relation to the total number of axons in ROI and ROA, as well as the PSO-predicted power consumption for each electrode configuration. ROI and ROA are in units of “number of axons” and power is in unit of “mA²” without the impedance.

Table 1

PSO Algorithm Parameter Values

Parameter	Value
Number of Particles	100
Generation Limit	200
Stalling Generation Constant	100
Cognitive Attraction Weight, c_1	4.5
Social Attraction Weight, c_2	0.2
Upper Inertia	0.9
Lower Inertia	0.4
Velocity Limit	Infinite
Additional Cognitive Weighting, r_{1ij}	[0,1] (random)
Additional Social Weighting, r_{2ij}	1
Lower Bound Current per Contact	-0.5 mA
Upper Bound Current per Contact	0 mA
Total Bounded Current	none