

# Toward Closed-Loop Optimization of Deep Brain Stimulation for Parkinson's Disease: Concepts and Lessons from a Computational Model

Xiao-jiang Feng<sup>§</sup>, Brian Greenwald, Herschel Rabitz\*  
Eric Shea-Brown<sup>§†</sup>, and Robert Kosut<sup>‡</sup>

Appears in *J. Neuroengineering* 4 (2007) L14-L21

Department of Chemistry, Princeton University, Princeton, NJ 08544

<sup>†</sup> Courant Institute of Mathematical Sciences and Center for Neural Science, New York University, NY 10012

<sup>‡</sup> SC Solutions, Sunnyvale, CA 94085

\* To whom correspondence should be addressed: Department of Chemistry, Princeton University, Princeton, NJ 08544. Phone: 609-258-3917, email: hrabitz@princeton.edu

<sup>§</sup> These authors contributed equally to this work.

# Abstract

Deep brain stimulation (DBS) of the subthalamic nucleus with periodic, high frequency pulse trains is an increasingly standard therapy for advanced Parkinson’s disease. Here, we propose that a closed-loop global optimization algorithm may identify novel DBS waveforms that could be more effective than their high-frequency counterparts. We use results from a computational model of the Parkinsonian basal ganglia to illustrate general issues relevant to eventual clinical or experimental tests of such an algorithm. Specifically, while the relationship between DBS characteristics and performance is highly complex, global search methods appear able to identify novel and effective waveforms with convergence rates that are acceptably fast to merit further investigation in laboratory or clinical settings.

## 1 Introduction

Deep brain stimulation (DBS) is an increasingly standard therapy for varied movement disorders [1]. In advanced Parkinson’s disease (PD), for instance, a high-frequency ( $> 100Hz$ ) stimulus is typically applied to the subthalamic nucleus (STN) or globus pallidus pars interna (GPi) of the basal ganglia [2, 3]. The therapeutic effects are often remarkable, alleviating many Parkinsonian motor symptoms and decreasing dependence on dopaminergic drugs [4–7].

Nevertheless, neither the mechanism of action for DBS nor the neurophysiology of the PD state are fully characterized [3, 8–21]. As a result, it is not known *a priori* which DBS waveform will be most effective for a given patient, and the standard DBS parameters (e.g., the frequency, duration, and amplitude of a square-wave pulse train) typically must be hand-tuned by neurologists [22]. Because both motor symptoms and side effects often follow predictable trends as the DBS parameters are varied within certain ranges [23], general clinical strategies for identifying appropriate parameters may be based on gradually increasing or decreasing individual parameters. In practice, this typically leads to effective parameter settings for high-frequency DBS after sampling an acceptable number of values [22].

However, it is possible that DBS inputs that differ in form from such high-frequency, periodic patterns could be more desirable for achieving clinical objectives. Specific alternatives to high-frequency

DBS patterns have been suggested via computational studies involving phase resetting and delayed feedback via DBS [24–27]. Part of the motivation is that if novel DBS waveforms were found that required lower average amplitudes than their high frequency counterparts, their use would prolong the life of stimulator batteries (which require surgical replacement). Such lower-amplitude patterns could also reduce leakage of DBS current into neighboring brain areas, potentially attenuating side effects of DBS. Perhaps more importantly, these alternative DBS waveforms could produce firing patterns in the basal ganglia which more closely resemble the nonparkinsonian state than those elicited by high-frequency DBS, which is believed to either suppress or entrain cell bodies and axons at close to the stimulator frequency. Finally, it is possible that alternative DBS waveforms would be useful for patients unresponsive to standard high-frequency DBS.

There is no guarantee, however, that novel DBS waveforms could be found by similar tuning strategies as described above for high-frequency DBS. As explored below, this is because the nuanced mechanisms by which such waveforms might act suggests that there will be a complex relationship between DBS parameters and their therapeutic effects, so that trends would be difficult to identify. Furthermore, grid search methods would be impossible if novel DBS waveforms were described by a large number of parameters (e.g., aperiodic inputs): the number of parameter combinations to be searched grows exponentially with the number of parameters. The alternative is to employ an appropriate algorithm to search for the optimal DBS settings giving the desired therapeutic effects. This algorithm needs to have four basic properties to be clinically applicable. First, its search/optimization needs to be model-free. In other words, its operation must directly utilize clinical observables that can be measured rapidly. Second, the algorithm must converge in an acceptable number of trial DBS applications. Third, the algorithm needs to have global search capability in order to fully explore the DBS parameter space for satisfactory therapeutic effects. Lastly, the algorithm should be able to operate with different types of DBS inputs (e.g., both periodic and aperiodic waveforms) to provide maximum flexibility to individual patients and neurologists.

In this article, we propose such an approach that integrates a nonlinear closed-loop learning algorithm with existing instrumentation with the aim of *optimally* achieving DBS therapy of Parkinsonian motor symptoms. The algorithm we suggest not only possesses all the properties above, but also provides several additional advantages, including favorable scaling with respect to the number of DBS

parameters and automatic balancing of practical constraints. Below, we describe the components and the operational procedure of this method in the eventual clinical applications that we envision, followed by an illustration of the algorithm based on a computational basal ganglia network model [20, 28]. Although this computational model is highly simplified in comparison with the real basal ganglia, using the learning algorithm to identify its optimal inputs illustrates several central issues relevant for future attempts at applying the algorithm to more complex models, laboratory experiments, or, eventually, in the clinical setting.

These central points, which are enumerated in boldface as they arise, are: (1) The dependence of basal ganglia network dynamics on DBS parameters can be highly complex, requiring nonlocal optimization algorithms; (2) Novel DBS waveforms, including those of low frequency, appear to exist which can return Parkinsonian network dynamics close to the normal state; (3) Stochastic DBS waveforms may be effective alternatives to periodic inputs; (4) The genetic algorithm (GA) automatically identifies *multiple* DBS waveforms of comparable efficacy; (5) Stochastic DBS waveforms are robustly effective as the frequency of burstlike Parkinsonian activity is varied in the computational model; (6) The GA requires evaluating only a few hundreds of parameter settings to converge in the computational model, and this number is fairly invariant to the number of DBS parameters.

Below, we demonstrate these features and discuss their consequences. First, however, we discuss several general concepts relevant to closed-loop optimization of DBS waveforms.

## 2 Closed-Loop Optimization of DBS Waveforms: General Concepts

**Algorithms:** In general, model-free optimization can be achieved by adaptive (closed-loop) methods. Since the landscape describing the dependence of therapeutic effects on DBS parameters can be “rugged” (i.e., containing multiple local extrema), search algorithms which do not depend on local estimates of slope are required. Genetic algorithms (GA) [29] are a class of such global methods and will be introduced below and applied to the basal ganglia model. After the optimal DBS settings have been identified, either these global algorithms or local methods such as linear feedback control [30] can

be utilized to continually adapt DBS settings to gradual physiological changes during chronic therapy.

**Assessing DBS effects:** During the optimization procedure, how would the efficacy of individual DBS waveforms be assessed? In a clinical setting, quantitative measures of PD symptoms such as tremor, upper limb rigidity, and bradykinesia could be used [23]. Such measurements could either be made by a neurologist or, in some cases, by an autonomous device, with either manual or automated interfacing with the search algorithm running on a personal computer. Depending on the implanted instrumentation, it is possible that statistical measures of the difference between firing patterns in the normal and Parkinsonian states could also be derived. In laboratory experiments, cellular recordings would likely be easier to obtain, providing ready assessment of DBS effects. In the computational model explored below, we use a statistical measure of simulated spike times in the basal ganglia nucleus Gpi.

**Parameterizing deterministic and stochastic DBS waveforms:** The standard square-wave DBS inputs are defined by pulse amplitude, period, and duration (as well as electrode index and polarity [22] in some cases). Alternative DBS waveforms can be parameterized in many ways: for example, via Fourier or Haar basis coefficients. In such cases, the resulting waveform would be periodic with some (possibly long) period.

Aperiodic, stochastic DBS waveforms are also possible. For instance, a stochastic DBS pulse train may be generated from a probability density function (PDF)  $P(d)$  (satisfying  $\int P(d) dd = 1$ ) which defines the probability for two adjacent DBS pulses spaced by a time delay  $d$ . To enable an efficient search, a piecewise constant PDF  $P(i), i = 1, 2, \dots, I$  can be used, where the  $i$ -th piece  $P(i)$  defines the probability for two adjacent DBS pulses to be separated by a time in the range  $[(i - 1)\Delta, i\Delta]$  (within which a time delay  $d$  is uniformly randomly selected), and  $\sum_{i=1}^I P(i) = 1$ . In the simulations below,  $P$  is defined by  $I = 10$  constants, so that the search algorithm needs to simultaneously optimize 12 parameters: ten  $P(i)$  values, the pulse amplitude  $i_D$ , and the pulse width  $\sigma_D$ . We note that the relatively high dimension of this parameter space necessitates an efficient optimization algorithm.

**Multiple solutions:** If the dependence of therapeutic effects on DBS parameters is complex, there may be many different waveforms that are comparably effective. Algorithms such as the GA can simultaneously identify such multiple waveforms, opening the possibility of a therapeutic schedule

which cycles through several of these.

### 3 The Genetic Algorithm

Here we describe in detail how a GA could be used to identify optimal DBS waveforms in individual patients (Fig. 1). This strategy is tested in a computational model in the following section.

GAs are randomized, global search/optimization algorithms whose operation mimics the process of natural evolution [29]. The efficiency of a GA may be affected by both its algorithmic details and the nature of the optimization problem that it is asked to solve, but its fundamental capabilities are generic. Specifically, a GA is applicable to almost any optimization problem (i.e., regardless of the nature of the optimization landscape, including the presence of locally optimal but globally sub-optimal solutions). Also, the convergence rates for identifying optimal solutions does not, in general, slow dramatically as the dimensionality of the parameter space increases. Finally, a GA can automatically identify multiple parameter settings giving comparable performance. All these properties are likely to be important for clinical identification of novel DBS waveforms, and are demonstrated in the computational model below.

In general, implementation of a GA would begin by applying  $N$  trial, time-dependent DBS waveforms  $I_{DBS}^i(t)$ ,  $i = 1, 2, \dots, N$ , to the target brain area (e.g., STN or GPi cells). A quantitative measure  $x^i \geq 0$  characterizing PD symptoms in response to each  $I_{DBS}^i(t)$  would then be assessed. Using the value of  $x^i$  as the “cost”  $J^i$  associated with each  $I_{DBS}^i(t)$  (i.e.,  $J^i = x^i$ ), the GA replaces a certain percentage of the “bad” DBS inputs (i.e., those associated with high  $J^i$  values if the goal is to minimize  $J$ ) by new waveforms generated from “crossover” and “mutation” operations among the “good” ones [29]. These new waveforms are then tested, potentially leading to lower values of  $J$ . This iterative process continues until one or more DBS waveform  $I_{DBS}^*(t)$  provides the desired effect. The concept can readily be extended to employ multiple response metrics  $x_1^i, x_2^i, \dots$  for a given waveform  $I_{DBS}^i(t)$ . The GA used in the computational study described below has a population size ( $N$ ) of 25, replacement rate of 40%, crossover rate of 70%, and mutation rate of 30%.

The trial waveform  $I_{DBS}^i$  generally may be described by a vector of parameters  $\mathbf{a}^i$ , so that optimiz-

ing the criterion  $J$  over all possible DBS signals corresponds to optimizing with respect to  $\mathbf{a}$ . In some cases, the choice of  $\mathbf{a}^i$  can be made so that the initial pattern  $I_{DBS}^i$  represents DBS inputs that are previously known to be generally effective. Alternatively, the initial  $\mathbf{a}^i$  value can be chosen randomly. While the latter choice might seem less efficient intuitively, starting from a random population of trial DBS waveforms can allow for a more thorough exploration of the DBS parameter space, which might reduce the bias inherited from previous knowledge and hence might locate novel DBS waveforms. The simulations presented below provide such an example.

The GA’s cost function  $J$  can also contain terms that quantify practical constraints, such as side effects or battery usage. Therefore,  $J$  can be generally expressed as  $J^i = x^i + wR$ , where  $R > 0$  denotes the practical constraints, and  $w \geq 0$  is a corresponding weight factor to balance the significance of  $x^i$  and  $R$ .

## 4 Illustration in a Computational Model

### 4.1 The network model and its properties

Here and in [31], the closed-loop DBS optimization procedure described above is demonstrated on the subthalamopallidal network model developed by Terman, Rubin et al [20, 28] (see those publications for a complete description of the model; the “structured, sparse” connectivity pattern is used here). In brief, the Terman, Rubin et al model consists of neurons belonging to three basal ganglia nuclei: the STN, GPe, and GPi. There are eight STN cells and eight GPe cells coupled as follows: each STN cell receives inhibition from two GPe cells, and each GPe cell receives excitatory input from one STN, a constant bias, and inhibition from two adjacent GPe neurons. There are also eight GPi neurons, each receiving inhibition from one GPe and excitatory input from one STN cell. Synaptic coupling is via currents with appropriate timescales. As described in [20, 28], all three cell types are described by single-compartment conductance-based models in the Hodgkin-Huxley formalism, with membrane currents representative of cells in the different nuclei (unlike more complex models of [18, 21], the Terman, Rubin et al. model neglects trans-synaptic effects of DBS currents and the possibility of distinct effects of DBS on cell bodies and axons). DBS waveforms enter homogeneously as current

inputs  $I_{DBS}(t)$  to STN cells, and we focus on the effects of these currents on Gpi firing patterns.

We emphasize that the main purpose of this paper is not to identify any particular clinically effective DBS parameter setting via simulations, but rather to propose an overriding optimization strategy schematically shown in Fig. 1 that can explore the DBS parameter space more thoroughly and more efficiently than hand-tuning or grid sampling approaches, and to illustrate that alternatives to standard high-frequency stimulation may produce distinct and potentially beneficial effects. As a result, we believe that the basic conclusions made in this article about the use of the procedure in Fig. 1 are generally valid despite the incompleteness of the model itself. The details of the model’s dynamical response to various forms of DBS inputs are addressed elsewhere [31].

Figure 2 shows the distinct firing patterns produced by this model in its normal and Parkinsonian states [28], as well as the dynamic response of the model under various DBS conditions. Under normal conditions, all of the eight GPi neurons display irregular firing times that are only weakly correlated (Figure 2(c)). By contrast, an increase, motivated by the underlying physiology, in the constant bias current  $I_{GPe}$  and a decrease in lateral synaptic strengths  $I_{GPe \rightarrow GPe}$  switches the network from the normal to the Parkinsonian state (Fig. 2(d)) where the GPi cells demonstrate bursting spike patterns with a common frequency and characteristic clustering [28]. We quantify these firing patterns through a measure  $Cor$  based on the Gpi cells’ auto- and cross-correlograms (as detailed in [31]). Specifically,  $Cor$  is a weighted sum of the standard deviation of the bin counts in the auto- and cross- correlograms from 0 to 800 ms (with weights 2.0 and 1.0, respectively, and bin width 10 ms) and the first bin count of the cross-correlogram (with weight  $2.0 \times 10^{-5}$ ). In general, desynchronized and irregular firing across the neurons (corresponding to the normal state) leads to low  $Cor$  values, while synchronized and bursting spike patterns results in high  $Cor$  values. Specifically,  $Cor = 0.1$  in the normal state and  $Cor = 1.5$  in PD. Therefore, minimizing  $J = Cor$  is expected to identify DBS waveforms that bring firing patterns closer to their normal state.

## 4.2 Optimization of DBS waveforms

Figure 2(a) illustrates the dependence of the measure  $Cor$  on changes in the standard (square-wave) DBS parameters of pulse period  $\rho_D$  and current amplitude  $i_D$ , with fixed pulse duration  $\sigma_D$ . Very

high  $Cor$  values are observed in the high frequency, high amplitude region of the parameter space, resulting from the largely synchronized GPi spikes induced by standard high-frequency DBS [20]. Other settings of the parameters typically lead to lower  $Cor$  values, and in general the ‘landscape’ of  $Cor$  as a function of parameters  $\rho_D$  and  $i_D$  alone is highly complex, suggesting that, especially when additional parameters such as pulse duration  $\sigma_D$  are introduced, it will be difficult to identify optimal parameter settings without the guidance of a global search algorithm. This brings us to our first major point: **(1) The neural response to DBS parameters can be highly complex, so that global optimization algorithms are required to identify optimal solutions.**

A GA is first employed to search for periodic DBS parameter settings (period, amplitude, and duration) that minimize  $J = Cor$  in the Parkinsonian network. Figure 2(f) shows the raster plot of an optimal solution found by the GA ( $Cor = 0.13$ ). The GPi firing becomes much less regular and than the PD state (Fig. 2(d)), although it is still more burstlike than the normal state (Fig. 2(c)). The frequency of the underlying DBS current (10Hz) is significantly lower than that used in standard high-frequency DBS. In addition, the mechanism of action for this low-frequency input is different from that of high-frequency DBS: the former partially restores the irregular and desynchronized firing of the normal state, while the latter entrains the network to the high-frequency DBS signal (Fig. 2(e) and [20]). Furthermore, the mean amplitude of the low-frequency inputs is substantially lower than for high-frequency DBS simulated within this model. This carries the potential benefits for stimulator battery life and minimization of side effects discussed in Section 1. This is the content of our second main point: **(2) Computational basal ganglia models suggest that novel DBS waveforms exist which may return Parkinsonian network dynamics close to the normal state and exploit mechanisms that differ from those of high-frequency DBS.**

Clinical studies [32–34] have found that low-frequency DBS either has little impact or can worsen Parkinsonian symptoms. This discrepancy with the present computational model raises two possibilities. First, the model could simply be inaccurate in predicting the effectiveness of low-frequency periodic inputs. Second, it could be that effective low-frequency parameter settings exist clinically but were not sampled. The complex dependence of DBS performance on its parameters (Fig. 2(a)) suggests that this is plausible. In either case, it is prudent to investigate aperiodic inputs, avoiding the possibility, as noted by [34], of purely periodic stimulation near multiples of the tremor frequency

enhancing pathological rhythms.

With this in mind, we next explore the optimization of stochastic DBS currents. The additional motivation for doing this is severalfold. First, it is intuitively reasonable that random DBS inputs might be better suited than periodic ones for inducing irregular firing patterns (note that the general goal is to desynchronize the highly clustered neurons in the PD state and to make their firing less burstlike), so that they may be useful for patients who are unresponsive to standard, periodic DBS patterns. Next, in some cases clinical responsivity to high-frequency DBS can decrease over time and disease progression, so that ever stronger high-frequency DBS currents are needed to provide stable therapeutic effects [4, 7, 35, 36]. While this phenomenon is not simulated in the model, it is possible that stochastic currents may be more robust to gradual physiological changes than periodic currents. Lastly, since introducing stochastic currents can call for a larger number of parameters than periodic ones, it serves as a good test of the GA’s scalability.

We used the parameterization of the stochastic DBS currents described in Section 2. In the GA optimization, a current penalty is included in the cost function (i.e.,  $J = Cor + w \int_{t=0}^T I_{DBS} dt$ ) to specifically minimize the DBS current amplitude while restoring approximately normal spiking patterns. The results, demonstrating substantial desynchronization of Gpi firing patterns with much lower overall current amplitude, are shown in Fig. 3. They illustrate our third main point: **(3) Stochastic DBS waveforms can be effective alternatives to periodic inputs.**

The three panels of Fig. 3 also illustrate our fourth main point: **(4) The optimization procedure can discover multiple DBS currents that are considerably different while yielding very similar cost function values  $J$ .** The existence of multiple optimal currents also occurs with periodic current optimizations above, but the number of comparable DBS currents is much smaller than with the stochastic inputs, presumably as a result of the lower number of parameters (three in this case) defining the DBS currents.

Any optimization algorithm used in clinical settings must converge robustly in the presence of gradual drift in, for example, the firing rate and burst frequency of underlying network. Such robustness is guaranteed if the DBS waveforms identified by the algorithm achieve the desired effect for Parkinsonian networks in a wide “operating range” (i.e., with a wide set of burst frequencies). In

fact, this class of broadly effective waveform is the only type of solution to which a GA can reliably converge when it is optimizing a nonstationary (drifting) system.

In Fig. 4, we demonstrate that the stochastic DBS currents of Fig. 3 are robust in this sense. Here, we first change the bias current intrinsic to GPe cells,  $I_{GPe}$ , from its nominal Parkinsonian-state value of  $-13.0 \text{ pA}/\mu\text{A}^2$  to  $-13.2 \text{ pA}/\mu\text{A}^2$ , so that the cells' burst frequency approximately doubles from 4 Hz (Fig. 2(d)) to 10 Hz (Fig. 4(d)). We then apply the same stochastic DBS currents found by the GA to be effective at the nominal 4 Hz frequency. The performance of these currents is not degraded at the new, higher frequency (Cor values remain at  $\sim 0.13$ ), illustrating the robustness of the DBS effects. In fact, for five values of  $I_{GPe}$  sampled between  $-13.0 \text{ pA}/\mu\text{A}^2$  to  $-13.2 \text{ pA}/\mu\text{A}^2$ , the maximum value of Cor (calculated for all three DBS currents in Fig. 3) was 0.16, 23% above its nominal value. This illustrates: **(5) In the computational model, stochastic DBS waveforms are robust to variation in the intrinsic burst frequency, suggesting their utility in the nonstationary settings expected clinically.** We note that continued drifting of parameter values following the initial optimization of the DBS can in principle be compensated for by standard adaptive control techniques [30] (see Conclusion).

Another important question is whether the clinical optimization procedure that we envision converges within an acceptable period of time. The GA algorithm used in the present modeling study required on the order of 200-500 experiments (conducted over 8-20 iterations), each using a distinct DBS waveform setting, to converge to an optimal solution (Fig. 2(b)). Since the input  $I_{DBS}$  is described by three parameters for the periodic currents and 12 for the stochastic currents, this figure demonstrates that the requisite number of iterations does not increase substantially with the number of DBS parameters. This is summarized by our final main point: **(6) The GA requires evaluating a fairly modest number of parameter settings to converge, and this number seems stable as the complexity of the DBS waveform increases.**

The corresponding number of parameter settings that would be required clinically may differ from the simulations above for several reasons, including the fact that further refinements of DBS inputs are possible beyond those explored here, and including the choices amongst combinations of electrode contacts and monopolar vs. bipolar stimulation [22]. Nevertheless, we note that it should in principle

be possible to evaluate a hundred to a few hundreds of different parameter settings within several hours, at least against criteria such as tremor or upper limb rigidity that can be assessed over seconds to tens of seconds (cf. [22, 23]). The efficiency of the optimization can also be improved by global-local hybrid search algorithms, such as the simplex method [37], or by combining the present approach with high-dimensional global mapping techniques [38] during iterative optimization.

## 5 Conclusion

This paper proposes that a closed-loop global optimization algorithm may be effective for identifying novel DBS waveforms that diminish rhythmic, burstlike activity characterizing the Parkinsonian basal ganglia. Some general issues relevant to eventual clinical or experimental tests of this algorithm were highlighted, including with an illustration of the optimization procedure in a computational model. A GA's model-free structure, global search capacity, and favorable scaling with the number of optimization parameters all make it well suited to the present application. We applied the concept to a biophysically-based neural network model to show the capability of identifying DBS waveforms (both periodic and non-periodic) that transform the burstlike, synchronized firing patterns characteristic of basal ganglia neurons in the Parkinsonian state into increasingly irregular, desynchronized patterns that more closely resemble those of neurons in the normal state.

The specific waveforms identified in the simulations of the paper are not intended for direct transfer to clinical use due to the model simplifications. However, this is a separate issue from the effectiveness of the algorithm in identifying alternative DBS waveforms. As illustrated in the previous sections, the favorable features of the algorithm make it especially attractive when a large number of parameters are required to represent the DBS waveforms (which almost inevitably lead to a highly complex relationship between DBS inputs and their effects). In addition, the same strategy can also be applied to enable automatic and efficient stabilization of the therapeutic effects during chronic treatments. In this case, the stimulator would operate initially using the algorithm-discovered optimal waveforms from the clinic (or standard high-frequency inputs, if alternative waveforms do not provide clear advantages [23]). Any change in the PD measure (e.g., due to disease progression) over a pre-defined threshold will then activate the search algorithm, which automatically adjusts the DBS settings of

the stimulator to return the PD measure to normal values. Since changes of this kind usually are not expected to be abrupt in nature, the operation likely may be achieved by linear feedback algorithms [30] hard-wired into a carry-on stimulator, thereby lessening the need for clinical visits. We hope that the general algorithm and the issues raised in this paper will inspire experimental tests of the closed-loop optimization strategy as the next step toward its potential clinical implementation.

## Acknowledgment

The authors acknowledge support from the National Science Foundation, and the USEPA-funded Environmental Bioinformatics and Computational Toxicology Center (ebCTC) under STAR Grant number GAD R 832721-010. Eric Shea-Brown is also supported by a Career Award at the Scientific Interface from the Burroughs Wellcome Fund. We thank Drs. Jonathan Rubin, David Terman, Karen Sigvardt, and Vikki Wheelock for important comments, advice, and references throughout this project.

## References

- [1] R.E. Gross and A.M. Lozano. Advanced in neurostimulation for movement disorders. *Neurol. Res.*, 22:247–258, 2000.
- [2] The deep brain stimulation for Parkinson’s disease study group. Deep-brain stimulation of the subthalamus nucleus or the pars interna of the globus pallidus in parkinson’s disease. *New England J. Med.*, 345:956–963, 2001.
- [3] A. Benabid. Deep brain stimulation for Parkinson’s disease. *Curr. Opin. Neurobiol.*, 13:696–706, 2003.
- [4] P. Krack, A. Batir, N. van Blercom, S. Chabardes, V. Fraix, and et al. C. Ardouin. Five-year follow-up of bilateral stimulation of the subthalamic nucleus in advanced Parkinson’s disease. *N. Engl. J. Med.*, 349:1925–1934, 2003.
- [5] G. Kleiner-Fisman, D.N. Fisman, E. Sime, J.A. Saint-Cyr, A.M. Lozano, , and A.E. Lang. Long-

- term of bilateral deep brain stimulation of the subthalamic nucleus in patients with advanced Parkinson's disease. *J. Neurosurg.*, 99:489–495, 2003.
- [6] M.C. Rodriguez-Oroz, I. Zamarbide, J. Guridi, M.R. Palmero, and J.A. Obeso. Efficacy of deep brain stimulation of the subthalamic nucleus in Parkinson's disease 4 years after surgery: double blind and open label evaluation. *J. Neurol. Neurosurg. Psychiatry*, 75:1382–1385, 2004.
- [7] M.C. Rodriguez-Oroz, J.A. Obeso, A.E. Lang, J.L. Houeto, and et al. P. Pollak. Bilateral deep brain stimulation in Parkinson's disease: a multicentre study with 4 years follow-up. *Brain*, 128:2240–2249, 2005.
- [8] W. Olanow, M. Brin, and J. Obeso. The role of deep brain stimulation as a surgical treatment for Parkinson's disease. *Neurology*, 55 (Supp.6):S60–S66, 2000.
- [9] A. Benabid, A. Koudsie, A. Benazzouz, B. Piallat, P. Krack, P. Limousin-Dowsey, J. Lebas, and P. Pollak. *Advances in neurology, Vol 86: Parkinson's disease*, chapter Deep brain stimulation for Parkinson's disease. Lippincott Williams & Wilkins, Philadelphia, 2001.
- [10] A. Benazzouz, D. Gao, Z. Ni, B. Piallat, R. Bouali-Benazzouz, and A. Benabid. Effect of high-frequency stimulation of the subthalamic nucleus on the neuronal activities of the substantia nigra pars reticulata and the ventrolateral nucleus of the thalamus. *Neuroscience*, 99:289–295, 2000.
- [11] T. Boraud, E. Bezard, B. Bioulac, and C. Gross. High frequency stimulation of the internal globus pallidus (gpi) simultaneously improves parkinsonian symptoms and reduces the firing frequency of gpi neurons in the mptp treated monkey. *Neurosci. Lett.*, 215:17–20, 1996.
- [12] J. Dostrovsky, R. Levy, J. Wu, W. Hutchison, R. Tasker, and A. Lozano. Microstimulation-induced inhibition of neuronal firing in human globus pallidus. *J. Neurophysiol.*, 84:570–574, 2000.
- [13] G. Paul, T. Reum, W. Meissner, A. Marburger, R. Sohr, R. Morgenstern, and A. Kupsch. High frequency stimulation of the subthalamic nucleus in uences striatal dopaminergic metabolism in naive rats. *NeuroReport*, 11:441–444, 2000.

- [14] F. Windels, N. Bruet, A. Poupard, N. Urbain, G. Chouvet, C. Feuerstein, and M. Savasta. Effects of high frequency stimulation of subthalamic nucleus on extracellular glutamate and gaba in substantia nigra and globus pallidus in the normal rat. *Eur. J. Neurosci.*, 12:4141–4146, 2000.
- [15] M. Anderson, N. Postpuna, and M. Ruffo. Effects of high frequency stimulation in the internal globus pallidus on the activity of thalamic neurons in the awake monkey. *J. Neurophysiol.*, 89:1150–1160, 2003.
- [16] T. Hashimoto, E. Elder, M. Okun, S. Patrick, and J. Vitek. Stimulation of the subthalamiaic nucleus changes the firing pattern of pallidal neurons. *J. Neurosci.*, 23:1916–1923, 2003.
- [17] N. Maurice, A. Thierry, J. Glowinski, and J. Deniau. Spontaneous and evoked activity of substantia nigra pars reticulata neurons during highfrequency stimulation of the subthalamic nucleus. *J. Neurosci.*, 23:9929–9936, 2003.
- [18] C. McIntyre, W. Grill, D. Sherman, and N.V. Thakor. Cellular effects of deep brain stimulation: model-based study of activation and inhibition. *J. Neurophysiol.*, 91:1457–1469, 2004.
- [19] E. Montgomery and K. Baker. Mechanisms of deep brain stimulation and future technical developments. *Neurol. Res.*, 22:259–266, 2000.
- [20] J. Rubin and D. Terman. High frequency stimulation of the subthalamic nucleus eliminates pathological thalamic rhythmicity in a computational model. *J. Comput. Neurosci.*, 16:211–235, 2004.
- [21] P. Hahn, D. Lee, G. Russo, J. Vitek, and C. McIntyre. Stimulation on a model of subthalamopallidal network activity. *Soc. Neurosci. Abstr.*, page 331.6, 2005.
- [22] J. Volkmann, J. Herzog, F. Kopper, and G. Deuschl. Introduction to the programming of deep brain stimulators. *Movement Disorders*, 17 (Suppl.3):S181–S187, 2002.
- [23] M. Rizzone, M. Lanotte, B. Bergamasco, A. Tavella, E. Torre, G. Faccani, A. Melcarne, and L. Lopiano. Deep brain stimulation of the subthalamic nucleus in Parkinson’s disease: effects of variation in stimulation parameters. *J. Neurol. Neurosurg. Psychiatry*, 71:215–219, 2001.

- [24] P.A. Tass. *Phase resetting in medicine and biology. Stochastic modeling and data analysis.* Springer, Berlin, 1999.
- [25] P.A. Tass. Desynchronizing double-pulse phase resetting and application to deep brain stimulation. *Biol. Cybern.*, 85:343–354, 2001.
- [26] O.V. Popovych, C. Hauptmann, and P.A. Tass. Effective desynchronization by nonlinear delayed feedback. *Phys. Rev. Lett.*, 94:164102–1–4, 2005.
- [27] C. Hauptmann, O.V. Popovych, and P.A. Tass. Delayed feedback control of synchronization in locally coupled neuronal networks. *Neurocomputing*, 65:759–767, 2005.
- [28] D. Terman, J. Rubin, A. Yew, and C.J. Wilson. Activity patterns in a model for the subthalamo-pallidal network of the basal ganglia. *J. Neurosci.*, 22:2963–2976, 2002.
- [29] D. Goldberg. *Genetic algorithms in search, optimization and machine learning.* Addison-Wesley, Boston, MA., 1989.
- [30] G. Franklin, J.D. Powell, and A. Emami-Naeini. *Feedback control of dynamic systems.* Prentice Hall, New Jersey, 2005.
- [31] X. Feng, B. Greenwald, H. Rabitz, E. Shea-Brown, and R. Kosut. Optimal deep brain stimulation of the subthalamic nucleus – a computational study, *submitted to J. Comput. Neurosci.*
- [32] L. et al. Wojtecki. Frequency-dependent reciprocal modulation of verbal fluency and motor functions in subthalamic deep brain stimulation. *Arch. Neurol.*, 63:1273–1273, 2006.
- [33] E. Moro, R. J.A. Esselink, J. Xie, M. Hommel, A. L. Benabid, and P. Pollak. The impact on parkinsons disease of electrical parameter settings in stn stimulation. *Neurology*, 59:706–713, 2002.
- [34] L. et al Timmerman. The impact on parkinsons disease of electrical parameter settings in stn stimulation. *Movement Disorders*, 19:1328–1333, 2004.

- [35] K. Lyons, W. Koller, S. Wilkinson, and R. Pahwa. Long term safety and efficacy of unilateral deep brain stimulation of the thalamus for parkinsonian tremor. *J Neurol Neurosurg Psychiatry.*, 71:682–4, 2001.
- [36] M. Hariz, P. Shamsgovara, F. Johansson, G. Hariz, and H. Fodstad. Tolerance and tremor rebound following long-term chronic thalamic stimulation for parkinsonian and essential tremor. *Stereotact. Funct. Neurosurg.*, 72:208–18, 1999.
- [37] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical recipes in C, second edition.* Cambridge University Press, NY, 1992.
- [38] G. Li, C. Rosenthal, and H. Rabitz. High dimensional model representations. *J. Phys. Chem. A*, 105:7765–7777, 2001.




Figure 1: Schematic of the closed-loop DBS current waveform optimization in clinical applications. A set of trial DBS currents  $I_{DBS}^i, i = 1, 2, \dots, N$  (with each member defined by the parameter vector  $\mathbf{a}^i$ ) are first generated by the stimulator and applied individually to the target brain area. The physiological responses  $\{x^i\}$  of the patient are then recorded for each current. The genetic algorithm evaluates these responses together with practical stimulator constraints  $R$  and uses them to guide the selection of the next set of trial parameters  $\{\mathbf{a}^i\}$  that may lead to better responses  $\{x^i\}$ . This iterative process continues until one or more DBS currents, represented by  $\mathbf{a}^*$  are found that satisfactorily achieve clinical objectives.




Figure 2: Left panel: (a) The ‘rugged’ dependence of  $Cor$ , the statistical measure of GPI cells’ spike patterns [31], on DBS pulse period  $\rho_D$  and current amplitude  $i_D$ . The impulse duration  $\sigma_D$  is fixed at 0.6ms. The region where  $Cor > 0.7$  is truncated to show the complex landscape at lower frequencies. The minimum  $I_{DBS}$  amplitude in the plot is  $10 pA/\mu m^2$ ;  $Cor$  values gradually increase to the unstimulated value of 1.5 for still lower amplitudes. (b) Convergence of the GA cost function over the iterations. Solid line: cost function  $J = Cor$  vs. GA iteration for optimization over the three standard DBS parameters – pulse frequency, duration, and amplitude. The GA converges after approximately 10 iterations. Dashed line: optimization of the 12 parameters describing stochastic DBS currents ( $J = Cor + w \int_{t=0}^T I_{DBS} dt$ ). The GA requires  $\sim 15$  iterations. Right panel: Raster plot of the spike times for the eight GPI cells in the (c) normal state, (d) PD state, (e) PD state with high-frequency DBS current ( $i_D = 200 pA/\mu m^2$ ,  $\rho_D = 6 ms$ ,  $\sigma_D = 0.6 ms$ , and average  $\langle I_{DBS} \rangle = 20 pA/(\mu m^2 ms)$ ), and (f) PD state with GA-optimized low-frequency DBS current ( $\rho_D = 100 ms$ ,  $i_D = 40 pA/\mu m^2$ ,  $\sigma_D = 0.6 ms$ , giving  $\langle I_{DBS} \rangle = 0.24 pA/(\mu m^2 ms)$ ).




Figure 3: Optimized stochastic DBS currents with  $J = Cor + w \int_{t=0}^T I_{DBS} dt$ . Panels (a), (b), and (c) illustrate three currents with comparable  $J$  values found from the optimization algorithm. For each panel, plots from the top are: the piece-wise probability distribution function used to generate the stochastic current (see section 2 for details), the corresponding DBS current, and the resulting GPi raster plot.  $Cor \simeq 0.13$  and averaged current values  $\langle I_{DBS} \rangle \simeq 0.4 pA/(\mu m^2 ms)$  for all three cases.




Figure 4: Robustness of the optimal DBS currents. Panels (a), (b), and (c) show the GPi raster plots for the three stochastic DBS currents in Fig. 3, when  $I_{GPe}$  is changed from  $-13.0 pA/\mu A^2$  to  $-13.2 pA/\mu A^2$ . Panel (d) shows the GPi raster plot of the PD state with  $I_{GPe} = -13.2 pA/\mu A^2$  and no DBS input. The firing frequency of GPi cells increases from  $\sim 4$  Hz in Fig. 2(d) to  $\sim 10$  Hz here.