Optimal Deep Brain Stimulation of the Subthalamic Nucleus – a Computational Study

Xiao-jiang Feng[†], Brian Greenwald, and Herschel Rabitz
Department of Chemistry, Princeton University, Princeton, NJ 08544

Eric Shea-Brown[†] *
Courant Institute of Mathematical Sciences
and Center for Neural Science,
New York University
251 Mercer Street, New York, NY 10012

Robert Kosut
SC Solutions, Sunnyvale, CA 94085

December 20, 2006

[†] These authors contributed equally to this work.

*To whom correspondence should be addressed: 251 Mercer Street, New York, NY 10012, Phone: (212)998-3132, Fax: (212) 995-4011, Email: ebrown@math.nyu.edu

Abbreviated title: An Algorithm for Optimal Deep Brain Stimulation

Keywords: Deep Brain Stimulation, Parkinson's disease, basal ganglia, subthalamic nucleus, numerical optimization, neural network model

Number of figures: 12

Number of tables: 0

Number of pages: 38 (double spaced)

Abstract

Deep brain stimulation (DBS) of the subthalamic nucleus, typically with periodic, high frequency pulse trains, has proven to be an effective treatment for the motor symptoms of Parkinson's disease. Here, we use a biophysically-based model of spiking cells in the basal ganglia (Terman et al., 2002) to provide computational evidence that alternative temporal patterns of DBS inputs might be equally effective as the standard high-frequency waveforms, but require lower amplitudes. Within this model, DBS performance is assessed in two ways. First, we determine the extent to which DBS causes Gpi synaptic outputs, which are burstlike and synchronized in the unstimulated Parkinsonian state, to cease their pathological modulation of simulated thalamocortical cells. Second, we evaluate how DBS affects the GPi cells' auto- and cross-correlograms. In both cases, a nonlinear closed-loop learning algorithm identifies effective DBS inputs that are optimized to have minimal strength. The network dynamics that result differ from the regular, entrained firing which some previous studies have associated with conventional high-frequency DBS. This type of optimized solution is also found with heterogeneity in both the intrinsic network dynamics and the strength of DBS inputs received at various cells. Such alternative DBS inputs could potentially be identified, guided by the model-free learning algorithm, in experimental or eventual clinical settings.

Introduction

During high-frequency deep brain stimulation (DBS) therapy for Parkinson's disease (PD), rhythmic ($> 100\ Hz$) pulsatile voltage transients are typically applied to the subthalamic nucleus of the basal ganglia. The therapeutic effects are often dramatic, alleviating motor symptoms and decreasing dependence on dopaminergic drugs (Krack et al., 2003; Kleiner-Fisman et al., 2003; Rodriguez-Oroz et al., 2004; Rodriguez-Oroz et al., 2005; Deep Brain Stimulation for Parkinson's disease study group, 2001; Benabid, 2003), and the corresponding

physiological effects are under intensive study – specifically, how DBS modulates activity in the target areas.

One possible answer is suggested by the similar effects of GPi lesions and DBS of the STN: that DBS suppresses firing (STN excites Gpi; see Fig. 1) (Benabid, 2003; Olanow, Brin, and Obeso, 2000; Benabid et al., 2001; Benazzouz et al., 2000; Boraud et al., 1996; Beurrier et al., 2001; Benabid, 2003). However, some experiments indicate that high-frequency DBS can *enhance* output of the stimulated areas (Windels et al., 2000; Hashimoto et al., 2003; Maurice et al., 2003). This discrepancy is addressed by the computational modeling of (McIntyre et al., 2004), which indicates that DBS can suppress firing of cell bodies while exciting axons at approximately the frequency of the DBS drive.

If high frequency DBS enhances synaptic output, its therapeutic effect cannot be due to its mimicking a lesion. An alternative hypothesis, and that adopted here, is that DBS otherwise modulates the activity of the BG network to "mask" the pathological firing patterns that characterize the Parkinsonian state (Benabid, 2003; McIntyre et al., 2004; Hashimoto et al., 2003; Montgomery and Baker, 2000; Rubin and Terman, 2004). Specifically, GPi neurons in the Parkinsonian vs. normal state are more synchronized, rhythmic, and burstlike ((Terman et al., 2002; Bergman et al., 1998; Magnin, Morel, and Jeanmonod, 2000; Nini et al., 1995)). Using a conductance-based network model, (Rubin and Terman, 2004) demonstrate that periodic, high-frequency DBS can elicit similarly periodic, high-frequency firing of STN and GPi cells, thereby replacing Parkinsonian firing with a mask of tonic activity (cf. (Hahn et al., 2005)). These authors show that this tonic Gpi activity restores the simulated function of downstream (thalamocortical) cells.

Other computational studies (see the Discussion Section) seek DBS inputs that achieve similar "masking" effects but have alternative temporal patterns, often requiring lower amplitudes than high-frequency DBS (Tass, 1999; Tass, 2001; Rosenblum and Pikovsky, 2004; Popovych, Hauptmann, and Tass, 2005; Hauptmann, Popovych, and Tass, 2005). As these authors emphasize, considering alternatives to high-frequency DBS is important for several reasons. If clinically effective lower-amplitude DBS inputs were identified, their use would pre-

serve stimulator batteries (which require surgical replacement), and might also lessen DBS side effects. Additionally, an open question is whether alternative patterns would be effective in patients who do not respond effectively to high-frequency DBS (Okun et al., 2005), or whose symptoms nevertheless worsen over time (Krack et al., 2003).

We propose and computationally test a new approach to identifying such alternatives to high-frequency DBS: using a nonlinear closed-loop learning algorithm to search for DBS inputs that *optimally* alleviate Parkinsonian firing patterns with the minimal possible amplitudes. The proposed algorithm conducts a global search to identify optimal DBS inputs drawn from any user-specified class of possible waveforms, and hence has the capability to identify highly novel solutions. We apply the algorithm to the computational subthalamopallidal model of (Terman et al., 2002). Our goals are twofold: to determine whether alternative DBS patterns exist within this model that have comparable or superior performance to high-frequency inputs, and whether closed-loop learning algorithms can effectively search the space of possibilities to identify these optimal DBS inputs. We find positive answers on both accounts.

The balance of the paper proceeds as follows. We first review the model of (Terman et al., 2002), and then undertake a systematic study of the influence of standard DBS parameters (cf. (Rubin and Terman, 2004; Rizzone et al., 2001)). Surprisingly, we find that certain DBS currents with relatively low frequency and amplitude desynchronize GPi neurons, disrupting rhythmic patterns characteristic of PD and restoring function in stimulated thalamocortical cells. Subsequent closed-loop optimization identifies DBS currents which simultaneously achieve this effect and minimize average levels of DBS currents. Next, to relax assumptions about thalamocortical function, a general statistical measure of the GPi cells' firing pattern is devised as the criterion to distinguish normal and Parkinsonian states. This measure is then utilized in the closed-loop learning algorithm to directly identify DBS currents that optimally transform the synchronized, burstlike firing dynamics of the Parkinsonian cells into firing patterns whose statistics match those of normal cells. In this process, we consider both standard periodic DBS inputs and more general DBS currents whose pulse timings are defined by either an optimally determined 12-parameter probability distribution function, or by a 16-parameter

discrete, periodic (Haar) basis. The effectiveness of these latter types of DBS input is of interest, because it suggests evaluating a greater variety of DBS patterns in experimental settings. We also demonstrate that optimal solutions can be found in the presence of heterogeneity in both the intrinsic network dynamics and the strength of DBS inputs received at various cells.

Model and Methods

The network model

The subthalamopallidal network model we use was developed and compared with experimental firing patterns in (Terman et al., 2002) and applied to a study of high-frequency DBS in (Rubin and Terman, 2004). It consists of model neurons belonging to three basal ganglia nuclei: the STN, GPe (external segment of globus pallidus), and GPi, as well as thalamocortical relay (TC) cells. We adopt the "sparsely connected, structured" architecture (Fig. 1) to reproduce both tonic and bursting firing modes, as in (Terman et al., 2002). There are eight STN cells and eight GPe cells: each STN cell receives inhibition from two GPe cells as well as the external DBS current, and each GPe cell receives excitatory input from one STN, an input from the (unmodelled) striatum, 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 (Rubin and Terman, 2004). Finally, two TC cells are included in the model, each receiving synaptic inhibition from four GPi cells and an excitatory input representing sensorimotor signals (Rubin and Terman, 2004).

The details of network connectivity are described in (Terman et al., 2002; Rubin and Terman, 2004) and Fig. 1. All of the network simulations presented here are run for 2500ms (to allow the effects of transients and initial conditions to decay) before plots are produced (over the following 800 ms) and statistics are computed over the following 6500 ms.

Conductance-based neuron equations

As in (Terman et al., 2002; Rubin and Terman, 2004), all four cell types are described by single-compartment conductance-based neuron models motivated by the underlying physiology. STN, GPe, GPi, and TC are used as subscripts to denote variables and functions corresponding to the various cell types. The synaptic current from cell type α to β is denoted as $I_{\alpha\to\beta}$. The parameter settings, including initial conditions, are as described in (Terman et al., 2002; Rubin and Terman, 2004).

The voltage of the STN neurons is modelled via

$$C_m\frac{dv_{STN}}{dt} = -I_L - I_K - I_{Na} - I_T - I_{Ca} - I_{AHP} - I_{GPe\to STN} + I_{STN} + I_{DBS}, \tag{1}$$

where I_L is the leak current, I_K is the potassium current, I_{Na} is the sodium current, I_T is a low-threshold T-type Ca^{2+} current, I_{Ca} is the high-threshold Ca^{2+} current, I_{AHP} is a Ca^{2+}-activated, voltage-independent afterhyperpolarization K^+ current, $I_{GPe\to STN}$ is the synaptic input from GPe to STN, I_{STN} is a constant depolarizing current, and I_{DBS} is the DBS current input. Most of these currents (and those below) are determined by the standard auxiliary differential equations; see (Terman et al., 2002) for details.

The voltage of the GPe neurons follow

$$C_m\frac{dv_{GPe}}{dt} = -I_L - I_K - I_{Na} - I_T - I_{Ca} - I_{AHP} - I_{GPe\to GPe} - I_{STN\to GPe} + I_{GPe}, \tag{2}$$

where $I_{GPe\to GPe}$ represents the interpallidal inhibition between GPe cells, $I_{STN\to GPe}$ is the synaptic excitation from STN to GPe, and I_{GPe} is a constant depolarizing current. The voltage dynamics of the GPi cells is similar:

$$C_m\frac{dv_{GPi}}{dt} = -I_L - I_K - I_{Na} - I_T - I_{Ca} - I_{AHP} - I_{GPe\to GPi} - I_{STN\to GPi} + I_{GPi}, \tag{3}$$

where $I_{GPe\to GPi}$ and $I_{STN\to GPi}$ denote the synaptic input from GPe and STN to GPi, respectively, and I_{GPi} is a constant depolarizing current. The TC cells are modelled via

$$C_m\frac{dv_{TC}}{dt} = -I_L - I_K - I_{Na} - I_T - I_{GPi\to TC} + I_{SM}, \tag{4}$$

where $I_{GPi\to TC}$ is the synaptic input from GPi cells to TC, and I_{SM} represents the sensorimotor input to the thalamus and is modeled as

$$I_{SM} = i_{SM} H(sin(2\pi t/\rho_{SM})) \times [1 - H(sin(2\pi(t + \delta_{SM})/\rho_{SM}))], \tag{5}$$

where H is the Heaviside step function; that is, $H(x) = 0$ if $x < 0$ and $H(x) = 1$ if $x \geq 0$ (Rubin and Terman, 2004). Here, ρ_{SM} is the period of I_{SM} and i_{SM} is its amplitude. In the simulations, the parameter values for I_{SM} are $i_{SM} = 10.0 pA/\mu m^2$, $\rho_{SM} = 50.0 ms$ and $\delta_{SM} = 5.0 ms$.

Although we use overlapping notation above, the membrane currents (e.g. I_T) differ as appropriate for the distinct types of neurons. Their voltage and time dependence is as described in (Terman et al., 2002; Rubin and Terman, 2004); the code used here, containing all parameter settings and equations, is available upon request.

Parameterizing DBS waveforms

Standard DBS waveforms

The periodic DBS waveforms typically used in clinical settings are modeled, as in (Rubin and Terman, 2004), by applying to STN cells a square-wave pulse train described by

$$I_{DBS} = i_D H(sin(2\pi t/\rho_D)) \times [1 - H(sin(2\pi(t + \sigma_D)/\rho_D))] \ . \tag{6}$$

Here there are three parameters: i_D is the stimulation amplitude, ρ_D is the stimulation period, and σ_D is the duration of each impulse.

Stochastic DBS waveforms

We also consider DBS waveforms consisting of pulse trains with *random* inter-pulse delays. Specifically, we specify a probability density function (PDF) $P(d)$ (satisfying $\int P(d)\,\mathrm{d}d = 1$) which defines the probability for two adjacent DBS pulses to be spaced by a time delay d. A piecewise constant PDF $P(i), i = 1, 2, \ldots, I$ is 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 randomly selected), and $\sum_{i=1}^{I} P(i) = 1$. Here, P is defined

by $I = 10$ constants, with $\Delta = 5$ms (so that the maximum pulse delay is 50ms). Such DBS waveforms are described by 12 parameters: ten $P(i)$ values, the pulse amplitude i_D, and the pulse width σ_D.

Nonpulsatile DBS waveforms

We also consider more general(periodic) DBS waveforms represented by a Haar basis. First, a parameter τ is used to define the period of summed basis functions so that for $i\tau \leq t < (i+1)\tau \leq T$ $(i = 0, 1, \ldots)$, $x = (t - i\tau)/\tau$. Within each period, the DBS current is

$$I_{DBS}(x) = \sum_{j=0}^{P} \sum_{k=0}^{2^j - 1} c_{jk} \psi_{jk}(x), \tag{7}$$

where $\psi_{jk}(x) \equiv \psi(2^j x - k)$ and

$$\psi(x) = \begin{cases} 1 & 0 \leq x < 1/2 \\ -1 & 1/2 \leq x < 1 \\ 0 & \text{otherwise} \end{cases}$$

To define the $j - th$ order Haar basis, 2^j parameters are needed. Haar bases up to the third order (i.e., $P = 3$) are used in our simulations, thereby the total number of GA parameters is $1 + 2 + 4 + 8 + 1 = 16$, where one parameter determines the period τ.

The closed-loop optimization algorithm

We use a Genetic Algorithm (GA) (Goldberg, 1989), a type of global optimization method, to iteratively search for periodic or aperiodic DBS currents that alleviate the simulated Parkinsonian condition in the network model, where we consider different criteria that assess the extent to which this objective has been achieved. In the first step of the GA, N different trial DBS currents $I^i_{DBS}(t)$, $i = 1, 2, \ldots, N$ are applied to the STN cells, and a quantitative measure $x^i \geq 0$ characterizing PD symptoms in response to each $I^i_{DBS}(t)$ is computed from the simulated neural dynamics. Two different measures are used here. The first measure, *Rel*, represents the reliability of TC cells in transmitting sensorimotor signals (as in (Rubin and

Terman, 2004)); the second measure, Cor, compares statistical properties of GPi cells in the presence of DBS stimulation with those in the normal state. The details will be described in the Results Section. The possible currents I^i_{DBS} are usually described by vectors of parameters $\mathbf{a}^i$, so that optimizing the criterion x over all possible DBS currents corresponds to optimizing with respect to $\mathbf{a}$.

The overall quality of each I^i_{DBS} is determined by a cost function J^i which guides the GA optimization. In the simplest case, $J^i = x^i$. In real applications, the algorithm also needs to take into account the practical constraints and requirements, such as the desirable property of minimizing the average and/or peak magnitude of DBS current to alleviate side effects and reduce battery usage. An extra term $R \geq 0$ representing such factors to be minimized can be added to the cost function, so that $J^i = x^i + wR$, where $w \geq 0$ is a weight parameter. The GA then compares J^i for all trial DBS currents and selects a certain percentage with the best effects (i.e., the minimal J^i values) to generate the next set of N trial DBS currents by "crossover" and "mutation" operations (Goldberg, 1989). This iterative optimization process continues until one or a few DBS currents are found that achieve satisfactory outcomes. Because of the computational demands of numerically solving the differential equations introduced above, the GA is implemented on a parallel computer cluster with 26 nodes. The GA used here has a population size (N) of 25, replacement rate of 40%, crossover rate of 70%, and mutation rate of 30%. The advantages of the GA relative to grid-sampling or local search algorithms will be discussed in the Results Section.

Results

Normal and Parkinsonian states, and the effect of high-frequency DBS

It was demonstrated in (Terman et al., 2002; Rubin and Terman, 2004) that the model described above can produce several different firing patterns, depending on parameter settings representing normal vs. Parkinsonian states (see below). Fig. 2 and Fig. 3 (a) and (b) illustrate

the behavior of the full network in the normal state. The voltage traces of the first and second (according to the indexing in Fig. 1) neuron of each cell type are plotted. STN, GPe, and GPi cells all display irregular firing times that are only weakly correlated (Fig. 2), and TC cells respond faithfully to the excitatory sensorimotor input (Fig. 3 (b)), meaning, as in (Rubin and Terman, 2004), that only one TC voltage spike appears "immediately" after each I_{SM} impulse, and that there are no extraneous TC voltage spikes (i.e., those not triggered by an I_{SM} impulse). As in (Rubin and Terman, 2004), this reliability is quantified by Rel, the ratio of I_{SM} pulses that lead to one spike in a TC cell *only* within 0.25ms after the end of the pulse; for the normal state, $Rel = 0.99$. This desynchronized and irregular normal state is also illustrated by the spike time raster plot of all eight GPi cells in Fig. 3(a). Similar plots will be used below to represent and distinguish different states of the network's dynamics.

As in (Rubin and Terman, 2004), a physiologically motivated increase in the constant bias current I_{GPe} and a decrease in lateral synaptic strengths $I_{GPe \to GPe}$ switches the network from the normal to the Parkinsonian state. Fig. 3(c) shows that GPi cells demonstrate bursting spike patterns with a characteristic clustering (see (Terman et al., 2002) and references therein); STN and GPe cells display similar clustering and bursting. These groups are organized as follows (see Fig. 1): cells $\{1, 5\}, \{2, 6\}, \{3, 7\}, \{4, 8\}$ are almost exactly synchronized, and the first and second (and third and fourth) of these pairs themselves fire periodically with only a small phase lag. This pairing may be understood directly from the network architecture: each cell within a pair of STN cells (STN cells $\{1, 5\}$, etc.) receives the same input from other cells in the network, no matter how desynchronized these other cells might be (although the grouping is preserved upon introducing mild network heterogeneity, a case we also consider below). Therefore, if the pair of STN cells tends to entrain to this input, the pair will become synchronized; the simulations show that this is what occurs in the simulated Parkinsonian state (while entrainment does not occur in the normal condition). The resulting partially synchronized, burstlike output of GPi cells compromises the faithful relay of the sensorimotor input by model TC neurons, as bursts of incoming inhibition suppress TC cells' membrane voltage, leading (via thalamic I_T currents) to rebound bursts of TC spikes which follow offset of this inhibition but which are

not related to the I_{SM} signal (see (Rubin and Terman, 2004), and the present Fig. 3(d)). The resulting reduced reliability of TC cell firing is $Rel = 0.43$.

High-frequency DBS of the STN cells, as is typically clinically, is modeled via Eqn. (6). Three parameters characterize such square-wave, periodic wavefroms: the amplitude i_D, period ρ_D, and impulse duration σ_D. Fig. 3(e) shows the dynamics of the GPi neurons during high-frequency DBS. Despite the high frequency inhibition from the GPi, the TC cells completely recover the ability to reliably transmit sensorimotor signals (Fig. 3(f)): the corresponding Rel value is ~ 1.0. As explained in (Rubin and Terman, 2004), this reliable transmission is restored because the intrinsic dynamics of the model network is essentially entrained by the strong, high frequency DBS input, which elicits regular, high-frequency firing in STN and hence in Gpe and Gpi cells. We note that DBS induces similar neural dynamics in the network model of (Hahn et al., 2005). This entrainment eliminates (or, in the words of (McIntyre et al., 2004), "masks") the Parkinsonian bursts of inhibition to TC cells that, as explained above, compromise their transmission of model sensorimotor inputs.

Influence of standard DBS parameters: from entrained to desynchronized network dynamics

Rubin et al. (Rubin and Terman, 2004) studied how the transmission of sensorimotor signals I_{SM} in their model depends on the parameters of simulated DBS currents. We start with a similar analysis, but over increased dynamic ranges for the standard DBS parameters and while covarying the stimulation period ρ_D and amplitude i_D (while pulse duration σ_D remains fixed). Fig. 4 illustrates the dependence of TC cell transmission reliability on the DBS pulse period ρ_D and amplitude i_D, via the measure Rel described above. It can be seen that reliable signal transmission of the TC cells mostly occurs in the high frequency (low period), high amplitude region of the parameter space, which is consistent with the results in (Rubin and Terman, 2004). However, high reliability values are also found in a small region corresponding to low impulse frequencies and moderate amplitudes: this is the "peak" around $\rho_D \approx 80\ ms$, $i_D \approx 80\ pA/\mu m^2$ in Fig. 4. The network dynamics in response to DBS with these parameters is

shown in Figs. 5 (a) to (c). Note that, unlike the high-frequency DBS of Fig. 3(e-f), Gpi spike times are not entrained to the DBS inputs. Rather, the lower-frequency DBS input results in a more subtle declustering of the GPi cells, relative to the unstimulated Parkinsonian state. Note that the cell pairs synchronized in the Parkinsonian condition (cells $\{1,5\},\{2,6\},\{3,7\},\{4,8\}$) remain essentially synchronized following the application of the DBS current, but these pairs become largely desynchronized relative to one another. This restores the transmission reliability *Rel* to a value of 0.95, as the summed inputs to the TC cells are more smoothly distributed over time. Furthermore, the average DBS current $\langle I_{DBS}\rangle$ (i.e., the time average of the I_{DBS} shown in Fig. 5(a)) that achieves this desynchronization and attendant reliability is $\langle I_{DBS}\rangle = 0.6pA/(\mu m^2 ms)$, substantially lower than the value $\langle I_{DBS}\rangle = 20pA/(\mu m^2 ms)$ of the high-frequency current in Fig. 3 (e-f).

The mechanism for this network desynchronization is not immediately evident. However, an expansion (Fig. 6 (a)) of the temporal scale of the voltage profiles for two STN cells in the Parkinsonian state indicates that, although the two cells demonstrate very similar bursting patterns, there is still a small time lag (phase shift) between their spikes. The standard high-frequency DBS technique employs strong, high frequency current inputs, which force the STN neurons to be phase-locked with the high-frequency DBS current (Fig. 6 (b) and (c)). In contrast, the low frequency DBS impulses in Fig. 5 (a) do not always trigger voltage spikes at their onset. Instead, such DBS input regulates the subthreshold dynamics of distinct neurons in a slightly different way due to the time lag, and this difference propagates over time in a complex, nonlinear fashion, resulting in the desynchrony among cell pairs seen in (Fig. 6 (d) and (e)). Again, we note that the cells within a given pair (the pairs being $\{1,5\},\{2,6\},\{3,7\},\{4,8\}$, as above) remain synchronized, as the STN cells within each pair are entrained to identical synaptic inputs (being connected to the same other cells in the network) and identical I_{DBS} currents. Overall, we emphasize that the present modelling result should not be interpreted as a suggestion that these same low-frequency DBS currents will be effective in clinical applications, but as illustrating the general possibility of DBS inputs that produce desynchronized dynamics differing substantially from entrained network responses to high-frequency DBS.

In the laboratory setting, the currents induced in different neurons by DBS depend not just on stimulator parameters but also on factors including properties of the extracellular medium and the orientation of target cells (e.g. (McIntyre and Grill, 2002; McIntyre et al., 2004)). To test the robustness of DBS effects against such factors, we first smooth the DBS pulse train by applying a Gaussian filter with standard deviation of 2ms to the square waveform considered above. We then include heterogeneity in the DBS input by differentially and randomly scaling the current impulse i_D received by each STN neuron. Fig. 7 (a) to (c) shows that these factors do not alter the entraining effect of high-frequency DBS currents. Similarly, the network's response to the low-frequency DBS input of Fig. 5 (a) is robust to the smoothing and heterogeneity: the input disrupts the regular alternating, clustered Gpi firing patterns of the DBS state (Fig. 7 (d) to (f)), yielding $Rel = 0.94$, although cells do drift in and out of different types of overlapping-in-time clusters on longer timescales. Note that this heterogeneity is insufficient to substantially desynchronize neurons within cell pairs (i.e. cells $\{1, 5\}$, etc. are still synchronized, due to common synaptic inputs); further simulations (not shown) with additional heterogeneity at a level of $50 - 80$ percent do begin to desynchronize the cells *within* these pairs, a result that may be roughly explained by considering the relative orders of magnitude of DBS and synaptic currents.

Optimization of periodic DBS currents

The observation that DBS inputs can interact with intrinsic network dynamics to desynchronize and make more irregular firing in the simulated Parkinsonian state suggests that PD symptoms might be treated by introducing carefully tuned DBS currents that evoke firing patterns resembling the normal state (hence restoring, for example, normal functioning of thalamocortical relay cells). This is in contrast to the present model of the standard high-frequency DBS procedure, which also restores normal functioning of thalamocortical relay cells, but via entraining the subthalamopallidal network to rapid and regular firing that differs substantially from its normal-state activity.

Despite the fact that a fairly coarse grid sampling of DBS frequency and amplitude identified

novel, lower-frequency parameter settings that give satisfactory TC cell transmission reliability *Rel* for the computational model of Terman et al. used here, it is not known whether such a sampling strategy, or even a local optimization algorithm such as gradient descent, would be effective in identifying novel patterns of DBS inputs in experimental or clinical applications. This is because the efficacy of such strategies would depend upon a relatively simple relationship between levels of clinical PD symptoms and the full range of possible DBS parameter settings. To demonstrate how such a problem might eventually be overcome, and to identify further DBS patterns that are effective within the present computational model, we employ the Genetic Algorithm (GA) as a closed-loop learning technique to identify optimal DBS parameter settings that best achieve simulated physiological outcomes. The GA's global search capability is especially important in this application, as Fig. 4 indicates that multiple local minima/maxima can exist even in a simple two-parameter case, so that local search algorithms can be easily trapped without finding the best solutions. Another overriding advantage of the GA in a clinical setting is its model-free operation. The GA operation requires only knowledge of the trial DBS parameter settings and measurement of the corresponding neurophysiological properties, thus it can be directly integrated with the instrument capabilities in clinical applications. Issues related to GA's practical applications are addressed elsewhere (Feng et al., in preparation).

We first illustrate the GA based closed-loop search over the three standard DBS parameters of pulse period ρ_D, amplitude i_D, and duration σ_D to identify values that maximize TC reliability *Rel* (note that all three of the standard DBS parameters are covaried, in contrast to the two-parameter grid search of Fig. 4). A current cost is subtracted to reduce the energy usage of the DBS input. Consequently, the function which the GA seeks to maximize is

$$J = Rel - w \int_{t=0}^{T} I_{DBS} dt, \tag{8}$$

where $\int_{t=0}^{T} I_{DBS} dt$ is the total DBS current over the duration of the DBS stimulation ($T = 6500ms$) and the weight parameter $w \approx 0.001$. The GA quickly finds several DBS parameter settings with high J values. Fig. 5 (d) shows one solution that achieves both reliable TC cell signal transmission and substantial current reduction. Specifically, while the reliability

Rel is slightly worse in comparison with the DBS current of Fig. 5 (a) identified from the two-parameter grid search above ($Rel = 0.95$ vs $Rel = 0.90$), the averaged current $\langle I_{DBS} \rangle$ in Fig. 5 (c) is about three times lower than that in Fig. 5 (a) and 100 times lower than that for high-frequency DBS in Fig. 3 (e). This suggests that the two terms in J can have significant competing effects. It is also notable that in both cases with low-frequency DBS, the cells remain in the four pairs discussed above, with spiking of each pair being relatively irregular in time and the different pairs being desynchronized. This pairing is a common characteristic of DBS solutions found by the GA.

Statistical measure of Gpi spiking patterns

In addition to the desynchronizing DBS currents such as that shown in Fig. 5 (c), the GA utilized in optimizing the periodic DBS also identifies multiple currents that give reliable transmission (i.e., $Rel \approx 1$), but where the GPi cells are mostly synchronized (data not shown), similar to the effects of high frequency currents discussed above. This shows that, when *Rel* is used as the cost function, the GA is not selective for tonic and regular, more synchronized vs. irregular, less synchronized Gpi firing patterns. To allow for such selectivity, we introduce a new measure of the network dynamics based directly on the statistical properties of these firing patterns. The TC cells are therefore not included in the simulations upon which these implementations of the GA are based.

Fig. 8 (a) shows the synaptic (conductance) output from GPi cells one and two in the normal state (see (Rubin and Terman, 2004) for details). To calculate the new measure, which we call *Cor* (see Methods), the synaptic pulse time t_i^a (defined as when the i-th pulse in conductance of the a-th GPi cell's synapse reaches its maximum) is numerically determined. Then, we compute the pulse delay $d_{ij}^a = t_j^a - t_i^a (i < j)$ between all pairs of these pulse times for the same GPi cell, and quantify the number k_m^a of d_{ij}^a values located in each of a set of discrete "interspike interval bins" $[t_m, t_m + \Delta]$, where t_m represents the m-th ($m = 1, 2, \ldots, M$) bin with width $\Delta = 10ms$, and the full set of bins cover the period over which the network is simulated (following an initial transient, as described above). A plot of k_m^a counts against the corresponding time t_m

gives a discrete measure of the autocorrelation for the a-th GPi cell's spike train (Fig. 8(b)). A Gaussian filter of width 10ms is then applied to smooth this autocorrelation function, resulting in $k_m^{*,a}$ for each bin. The same procedure is then executed to calculate the crosscorrelation $k_m^{*,a,b}$ of the synaptic spike time delay between two different GPi cells a and b (Fig. 8(c)). Cor is then represented as

$$Cor = \alpha_1 \frac{1}{N} \sum_{a=1}^{N} [\frac{\sum_{m=1}^{M} (k_m^{*,a} - \bar{k}_m^{*,a})^2}{M}]^{1/2} + \alpha_2 \frac{2}{N(N-1)} \sum_{1 \le a < b \le N} [\frac{\sum_{m=1}^{M} (k_m^{*,a,b} - \bar{k}_m^{a,b})^2}{M}]^{1/2} + \alpha_3 \bar{k}_0^{a,b}, \tag{9}$$

where $\bar{k}_m^{*,a}$ and $\bar{k}_m^{*,a,b}$ are the mean values of $k_m^{*,a}$ and $k_m^{*,a,b}$, respectively. The first and second terms in Eq. (9) represent the standard deviation of the bin counts for the auto- and cross-correlations, repectively, normalized over all GPi cells. That is, the first term measures the extent to which firing of the Gpi cells deviates from a uniform (Poisson) process, while the second measures the extent to which pairs of Gpi cells are synchronized, or display preferred phase shifts in their firing times. A third term $\bar{k}_0^{a,b}$ is added as a separate measure of purely synchronized Gpi firing. This third term serves to make the GPi firing statistics more similar to those of the normal state, as normal cells have fewer counts in this first bin (Fig. 8(c)). α_1, α_2, and α_3 are the positive weighing paramters for the three terms, respectively, and their values in the simulations are $\alpha_1 = 2.0$, $\alpha_2 = 1.0$, and $\alpha_3 = 2.0 \times 10^{-5}$.

Fig. 8 reveals the expected substantial differences in these statistical measures of GPi firing patterns between normal and Parkinsonian states. In the normal state, the auto- and cross-correlation values are roughly similar for most spike delays, indicating that spike times are approximately uniformly distributed both within and across GPi cells (although a moderately preferred interspike interval is apparent). For the normal state, $\bar{k}_0^{a,b}$ is also low in magnitude, showing an anticorrelation in synaptic pulse times among GPi cells. In contrast, in the Parkinsonian state both the auto- and cross-correlation curves show strong periodic oscillations due to the synchronized bursting dynamics of GPi cells; the relatively high value of $\bar{k}_0^{a,b}$ also follows from the synchrony. Evidently, the normal state has lower values for all three terms in Eq. (9) than the Parkinsonian state, so that Cor is a good indicator to distinguish between the two states; below, the GA will be used to find DBS currents that minimize Cor. Other measures

of network dynamics, including those based more directly on the power spectrum of the GPi spike pattern, have been tested. We find that the correlation measure Cor is better suited to extract statistical properties of Gpi synaptic output that distinguish normal and Parkinsonian firing, while filtering out the less important details related to spike shape and amplitude, and, additionally, that Cor produces a stable characterization of network dynamics with a limited amount of data in the form of conductance spikes.

Fig. 9 illustrates the dependence of the measure Cor on changes in the standard DBS parameters of pulse period ρ_D and amplitude i_D, with (as Fig. 9) fixed pulse duration σ_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. Other settings of the parameters generally lead to lower Cor values, and in general the 'landscape' of Cor as a function of parameters ρ_D and i_D alone is highly complex, suggesting that, especially when additional parameters such as pulse duration σ_D are introduced, it will be difficult to identify optimal parameter settings (giving the lowest possible values of Cor) without the guidance of a global search algorithm.

Based on the observation mentioned above that the normal state has lower Cor value than the Parkinsonian state, we next apply the GA to search for periodic DBS parameter settings (period, amplitude, and duration) that minimize $J = Cor$ in the Parkinsonian network. Note that, in contrast to the case of maximizing reliability Rel, a current cost is no longer necessary in minimizing Cor, because the GA search *a priori* avoids the high frequency region, which is usually associated with high Cor values (Fig. 9). Fig. 10 (left panel) shows an optimal solution found by the GA. Both the auto- and cross-correlation curves are roughly similar to the normal state, although Gpi firing is more burstlike. Similar to Fig. 5, the eight GPi neurons fire in four pairs that are desynchronized with respect to one another.

Optimization of stochastic DBS currents

We next explore stochastic DBS currents to study whether additional effective inputs can be found by enlarging the space of possible currents. The primary motivation resides in the argu-

ment that random DBS inputs might be better suited than periodic ones for inducing irregular firing patterns. Beyond insights for this model, it is intuitively reasonable that stochastically defined inputs could provide a useful alternative in cases where patients are unresponsive to standard, periodic DBS currents or the effects of periodic DBS decay as the disease progresses (cf. (Lyons et al., 2001; Hariz et al., 1999)), although this question clearly must be adjudicated experimentally.

The GA is employed to identify random DBS currents that optimally return Gpi firing patterns to their normal state, as measured by minimal values of the measure Cor. In general, such a random current can be represented in many different ways. Here the procedure consists of generating each DBS pulse train from a probability density function (PDF) which determines the probability of interpulse delays (see Methods). To enable an efficient GA search, we use a 10-parameter, piecewise-constant PDF. The GA then optimizes 12 parameters: the 10 constants defining the PDF, the pulse amplitude i_D, and the pulse width σ_D. Comparing with the three-parameter case of purely periodic square wave DBS currents, the advantages the GA are clear in the present case because the high dimensionality makes grid search methods unfeasible (cf. (Rizzone et al., 2001)).

GA optimization of stochastic DBS currents is performed both with and without the addition of a current cost, as considered above. Due to the randomness in generating the DBS pulse train from realization-to-realization and the fact that J is measured from finite-length (6500 ms) simulations, the minimum J values fluctuate throughout the GA optimization (see Fig. 10 (e)). However, the GA is still able to quickly converge to DBS parameter settings that desynchronize the GPi cells (both with and without the current cost), confirming that stochastically defined DBS inputs can be effective in this manner if their parameters are properly chosen. Nevertheless, stochastic DBS inputs still produce Gpi firing that is more burstlike than in the normal state, as clusters of transiently (nearly) synchronized cell pairs continuously and irregularly form and dissipate over timescales of 100ms. Fig. 10 (right panel) shows the case including the current cost (i.e., $J = Cor + w \int_{t=0}^{T} I_{DBS} dt$), which also minimizes the DBS current amplitude (data for similar results without current cost not shown). Similar to the left panel of Fig. 10,

clustering and bursting lead to periodic oscillations evident in the correlograms, but which are nevertheless more moderate than those observed in the unstimulated PD state (Fig. 8, right panel). Additionally, the optimal DBS currents obtained via GA optimization of the *Cor*-based cost function also produce highly reliable signal relay by the TC cells ($Rel > 0.9$), further confirming the utility of this measure for seeking optimal DBS currents. (We note that work in progress (D. Terman, et al., unpublished observations) studies the relation between Gpi firing patterns and TC reliability in greater detail.)

DBS currents with alternative waveforms

As emphasized in (Tass, 2001), it is possible that currents deviating from the standard 'pulsatile' form traditionally used in DBS may be more effective in realizing clinical objectives. Here we choose a discrete, periodic (Haar) basis to represent general DBS waveforms (see Methods), and we explore computationally whether the GA can identify such general waveforms that achieve the objectives quantified above.

Fig. 11(left, right) shows optimal solutions identified using the *Rel* (with current cost $w = 10^{-5}$) and *Cor* cost functions, respectively. The left panel of Fig. 11 demonstrates a substantial improvement in TC cell signal transmission reliability compared with the unstimulated PD state ($Rel = 0.98$ vs. $Rel = 0.43$). This follows even though Gpi cells are mostly synchronized under this DBS input, because their firing pattern is more tonic than burstlike. The right panel, for the *Cor* cost, shows that the synchronized, periodic bursts of the PD state have been reshaped to be more irregular in duration, although they are still largely coordinated across cells. These results indicate that DBS waveforms other than sparse pulses can modulate the otherwise bursting and clustered firing patterns of the network in diverse ways depending on the optimization objective specified. Nevertheless, the net current used in the present DBS patterns is much greater than for the standard pulsatile inputs ($\langle I_{DBS}\rangle = 24pA/(\mu m^2 ms)$ and $\langle I_{DBS}\rangle = 90pA/(\mu m^2 ms)$ for the *Rel* and *Cor* cost, respectively), indicating that sustained input currents may not be as efficient in achieving a given type of network state.

Optimal DBS in heterogeneous networks

Here, we add heterogeneity to the intrinsic dynamics of the STN-Gpe network, and explore how this affects the form and quality of the optimal DBS solutions found by the GA. Specifically, we first introduced a 1% multiplicative heterogeneity (that is, independent, random prefactors with mean 1 and standard deviation 0.01) to each intrinsic current in each STN and GPe cell, as well as to all synaptic currents among STN and GPe neurons (i.e., $I_{GPe\to GPe}$, $I_{GPe\to STN}$, and $I_{STN\to GPe}$). We also included prefactors for the DBS currents to STN cells, with 10% random heterogeneity in cell-to-cell values, as in Fig. 7. Additionally, the values of the synaptic and intrinsic currents $I_{GPe\to GPe}$ and I_{STN} were adjusted to ensure that (*i*) the intrinsic dynamics of the network is Parkinsonian in character (i.e., the cells burst in synchronized clusters, as in Fig. 3(c)) and (*ii*) high frequency DBS (as in Fig. 3(e)) restores the faithful transmission of I_{SM} signals.

The GA was then used, as above, to search for periodic, stochastic, and alternative (defined via the Haar basis) DBS currents that optimize either the *Rel* or the *Cor* cost function. Fig. 12 shows one DBS optimal input of each type, for the case where *Cor* is optimized (as above, with no extra term penalizing total current); data for *Rel* is not shown here. For all three types of DBS, optimized inputs modified the the highly rhythmic auto- and cross-correlation statistics of the unstimulated Parkinsonian state, producing lower values of *Cor*. For the periodic and alternative DBS inputs, firing times are largely synchronized across Gpi cells (evidenced by the similarity between auto- and cross-correlograms in panels A and C), although they are not directly entrained to the inputs. In comparing with analogous optimal DBS currents found without network heterogeneity (Fig. 10 and Fig. 11), the most substantial enhancement of DBS effects due to the presence of heterogeneity is for the case of stochastic DBS (panel B of Fig. 12); auto- and cross-correlograms in this case also most closely resemble those of the normal state. This raises the possibility that a stochastic class of input may also be the best equipped to exploit the far more complex forms of neuronal heterogeneity that will be encountered in experimental or clinical settings, and in any case illustrates that the efficacy of the GA in identifying novel DBS patterns without relying on a purely homogeneous network structure.

Discussion

Summary

In the present paper, we test a novel strategy for identifying optimal DBS waveforms for the treatment of Parkinson's Disease, using a computational network model for the basal ganglia previously developed by Terman and collaborators (Terman et al., 2002; Rubin and Terman, 2004). We employ two statistical measures that gauge the effect of the DBS inputs on the network: *Rel* (following (Rubin and Terman, 2004)), which assesses the impact of Gpi firing on the reliability of simulated thalamocortical transmission, and Cor, based on the auto- and crosscorrelation of simulated Gpi cell firing times. A Genetic Algorithm (GA) global optimization method searches the multidimensional parameter space describing DBS inputs to seek parameter settings that desynchronize firing times between clusters of Gpi cells and cause firing within these clusters to become more irregular, while also minimizing total DBS current flux $\langle I_{DBS} \rangle$. These optimal DBS inputs differ substantially from standard DBS inputs in both their waveform (e.g., being stochastic or of lower frequency) and their impact on the network (often desynchronizing rather than entraining cells).

Three types of DBS waveforms are tested: three-parameter periodic square pulses with constant spacing (of the square wave form used clinically), 12-parameter stochastic inputs generated from a piecewise constant probability distribution function, and 16-parameter nonpulsatile waveforms generated from a discrete, periodic Haar basis. The GA converges to optimal parameter settings for all three types of waveforms, which achieve satisfactory desynchronization and regularization of Gpi firing times. This corresponds to minimized values of Cor. Moreover, the optimized square wave and stochastic waveforms are sparse in time, and therefore require average I_{DBS} amplitudes substantially lower than those used in models of high-frequency DBS. In all three cases, we observe that the DBS inputs that minimize Cor also maximize the reliability measure *Rel*. The converse is not true: DBS waveforms yielding high *Rel* values do not necessarily desynchronize the Gpi cells, and hence do not necessarily give low values of Cor. Finally, when heterogeneity was introduced into both network and DBS parameters, we

found that the stochastic waveforms were most effective in minimizing Cor. This illustrates a potential advantage of this novel class of DBS waveform, which bears further exploration in more sophisticated models and in experimental settings.

Relationship to previous computational studies

Previous studies (Tass, 1999; Tass, 2001; Rosenblum and Pikovsky, 2004; Popovych, Hauptmann, and Tass, 2005; Hauptmann, Popovych, and Tass, 2005) have also used computational models to explore alternative DBS patterns, with similar overall objectives of identifying effective currents with lower total amplitude. In these studies, the basal ganglia is approximated by a large ensemble of simplified oscillators that is often (but not always (Hauptmann, Popovych, and Tass, 2005)) *globally* coupled so as to produce synchronized dynamics representing the Parkinsonian state. Techniques of "demand controlled" stochastic phase-resetting (Tass, 1999; Tass, 2001; Tass, 2003), and linear (Rosenblum and Pikovsky, 2004) or nonlinear (Popovych, Hauptmann, and Tass, 2005; Hauptmann, Popovych, and Tass, 2005) delayed feedback of averaged oscillator states are used to identify DBS currents that desynchronize the model oscillators and differ in form from traditional high-frequency inputs. The models indicate that these alternative DBS currents may indeed be comparably effective, or even more so, in achieving desynchrony with respect to their high-frequency counterparts, while requiring a substantially lower current integral. Our study, therefore, adds to this previous computational evidence that alternatives to high-frequency DBS might exist in the Parkinsonian brain itself. However, differences in the underlying computational models among these previous studies and the present one preclude a direct comparison of the various types of currents they suggest.

Limitations and interpretation of the computational model

Our results are found using the computational subthalamopallidal network model developed in (Terman et al., 2002) and applied to studies of standard high-frequency DBS waveforms in (Rubin and Terman, 2004). While it is physiologically detailed and carefully fit to experiments, like any model it has limitations in describing the Parkinsonian basal ganglia *in vivo*. First,

the model was originally parameterized to match activity in slice experiments (Terman et al., 2002); while parameter adjustments to bring model dynamics closer to the *in vivo* range were explicitly made and described in (Rubin and Terman, 2004), some differences may be expected to remain. Moreover, the simulated nuclei represent only the indirect pathway from striatum to Gpi: backprojections from Gpe to striatum as well as the general closed loop structure of the cortico-striatal-thalamic network are neglectged. Within the indirect pathway, the model necessarily includes only a representative subnetwork of cells, with equal numbers belonging to each nucleus. Finally, following, e.g., (Rubin and Terman, 2004; Popovych, Hauptmann, and Tass, 2005), we study only the effects of DBS waveforms applied directly as (homogeneous or heterogeneous) intracellular currents. Thus, we neglect the important stage by which voltage transients at the DBS electrode are transduced to transmembrane currents, a process explicitly modeled in, e.g., (Hahn et al., 2005; McIntyre and Grill, 2002).

Nevertheless, the computational model of (Terman et al., 2002; Rubin and Terman, 2004) exhibits key features that make it an ideal for our study of the effects of novel DBS waveforms, and how closed-loop learning algorithms might identify and optimize such waveforms. The conductance-based equations for the individual neurons contain membrane currents which were matched in (Terman et al., 2002) to the distinct physiology of the different basal ganglia nuclei, producing, e.g., post-inhibitory rebound bursts. Moreover, these model neurons, when coupled (via experimentally-motivated synaptic dynamics) in the present sparse, regular network architecture, have been shown to display *both* archetypal normal-state and Parkinsonian modes of activity, according to the level of two network parameters whose values are known to change with neurodegeneration in PD (Terman et al., 2002). Thus, we believe that the computational model has the critical components of cellular physiology, network connectivity, and normal and Parkinsonian dynamics to serve as a first testing ground for the impact of *qualitatively* novel classes of DBS waveforms and algorithms.

This said, we do not expect that the DBS patterns found to be effective or optimal for the computational model of (Terman et al., 2002; Rubin and Terman, 2004) to be likewise optimal for more detailed models or in experimental or clinical settings. Indeed, given the model's

structure as highlighted above, such a quantitative match is extremely unlikely. However, the point of this paper is not to suggest specific DBS waveforms that could be applied, for example, clinically, but rather to provide computational evidence that (1) different types of DBS waveforms exist that could be equally or more effective than the standard high-frequency patterns, and (2) these novel waveforms can be identified by standard, model-independent learning algorithms.

Conclusion

The objective of this study is to test, using a biophysically-based model of the basal ganglia, whether closed-loop nonlinear learning optimization techniques can identify DBS waveforms that differ from standard high-frequency patterns and nonetheless modulate network dynamics according to various criteria. The learning algorithm we use requires no knowledge about the dynamics of the underlying system, so that it can in principle be used with models or even experimental or clinical systems of arbitrary biological scale, complexity, and detail. We hope that the positive findings here will inspire experimental tests of the present optimization strategy as a first step toward possible clinical implementation feasibility of an alternative strategy for alleviating the motor symptoms associated with PD and related neurodegenerative diseases. Further issues relevant to experimental or eventual clinical applications are discussed in an upcoming paper (Feng et al., in preparation).

Acknowledgements

We thank Dr. Jonathan Rubin for helpful discussions during the project and insightful comments on an earlier version of this manuscript. We also thank Dr. David Terman for his insights, as well as for providing the simulation code from the paper (Terman et al., 2002). We are indebted to Drs. Karen Sigvardt and Vikki Wheelock for their important comments, advice, and references throughout this project. E.B. was supported by a NSF Mathematical Sciences Postdoctoral Research Fellowship and holds a Career Award at the Scientific Interface

from the Burroughs Wellcome Fund. H.R. and X.F. acknowledge support from the National Science Foundation.

References

Benabid, A. (2003). Deep brain stimulation for Parkinson's disease. *Curr. Opin. Neurobiol.* 13: 696–706.

Benabid, A., A. Koudsie, A. Benazzouz, B. Piallat, P. Krack, P. Limousin-Dowsey, J. Lebas, and P. Pollak (2001). *Advances in neurology, Vol 86: Parkinson's disease*, chapter Deep brain stimulation for Parkinson's disease. Lippincott Williams & Wilkins, Philadelphia.

Benazzouz, A., D. Gao, Z. Ni, B. Piallat, R. Bouali-Benazzouz, and A. Benabid (2000). 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.

Bergman, H., A. Feingold, A. Nini, A. Raz, H. Slovin, M. Abeles, and E. Vaadia (1998). Physiological aspects of information processing in the basal ganglia of normal and parkinsonian primates. *Trends. Neurosci.* 21: 32–38.

Beurrier, C., B. Bioulac, J. Audin, and C. Hammond (2001). High-frequency stimulation produces a transient blockade of voltage-gated currents in subthalamic neurons. *J. Neurophysiol.* 85: 1351–1356.

Boraud, T., E. Bezard, B. Bioulac, and C. Gross (1996). 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.

Deep Brain Stimulation for Parkinson's disease study group (2001). Deep-brain stimulation of the subthalamic nucleus or the pars interna of the globus pallidus in Parkinson's disease. *New England J. Med.* 345: 956–963.

Feng, X., B. Greenwald, H. Rabitz, E. Shea-Brown, and R. Kosut (in preparation). Closed-loop optimal deep brain stimulation for treating parkinson's disease .

Goldberg, D. (1989). *Genetic algorithms in search, optimization and machine learning.* Addison-Wesley, Boston, MA.

Hahn, P., D. Lee, G. Russo, J. Vitek, and C. McIntyre (2005). Stimulation on a model of subthalamopallidal network activity. *Soc. Neurosci. Abstr.* p. 331.6.

Hariz, M., P. Shamsgovara, F. Johansson, G. Hariz, and H. Fodstad (1999). Tolerance and tremor rebound following long-term chronic thalamic stimulation for parkinsonian and essential tremor. *Stereotact. Funct. Neurosurg.* 72: 208–18.

Hashimoto, T., E. Elder, M. Okun, S. Patrick, and J. Vitek (2003). Stimulation of the subthalaic nucleus changes the firing pattern of pallidal neurons. *J. Neurosci.* 23: 1916–1923.

Hauptmann, C., O.V. Popovych, and P.A. Tass (2005). Delayed feedback control of synchronization in locally coupled neuronal networks. *Neurocomputing* 65: 759–767.

Kleiner-Fisman, G., D.N. Fisman, E. Sime, J.A. Saint-Cyr, A.M. Lozano, , and A.E. Lang (2003). Long-term of bilateral deep brain stimulation of the subthalamic nucleus in patients with advanced Parkinson's disease. *J. Neurosurg.* 99: 489–495.

Krack, P., A. Batir, N. Blercom, S. Chabardes, V. Fraix, and et al. C. Ardouin (2003). Five-year follow-up of bilateral stimulation of the subthalamic nucleus in advanced Parkinson's disease. *N. Engl. J. Med.* 349: 1925–1934.

Lyons, K., W. Koller, S. Wilkinson, and R. Pahwa (2001). Long term safety and efficacy of unilateral deep brain stimulation of the thalamus for parkinsonian tremor. *J Neurol Neurosurg Psychiatry.* 71: 682–4.

Magnin, M., A. Morel, and D. Jeanmonod (2000). Single-unit analysis of the pallidum, thalamus, and subthalamic nucleus in parkinsonian patients. *Neuroscience* 96: 549564.

Maurice, N., A. Thierry, J. Glowinski, and J. Deniau (2003). Spontaneous and evoked activity of substantia nigra pars reticulata neurons during highfrequency stimulation of the subthalamic nucleus. *J. Neurosci.* 23: 9929–9936.

McIntyre, C., W. Grill, D. Sherman, and N.V. Thakor (2004). Celluar effectsof deep brain stimulation: model-based study of activation and inhibition. *J. Neurophysiol.* 91: 1457–1469.

McIntyre, C.C. and W.M. Grill (2002). Extracellular stimulation of central neurons: influence of stimulus waveform and frequency on neuronal output. *J. Neurophysiol.* 88: 1592–1604.

Montgomery, E. and K. Baker (2000). Mechanisms of deep brain stimulation and future technical developments. *Neurol. Res.* 22: 259–266.

Nini, A., A. Feingold, H. Slovin, and H. Bergman (1995). Neurons in the globus pallidus do not show correlated activity in the normal monkey, but phase-locked oscillations appear in the mptp model of parkinsonism. *J. Neurophysiol.* 74: 18001805.

Okun, M., M. Tagliati, M. Pourfar, H. Fernandez, R. Rodriguez, R. Alterman, and K. Foote (2005). Management of referred deep brain stimulation failures: a retrospective analysis from 2 movement disorders centers. *Arch Neurol.* 62: 1250–5.

Olanow, W., M. Brin, and J. Obeso (2000). The role of deep brain stimulation as a surgical treatment for Parkinson's disease. *Neurology* 55 (Supp.6): S60–S66.

Popovych, O.V., C. Hauptmann, and P.A. Tass (2005). Effective desynchronization by nonlinear delayed feedback. *Phys. Rev. Lett.* 94: 164102–1–4.

Rizzone, M., M. Lanotte, B. Bergamasco, A. Tavella, E. Torre, G. Faccani, A. Melcarne, and L. Lopiano (2001). Deep brain stimulation of the subthalamic nucleus in Parkinson's disease: effects of variation in stimulation parameters. *J. Neurol. Neurosurg. Psychiatry* 71: 215–219.

Rodriguez-Oroz, M.C., J.A. Obeso, A.E. Lang, J.L. Houeto, and et al. P. Pollak (2005). Bilateral deep brain stimulation in Parkinson's disease: a multicentre study with 4 years follow-up. *Brain* 128: 2240–2249.

Rodriguez-Oroz, M.C., I. Zamarbide, J. Guridi, M.R. Palmero, and J.A. Obeso (2004). 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.

Rosenblum, M. and A. Pikovsky (2004). Delayed feedback control of collective synchrony: an approach to suppression of pathological brain rhythms. *Phys. Rev. E* 70: 041904–1–11.

Rubin, J. and D. Terman (2004). High frequency stimulation of the subthalamic nucleus eliminates pathological thalamic rhythmicity in a computational model. *J. Comput. Neurosci.* 16: 211–235.

Tass, P.A. (1999). *Phase resetting in medicine and biology. Stochastic modeling and data analysis*. Springer, Berlin.

Tass, P.A. (2001). Desynchronizing double-pulse phase resetting and application to deep brain stimulation. *Biol. Cybern.* 85: 343–354.

Tass, P.A. (2003). A model of desynchronizing deep brain stimulation with a demand-controlled coordinated reset of neural subpopulations. *Biol. Cybernetics* 89: 81–88.

Terman, D., J. Rubin, A. Yew, and C.J. Wilson (2002). Activity patterns in a model for the subthalamopallidal network of the basal ganglia. *J. Neurosci.* 22: 2963–2976.

Windels, F., N. Bruet, A. Poupard, N. Urbain, G. Chouvet, C. Feuerstein, and M. Savasta (2000). 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.

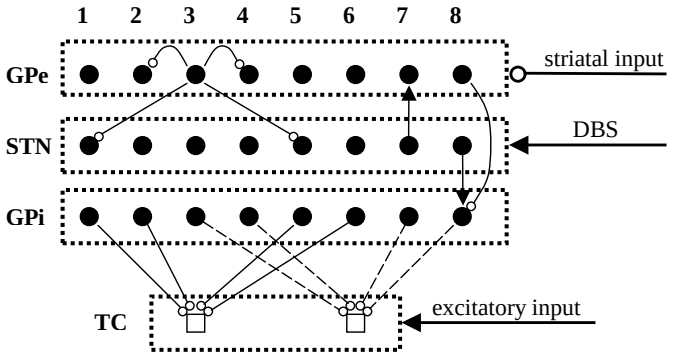

Figure 1: The "structured, sparsely-connected" network architecture adapted from (Terman et al., 2002; Rubin and Terman, 2004). Lines ending with arrows and open circles indicate excitatory and inhibitory synaptic connections, respectively. Each GPe neuron inhibits two immediate GPe neighbors as well as two STN neurons, skipping the three STN cells located nearest to it in the 'arrays' of cells. Each STN cell sends excitation to the nearest GPe cell in the array. In addition, GPe neurons receive simulated striatal inhibition, and excitatory DBS inputs are applied to STN neurons in some cases. Each GPi cell receives inhibition from the nearest GPe cell and excitation from the nearest STN cell. Each of two TC cells receives inhibitory input from the four GPi neurons shown, and the TC cells also (uniformly) receive model excitatory sensorimotor input. As in (Terman et al., 2002; Rubin and Terman, 2004), our network architecture has a periodic structure, so that cells 1 and 8 in each array are neighbors, etc.

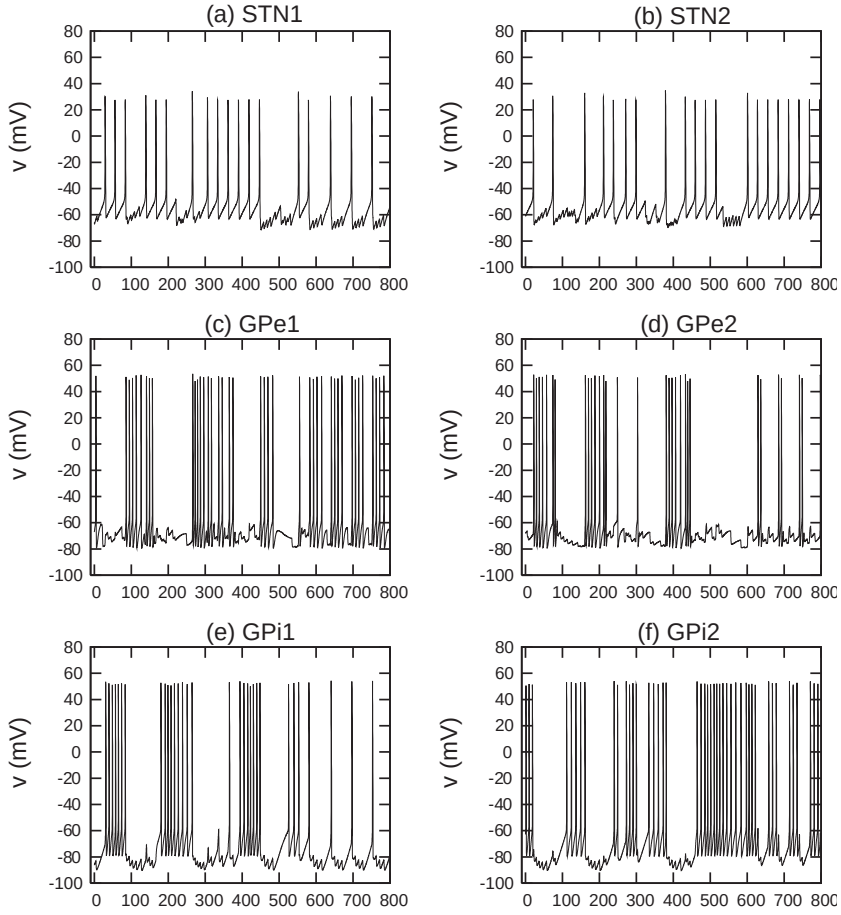

Figure 2: Voltage profiles for STN, GPe, and GPi neurons in the normal state. Two cells of each type are plotted. The firing patterns of each type of cells are irregular and uncorrelated.

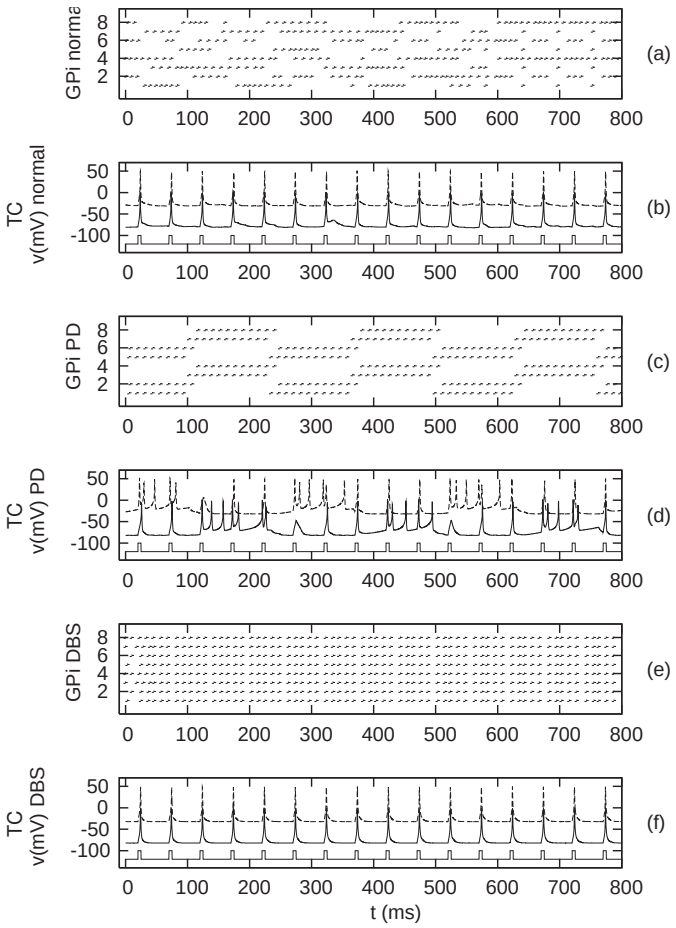

Figure 3: (a) Raster plot of the spike times for all eight GPi cells in the normal state. (b) Voltage traces of the two TC neurons in the normal state with I_{SM} displayed beneath the voltage traces; solid line is cell 1 of Fig. 1, dashed line is cell 2, shifted up by 50mV for clarity. Note that I_{SM} input is faithfully relayed. (c) Raster plot of GPi cells in the Parkinsonian state, showing clustered, synchronous bursting. (d) Voltage of TC cells in the Parkinsonian state, where the TC cells fail to relay the I_{SM} input faithfully. (e) Raster plot of the Parkinsonian GPi cells with high frequency DBS current applied ($i_D = 200pA/\mu m^2$, $\rho_D = 6ms$, $\sigma_D = 0.6ms$, and average $\langle I_{DBS} \rangle = 20pA/(\mu m^2 ms)$). The Gpi neurons are entrained to the DBS input, firing at half the DBS frequency (STN cells fire at the DBS frequency, see Fig 6(c)). (f) Voltage of TC cells shows that high frequency DBS has restored the faithful transmission of I_{SM} signals, as in (Rubin and Terman, 2004).

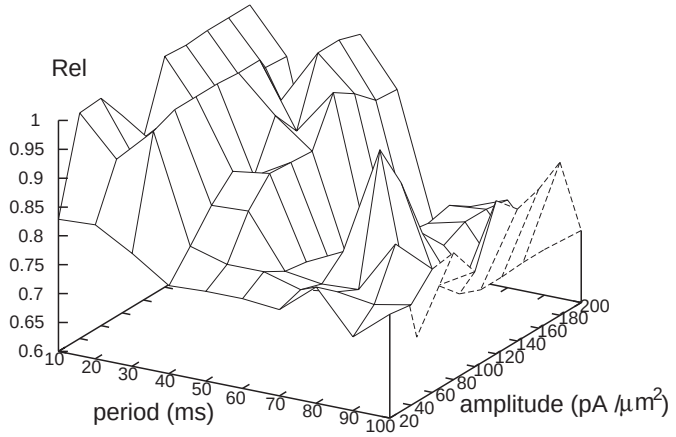

Figure 4: The dependence of TC transmission reliability Rel on DBS pulse period ρ_D and amplitude i_D. The impulse duration σ_D is fixed at 0.6ms. 100% reliability corresponds to $Rel = 1$.

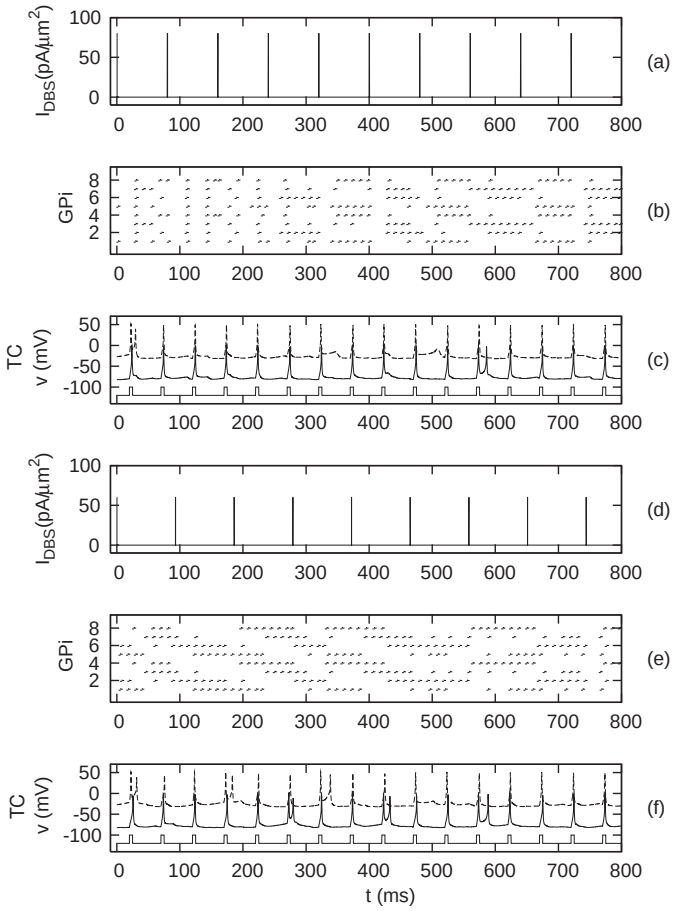

Figure 5: (a) A low frequency, periodic DBS current with parameters $\rho_D = 80ms$, $i_D = 80pA/\mu m^2$, $\sigma_D = 0.6ms$, giving time-average $\langle I_{DBS} \rangle = 0.6pA/(\mu m^2 ms)$. (b) Raster plot for the Parkinsonian GPi cells during application of the DBS current in (a), which largely desynchronizes them. (c) Faithful TC cell signal transmission is restored by the current in (a). (d) A GA-optimized periodic DBS current ($\rho_D = 93ms$, $i_D = 60pA/\mu m^2$, $\sigma_D = 0.3ms$, and $< I_{DBS} >= 0.19pA/(\mu m^2 ms)$). (e) Raster plot of the Parkinsonian GPi cells with the DBS current in (d). (f) TC voltage profile with the DBS current in (d).

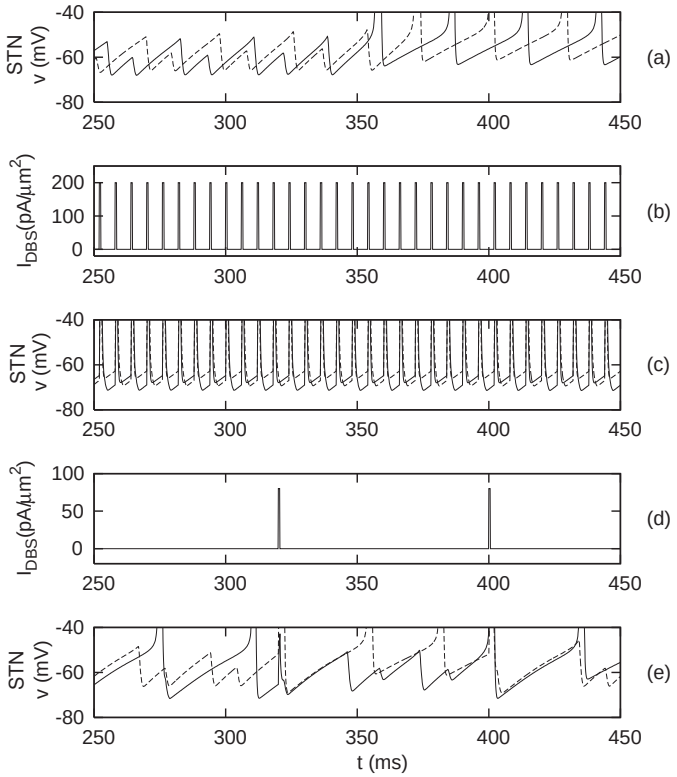

Figure 6: (a) A zoomed plot of the voltage traces for STN cell 1 (solid line) and 2 (dashed line) in the Parkinsonian state, showing subthreshold dynamics. Note that both the spike times and subthreshold fluctuations have a small time lag. (b) The high-frequency DBS applied to the PD cells as in Fig. 3 (e). (c) The voltage profiles of two Parkinsonian STN cells with the DBS current in (b). (d) The low-frequency DBS applied to the PD cells as in Fig. 5 (a). (e) The subthreshold voltage profiles of two Parkinsonian STN cells with the DBS current in (d).

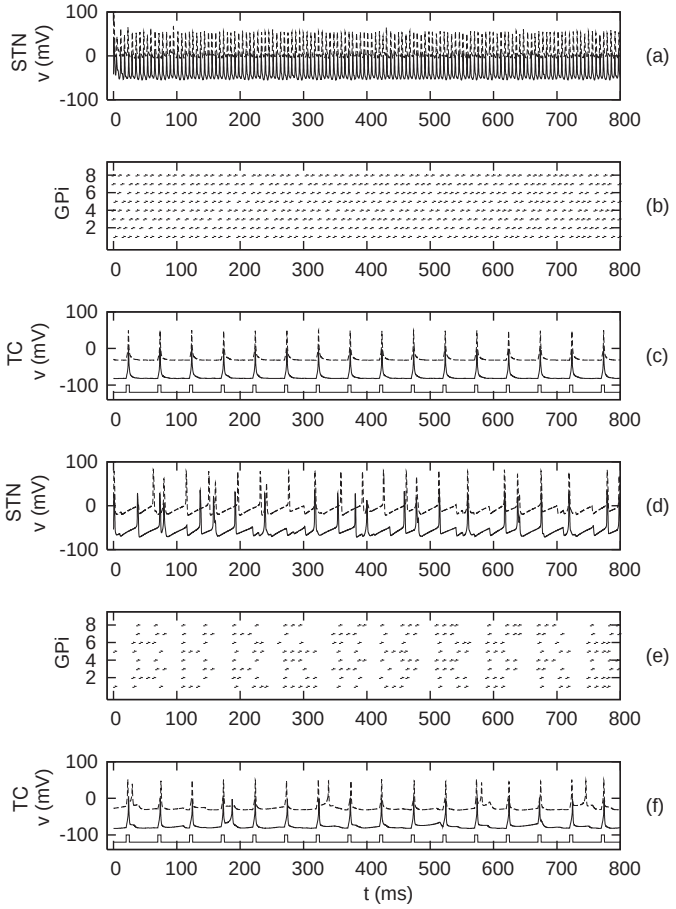

Figure 7: Robustness of DBS effects. A Gaussian filter of standard deviation 2ms is first applied to the DBS current to smooth the edges of the square pulses; i_D is increased to compensate for the loss in DBS amplitude due to filtering. The DBS currents received by each STN cell are then multiplied by an independent, random factor with mean 1 and standard deviation 0.1 (over the full extent of the DBS input), to model a 10% heterogeneity in DBS effects. (a) Voltage profiles of two STN cells following the filtered, heterogenous high-frequency DBS input in Fig. 3. (b) Raster plot of the Parkinsonian GPi cells with the high-frequency DBS current. (c) Voltage traces for the TC cells, demonstrating that faithful transmission of model sensorimotor inputs is maintained. (d) Voltage profiles of two STN cells following the heterogeneous, filtered lower-frequency DBS input of Fig. 5(a), showing that desynchronized GPi spiking persists among cell pairs. (e) Raster plot of the Parkinsonian GPi cells with the low-frequency DBS current. (f) Voltage profiles of TC cells show that transmission remains largely reliable.

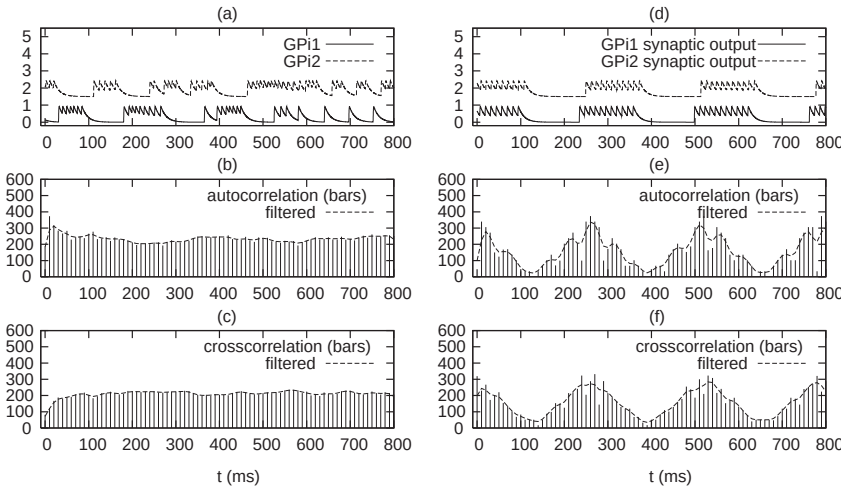

Figure 8: Left panel: normal cells; right panel: Parkinsonian cells. (a) and (d): the synaptic (conductance) output from two GPi cells; (b) and (e): the autocorrelation of GPi1's synaptic output; (c) and (f): the autocorrelation between GPi1 and GPi2's synaptic outputs. The calculation of the auto- and cross-correlation is described in Section "Statistical measure of Gpi spiking patterns."

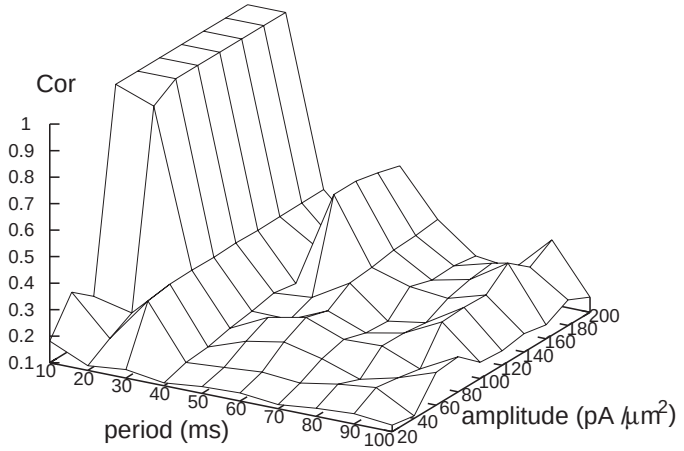

Figure 9: The dependence of Cor, the statistical measure of GPi cells' spike patterns, on DBS pulse period ρ_D and amplitude i_D. As in Fig. 4, the impulse duration σ_D is fixed at 0.6ms. The region where $Cor > 1.0$ is truncated to show the complex landscape at lower frequencies.

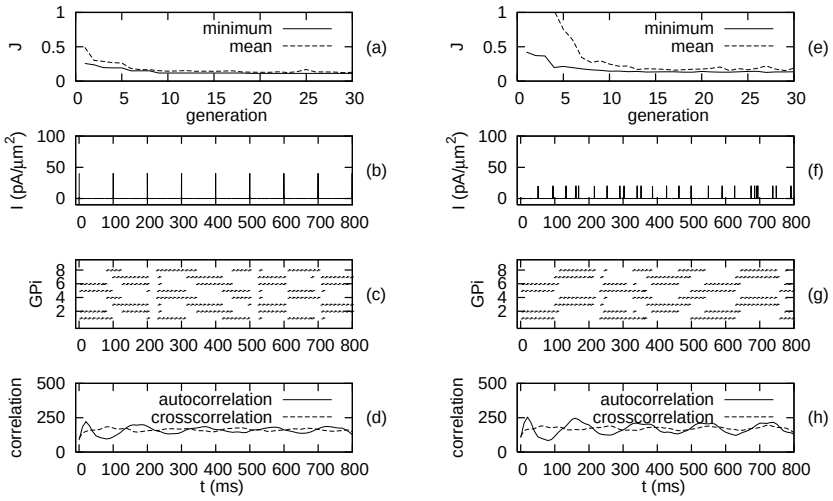

Figure 10: GA optimization of the statistical measure Cor. Left panel: periodic DBS, optimization with the GA cost function value $J = Cor$. Right panel: stochastic DBS, optimization with $J = Cor + w \int_{t=0}^{T} I_{DBS} dt$. (a) and (e): the evolution of J over generations of the GA; minimum and mean J values are given for the parameter settings tested at each generation. (b) and (f): the optimal currents found by the GA. DBS current parameters are $\rho_D = 100ms$, $i_D = 40pA/\mu m^2$, $\sigma_D = 0.6ms$, giving $< I_{DBS} >= 0.24pA/(\mu m^2 ms)$ for the periodic current and $i_D = 20pA/\mu m^2$, $\sigma_D = 0.4ms$, $< I_{DBS} >= 0.21$ for the stochastic current. (c) and (g): raster plot of GPi spike times. (d) and (h): GPi1's (filtered) spike time autocorrelation and GPi1-GPi2 (filtered) spike time crosscorrelation.

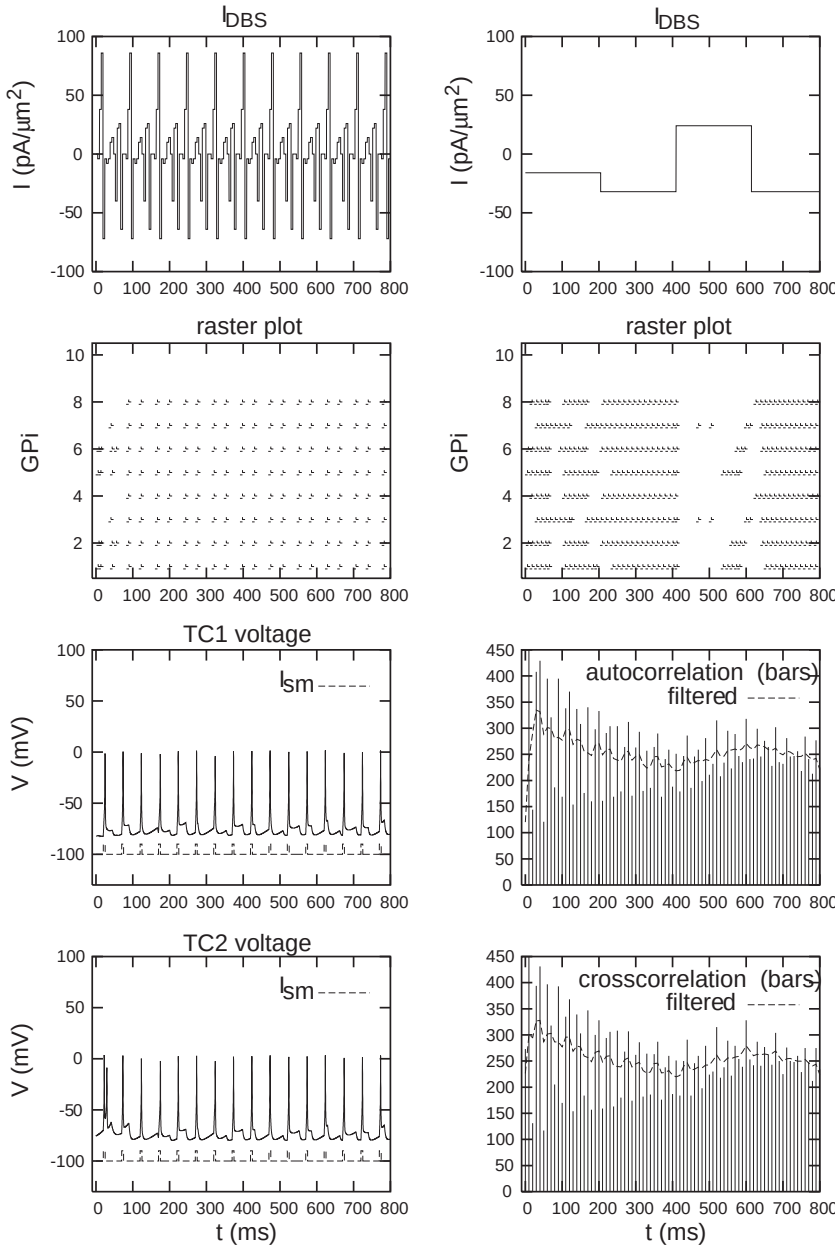

Figure 11: Optimization of I_{DBS} represented by discrete Haar basis (Eq. 7). Left panel: using the *Rel* cost function with current weight $w = 10^{-5}$. Plots from the top are: one optimal DBS current found by the GA, the corresponding GPi raster plot, spike pattterns of TC1 and TC2. Right panel: using the *Cor* cost function with no current cost. Plots from the top are: an optimal DBS current found by the GA, the corresponding GPi raster plot, GPi1 spike time autocorrelation, and $GPi1 - GPi2$ crosscorrelation.

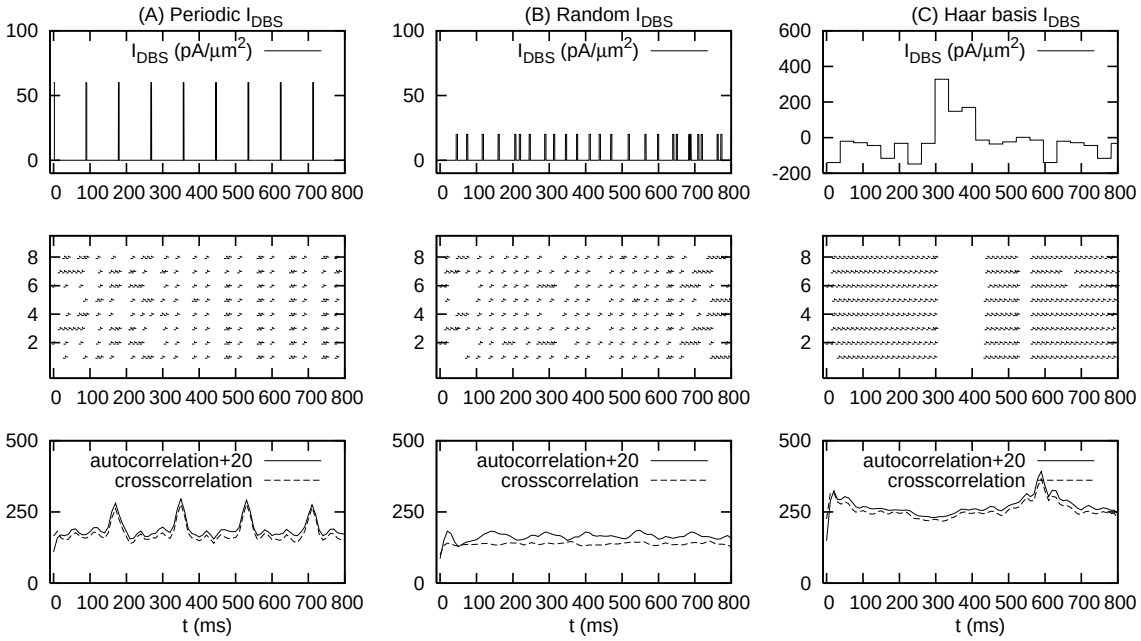

Figure 12: Optimized I_{DBS} currents in the case of 1% random heterogeneity in the intrinsic network currents of STN and GPe cells, as well as their synaptic interactions (see text). 10% heterogeneity was also included in the DBS input into different STN neurons. The *Cor* cost function with no current weight was used in all cases shown. Panel A) periodic DBS; panel B), stochastic DBS; panel C); DBS currents defined via the Haar basis. For each panel, plots from the top are: optimal DBS current found by the GA, the corresponding GPi raster plot, and autocorrelation of the first Gpi cell plotted with the crosscorrelation between spike times of Gpi cells 1 and 2. DBS current parameters are $\rho_D = 89ms$, $i_D = 60pA/\mu m^2$, $\sigma_D = 2.2ms$, giving $< I_{DBS} >= 1.5pA/(\mu m^2 ms)$ for the periodic current and $i_D = 20pA/\mu m^2$, $\sigma_D = 3.5ms$, $< I_{DBS} >= 2.0pA/(\mu m^2 ms)$ for the stochastic current.