Chapter 16
Neuroswarm: A Methodology to Explore the Constraints that Function Imposes on Simulation Parameters in Large-Scale Networks of Biological Neurons

David Gomez-Cabrero, Salva Ardid, Maria Cano-Colino, Jesper Tegnér and Albert Compte

Abstract Candidate mechanisms of brain function can potentially be identified using biologically detailed computational models. A critical question that arises from the construction and analysis of such models is whether a particular set of parameters is unique or whether multiple different solutions exist, each capable of reproducing some relevant phenomenology. Addressing this issue is difficult, and systematic procedures have been proposed only recently, targeting small systems such as single neurons or small neural circuits [16] (Marder and Taylor, Nat Neurosci 14:133–138, 2011), [1] (Achard and De Schutter, PLoS Comput Biol 2:e94, 2006). However, how to develop a methodology to address the problem of non-uniqueness of parameters in large-scale biological networks is yet to be developed. Here, we describe a computational strategy to explicitly approach this issue on large-scale neural network models, which has been successfully applied to computational models of working memory (WM) and selective attention [2] (Ardid, J Neurosci Off J Soc Neurosci 30:2856–2870, 2010), [3] (Cano-Colino et al., Cereb Cortex 24:2449–2463, 2014). To illustrate the approach, we show in this chapter how our strategy applies to the problem of identifying different mechanisms underlying visuospatial WM. We use a well-established biological neural circuit model in the literature [6] (Compte et al., Cereb. Cortex 10:910–923, 2000) as a reference point, which we then

D. Gomez-Cabrero (✉) · J. Tegnér
Unit of Computational Medicine, Center for Molecular Medicine, Department of Medicine, Karolinska Institutet, Solna, Sweden
e-mail: david.gomezcabrero@ki.se

D. Gomez-Cabrero · S. Ardid · M. Cano-Colino · A. Compte
Institut d'Investigacions Biomediques August Pi i Sunyer (IDIBAPS), Barcelona, Spain

S. Ardid
Center for Computational Neuroscience and Neural Technology (CompNet), Department of Mathematics and Statistics, Boston University, Boston, USA

M. Cano-Colino
Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Lisbon, Portugal

L. Geris and D. Gomez-Cabrero (eds.), *Uncertainty in Biology*,
Studies in Mechanobiology, Tissue Engineering and Biomaterials 17,
DOI 10.1007/978-3-319-21296-8_16

perturb by using the Swarm Optimization Algorithm. This algorithm explores the space of biologically unconstrained parameters in the model under the constraint of preserving a solution defined here as a network in which the activity of model neurons mimics the properties of neurons in the dorsolateral prefrontal cortex (dlPFC) of monkeys performing a visuospatial WM task [7] (Funahashi et al., J Neurophysiol 61:331–349, 1989). The results are: (1) identification of a set of model solutions, composed of alternative and, in principle, feasible and sufficient mechanisms generating WM function in a cortical network. In particular, we found that the dynamics of interneurons play a main role in distinguishing among potential circuit candidates. Secondly we uncovered compensatory mechanisms in a subset of the parameters in the model. In essence, the compensatory mechanisms we observe in the different solutions are based on correlations between sets of parameters that shift the local Excitatory/Inhibitory balance in opposite directions. In summary, our approach is able to identify distinct mechanisms underlying a same function, as well as to propose a dynamic solution to the problem of fine-tuning. Our results from the proposed workflow would be strengthened by additional biological experiments aimed to refine the validity of the results.

Keywords Prefrontal cortex · Workflow · Ensemble analysis · Working memory model · Neuroscience · Computational biology

16.1 Introduction

A branch of Computational Biology makes use of mathematical modelling (such as differential ordinary equation systems) to understand better the mechanisms of the biological system of interest. In those cases, models are tools to test and generate hypotheses, to then validate experimentally. But the use of models is not trivial and requires robust methodologies of data analysis, model generation [14, 22], parameter estimation [4, 23], and experimental design [27]. A particularly crucial decision in this process is how complexity and uncertainty are being considered during the modelling [11, 13, 15]. However, the major challenge in using computational models under uncertainty is the generation of relevant hypotheses that are not exclusively dependent on choices during modelling (such as parameter selection). In this chapter, we provide a methodology for robust hypothesis generation in the context of Neuroscience and under parameter uncertainty; relevant work in the book addressing similar challenges are described in [5, 15].

Mechanistic aspects of brain function can be studied with the use of biologically detailed computational models. Those models detail relations and/or interactions between entities through mathematical formulations that depend on a set of parameter values. The first major success in modelling was the classical model of Hodgkin and Huxley ([10] HHM) of the action potential. HHM was developed to provide a mechanistically unified system description by mathematically organizing experimental observations. Interestingly, HHM not only described known facts but also

allowed to generate predictions that only many years later, when technical development allowed, were experimentally validated [9]. The success of the model started what is known nowadays as computational biology. Typically, a general approach for modelling-based studies has three phases: (1) the collection of relevant experimental facts and expert knowledge to be considered, (2) the mathematical description of the system, and (3) the fine-tuning of the parameters to reproduce the "expected" behaviours.

A critical question that arises from that approach is to what extent conclusions depend on particular simulation parameters (see [5, 15]). However, to demonstrate whether a model is unique in reproducing some relevant phenomenology can be a tall order. Interestingly, the questions when more than one solution (parameter set) reproduce expected experimental behaviours (which we will term "solution" parameter sets), uncover many other critical questions such as: (1) are those parameter sets providing different biological mechanisms or predictions? [5], (2) do parameter sets group into families of solutions or do they provide a continuum of solutions? Following these questions [16] proposed the need to investigate in populations of parameter sets in order to discover compensatory mechanisms in neurons or circuits.

In the context of single neurons and small neural networks this problem has been addressed by constructing and analysing databases of models compatible with biological function [1, 20, 21]. Small networks of three interconnected cells could reproduce the rhythmic patterns of activity in the crustacean stomatogastric ganglion for millions of different, disparate parameter combinations [21]. In the analysis of a cerebellar Purkinje cell models authors found 20 different *solution* models (i.e. parameter sets compatible with experimental data) [1]. Furthermore, by studying the parameter landscape created by the good models it was found that the *parameter space of good models* could be defined by a set of "*loosely connected hyperplanes*" [1].

In the present chapter we present a computational strategy to explicitly explore, group and characterize parameter-sets on large-scale neural network models. This strategy is similar to that described in [16], but adapted to deal with the complexity and computational cost of large-scale neural simulations. We applied this strategy to study a specific cognitive function, visuospatial working memory, which can be modelled with a biological neural network [6] that mimics the properties and dynamics of neurons in the dorsolateral prefrontal cortex (PFC) of monkeys engaged in oculomotor delayed response tasks [7]. In such tasks, the monkey is required to retain the location of a visual cue during a delay period between the cue stimulus and the memory-guided saccadic response, and PFC neurons reflect this memorized information through selective persistent activation in the delay period [7].

The typical experimental design is depicted in Fig. 16.1a. The trial starts with a blank screen containing just a central cross on which the monkey fixates its gaze to initiate the trial ("pre-stimulus"). While fixating, a stimulus cue appears briefly in one of eight possible locations equidistant from the fixation point ("stimulus"). After cue presentation a delay period of a few seconds follows during which the monkey needs to remember the location of the previously presented cue ("delay period"). At the end of the delay period, the fixation cross disappears and the monkey makes a

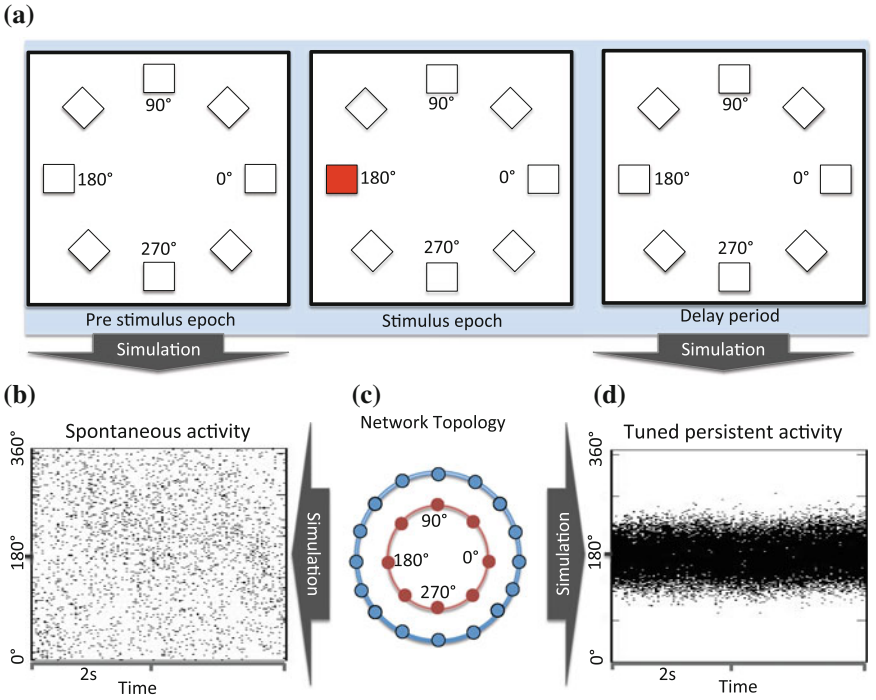

Fig. 16.1 The experiment, the model and the simulation results explained. **a** Three stages of the experiment. In a first stage there is no stimulus and a blank screen is observed. In a second Stage a visual stimulus is briefly flashed in one of 8 possible locations; finally in a third stage the stimulus disappears but it must be remembered during few seconds. The task consists in reporting the location of the briefly flashed stimulus after a delay of a few seconds and it thus require memorising this location. **c** Topology of the neuronal network used to model neural activity during this task. Neurons have a 1-D ring topology, and their position in the ring is associated to the possible location of the stimulus. There are two concentric rings, a ring of pyramidal (excitatory) neurons (in *blue*) and a ring of interneurons (inhibitory neurons, in *red*). The position of the two rings denotes the selectivity of excitatory and inhibitory neurons. Neurons are connected according to their relative angular distance on the ring. **b** Spontaneous firing pattern of the neurons over time. A point denotes a given neuron (y-axis) firing in a given time (x-axis). **d** Persistent firing pattern of the neurons associated with the location of the presented stimulus

saccadic eye movement to the location where he remembers the cue was presented. In experiments, PFC neurons show tuned persistent activity in the delay period of this behavioural protocol [7]. This experimental design can also be simulated using a network model of excitatory and inhibitory neurons [6]. By arranging neurons according to spatial selectivity (Fig. 16.1c), connectivity parameters can be tuned so that strong local excitation and strong global feedback inhibition combine to produce neuronal responses in line with experimental data. Figure 16.1b, d shows in rastergrams the activity of model pyramidal neurons in one tuned network (see next section) during the pre-stimulus and delay period epochs, respectively [6]. Notably, activity in the pre-stimulus epoch appears uniform and at low firing rates, and we

term this condition "spontaneous activity", and network activity in the delay period is tuned and sustained, and we call this "tuned persistent activity".

The methodology proposed here is similar to the one described in [8], which was used to analyse a computational model of atherosclerosis. We update the necessary steps to apply it in Neuroscience, in particular for the analysis of the PFC model (see also [2, 3]). Interestingly our approach is able to identify: (1) compensatory mechanisms, (2) characteristics of inhibitory neurons firing patterns that allow the grouping of solutions, and (3) significant opposite compensatory mechanisms in some of the groups identified. Therefore, our computational approach that explores the "solution space" identifies relevant mechanisms to be further tested by appropriate experimental designs.

Section 16.2 of the chapter presents a review of the PFC model. Section 16.3 details how feasible parameter sets are searched. The following section, Sect. 16.4, summarizes the integrative analysis of the parameter sets that are considered to be correct based on the experimental data. The final section, Sect. 16.5, presents the conclusions regarding the utility of the proposed approach and summarizes the biological results associated to PFC.

16.2 PFC-Working Memory Model

We used the PFC network model described in [6]. We refer the reader to this publication for a thorough account of the computational model and we provide here only a succinct description. The network contains 1,024 excitatory neurons and 256 inhibitory neurons modelled according to the leaky integrate-and-fire formalism [24]. Model neurons are arranged according to their preferred cue directions in a ring topology as shown in Fig. 16.1c and they are interconnected via conductance-based synapses with dynamics consistent with AMPA, NMDA and $GABA_A$ receptor mediated synaptic transmission in the cortex. Specifically, AMPA and $GABA_A$ synaptic conductances jump instantaneously when a pre-synaptic spike occurs and decay exponentially with time constant 2 ms for AMPA and 10 ms for $GABA_A$.

NMDA conductances are voltage dependent and their dynamics are defined by a rise time (set to 2 ms), a decay time (set to 100 ms) and a saturation term such that they become insensitive to high presynaptic firing rates.

$$\frac{ds}{dt} = -\frac{1}{\tau_s}s + \alpha_s x(1-s)$$

$$\frac{dx}{dt} = -\frac{1}{\tau_x}x + \sum_i \delta(t - t_i)$$

While the parameters defining these time dynamics are relatively well constrained by experimental data, the strength of the conductances are much more unconstrained

Notation	Description	Lower Bound	Upper Bound	Value of Reference
$G_{EE,AMPA}$ (nS)	Recurrent excitatory-to-excitatory synaptic conductance mediated by AMPAR channels	0	0.98	0
$G_{EE,NMDA}$ (nS)	Recurrent excitatory-to-excitatory synaptic conductance mediated by NMDAR channels	0	1.95	0.76
$G_{EI,AMPA}$ (nS)	Recurrent excitatory-to-inhibitory synaptic conductance mediated by AMPAR channels	0	0.98	0
$G_{EI,NMDA}$ (nS)	Recurrent excitatory-to-inhibitory synaptic conductance mediated by NMDAR channels	0	1.95	0.58
G_{IE} (nS)	Recurrent inhibitory-to-excitatory synaptic conductance mediated by GABAAR channels	0	7.81	2.67
G_{II} (nS)	Recurrent inhibitory-to-inhibitory synaptic conductance mediated by GABAAR channels	0	7.81	2.05
σ_{EE} (°)	Excitatory-to-excitatory width of the connectivity footprint	7.2	360	14.4 **
σ_{EI} (°)	Excitatory-to-inhibitory width of the connectivity footprint	7.2	360	Unstructured *
σ_{IE} (°)	Inhibitory-to-excitatory width of the connectivity footprint	7.2	360	Unstructured *
σ_{II} (°)	Inhibitory-to-inhibitory width of the connectivity footprint	7.2	360	Unstructured *
J^{+}_{EE}	Excitatory-to-excitatory strength of the stronger isodirectional connections	0.8	3	1.62
J^{+}_{EI}	Excitatory-to-inhibitory strength of the stronger isodirectional connections	0	2	1*
J^{+}_{IE}	Inhibitory-to-excitatory strength of the stronger isodirectional connections	0	2	1*
J^{+II}	Inhibitory-to-inhibitory strength of the stronger isodirectional connections	0	2	1*
$g_{ext,E}$ (nS)	External excitatory (AMPAR-mediated) conductance on excitatory cells	0	6	3.1
$g_{ext,I}$ (nS)	External excitatory (AMPAR-mediated) conductance on inhibitory cells	0	6	2.38

(a) Parameter Description

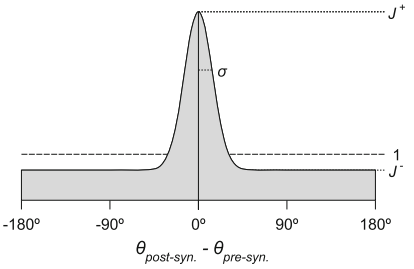

(b) Structured connectivity of the model.

Fig. 16.2 Parameters selected to vary in the PFC-model. **a** Parameter names, description and lower and upper bounds are included. (*) cross- and isodirectional components of these connections were equally strong (i.e. $J^{+} = 1$; in this case, the value of σ is irrelevant; see panel **b**). (**) corrected from [6]. **b** Structured connectivity of the model. The synaptic connection strength decreases with the difference in the preferred cues of two neurons, with strong interactions between neighboring neurons and weak interactions between more distant neurons

and need to be tuned to achieve the required function. We impose however a topographical constraint, so that the strength of synaptic conductances is a function of the difference in the preferred cues of the presynaptic and postsynaptic neurons. We specified this function to be a Gaussian, defined by three parameters (see Fig. 16.2b): the width σ, the tuning parameter J^{+}, and the overall strength g. It has been shown that if excitatory connections among excitatory neurons are such that synaptic strength decreases with the difference in the preferred cues of two neurons, with strong interactions between neighbouring neurons and weak interactions between more distant neurons, the network has regimes of operation compatible with working memory physiology: tuned persistent delay period activity (Fig. 16.1d) bi-stable with a low-rate, unstructured spontaneous activity (Fig. 16.1b).

For a given set of parameters the model is simulated for 5 s in repeated trials. Some trials are run without any phasic external stimulation to test for the stability of the spontaneous activity (Fig. 16.1b). In other trial simulations (persistent activity trials), a stimulus is applied by transiently injecting current to a subset of neurons after the first second of simulation. Every time a model is run a numerical seed is

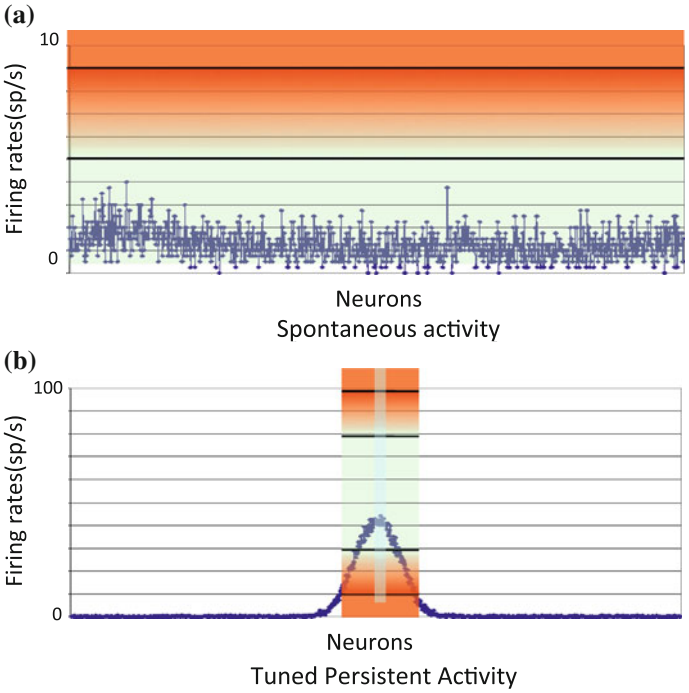

Fig. 16.3 Evaluation functions in plots. **a** Expected firing rate for spontaneous activity. The maximum firing rate is limited to 8 sp/s, but penalized if more than 5 sp/s. Also there is a penalty for the fraction of pyramidal neurons that are silent. **b** Expected firing rate pattern for persistent activity. The maximum firing rate is limited to 100 sp/s, and penalized if it is more than 80 sp/s. The minimum firing rate is 10 sp/s, and penalized if it is less than 30 sp/s

randomly selected that controls the timing of nonspecific external Poisson spiking activity that depolarizes the network neurons and set a general random background activity in the network. The output of the model for a given random seed is the timing of the spiking events for all the neurons included in the model (*rastergram*); examples of rastergrams with the desired patterns of activity are presented in Fig. 16.1 for persistent activity (Fig. 16.1d) and for spontaneous activity trials (Fig. 16.1b). The stationary patterns in these rastergrams can be summarized with *firing rates* (measures in *spikes per second*) computed in a final 4 s window of the simulation. Firing rate plots are shown in Fig. 16.3. For the same parameter set the rastergram and firing rates may differ between different trial simulations depending on the initial numerical seeds selected.

16.3 Neuroswarm: A Tool to Explore the Parameter Space

While the network solution of Compte et al. [6] (Fig. 16.1b, d) had qualitative features consistent with experimental data, some aspects of model function did not match quantitatively: the model displayed a large gap between persistent activity to preferred and non-preferred stimuli (>20 sp/s) while experimentally this is a narrower gap (<10 sp/s); the model required a very precise symmetry in the translationally invariant connectivity; connectivity parameters required a significant degree of fine tuning; and neuronal activity in the persistent state was more stable in the model than observed experimentally. These discrepancies could mark fundamental flaws in the model or some of them could be specific of the dynamical regime ensuing from our particular choice of parameters and could be alleviated in a different parameter regime. Here, we will design a protocol to address this question by exploring network behaviour in very different parameter configurations. The model has over 100 parameters, and we selected 16 that we considered to be unconstrained by experimental data and relevant candidates to regulate the expected behaviour of the model. The name of the parameters, their ranges and their lower and upper bounds are detailed in Fig. 16.2a; Fig. 16.2a includes also the values originally considered in [6].

A likely hypothesis is that there is no unique but several combination of these parameters that are able to produce network activity as in Fig. 16.1b, d, in qualitative agreement with experimental data [7]. To test this hypothesis we need to find different parameter sets that are in agreement with the observed experimental results. To this end, we need to define two major elements: (1) a method to evaluate the correspondence of network function with the expected dynamics (Fig. 16.1b, d) for each parameter set, and (2) a method to explore the parameter space to search efficiently for parameter sets that match the required function optimally. In addition, and considering the computational costs implicated in the exploration of this high-dimensional parameter space, we designed a heuristic approach to extend the search by exploring the linear relations between good candidate solutions. We termed our implementation of this procedure "*Neuroswarm*" because it used a Particle Swarm algorithm to search for optimal solutions, but it can be readily extended to other optimization procedures. In the following we describe each of the steps involved.

16.3.1 Evaluating a Parameter Set

We defined a set of network activity properties expected from a "feasible" parameter set, and we specified each such property in a mathematical fitness function of the simulated network activity. Thus, after each simulation we could compute one fitness value for each of the expected network activity properties so that we could then compute a single evaluation value for each parameter set as a weighted sum of all these fitness values. The overall evaluation function is a cornerstone in the searching process as it defines the landscape and therefore it is largely associated to the

difficulty of the search. Evaluations are designed to have linear components to allow the optimization over gradients. Fitness functions were computed from the spiking times of only excitatory neurons in each simulation where for every parameter set we ran five simulations with different initial random seeds. The five fitness functions that we employed were the following:

- FF1. *Maximum firing rate in spontaneous activity:* a maximum firing rate of 8 sp/s in *spontaneous activity*. A linear fitness function taking values from 0 to 1 starts penalizing maximum firing rates of 5 sp/s (penalization of 0) and penalizes with a maximum value of 1 firing rates of 8 sp/s and above. We ran 5 different network simulations with different random seeds and took the largest penalization as fitness value here. See Fig. 16.3a.
- FF2. *Percentage of silent neurons in spontaneous activity:* the amount of excitatory neurons being silent (firing rate of 0 sp/s) during *spontaneous activity* is penalized. The fitness function takes the fraction of silent excitatory neurons. The maximum fraction computed from 5 simulations is taken as the FF2 evaluation. See Fig. 16.3a.
- FF3. *Maximum and minimum firing rate in persistent activity:* for excitatory cells targeted by the stimulus in *persistent activity* trials, the preferred peak firing rate is in the range 30 sp/s to 80 sp/s. A linear fitness function starts penalizing maximum firing rates of 80 sp/s (penalization of 0) up to 100 sp/s (penalization of 1). Similarly, the linear fitness function starts penalizing minimum firing rates of 30 sp/s (penalization of 0) up to 10 sp/s (penalization of 1). See Fig. 16.3b. The penalization is computed for 5 simulations and the maximum is selected.
- FF4. *Asynchronous activity:* parameter sets that generate extremely synchronized activity patterns during *spontaneous* and/or *persistent activity* are also penalized. For 5 simulations of each type we compute the average binary Pearson correlation between all pairs of neurons computed over windows of 4 ms. The maximum among the 5 values is selected as FF4. In Fig. 16.4 we provide examples of what we consider asynchronous activity (upper panel) and synchronized activity (lower panel).
- FF5. *Homogeneity:* we run five simulations with different random seed for each of the two conditions of interest: *spontaneous activity* (Fig. 16.1b) and *persistent activity* (Fig. 16.1d), and we penalize networks that did not provide stable results across these five simulations in each case. The penalization is computed by subtracting to 1 the average of the p-values computed from comparing the firing rate distributions across pairs of simulations by the Kolmogorov-Smirnov test; a value close to 0 denoted highly correlated distributions.

Finally, each fitness function had an associated weight to compute the total fitness evaluation value (TFEV) as its weighted average. We found that it was necessary to assign the highest weights to FF3 (5×10^3) and FF2 (10^3) fitness functions described above in order to find satisfactory solutions; the weights assigned to other fitness function were 10^2, 10^2 and 10^0, respectively for FF1, FF4 and FF5. Thus, for each parameter set, 5 simulations were run in each condition (spontaneous activity and persistent activity), one fitness value was obtained by evaluating the ensuing spiking activity for each of 5 fitness function and this was all combined in one single

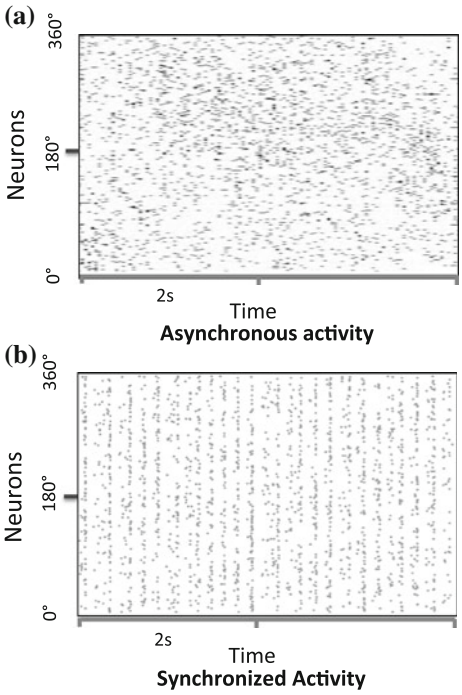

Fig. 16.4 Asynchronous versus Synchronized neuron activity. Rastergrams showing examples of what we consider **a** asynchronous and **b** synchronized activity

evaluation value (TFEV) that characterized how well activity for this parameter set matched the required working memory function depicted in Fig. 16.1. The objective was then to find those parameter sets that yielded the lowest TFEV.

16.3.2 Exploring the Parameter Space

We sampled the parameter space in order to find parameter sets which yielded simulations with stable spontaneous activity and tuned persistent activity as illustrated in Fig. 16.1 (we denote these sets as *feasible parameter sets, FPS*). We assessed network function for each parameter set using the fitness functions defined above. Because of non-linearities both in the model and in the parameter set evaluation, no exact algorithm can be used to find the parameters that minimize the evaluation value. We used instead a heuristic algorithm to search for FPS: the Particle Swarm Optimization Algorithm (PSO; [12]). As a black-box optimization algorithm, the PSO can operate with any fitness function and it was originally designed to search optimally in a hyperspace of real numbers [12]. More detailed description of the algorithm is provided in [4] elsewhere in this volume. In brief, each model instantiation defined by a particular choice of the values of 16 parameters (Fig. 16.2a) represents one particle in a 16-dimensional space in the PSO. We run simulations

for each network model (particle) and compute the TFEV value for that particle. By working simultaneously with a large number of particles, PSO establishes parameter updates that move all particles towards the particle with the best TFEV value, while also attracting each particle towards its previous best evaluation. This optimization process has been shown to be effective in several applications [19] and the authors of this chapter have shown its usefulness in several cases in Neuroscience [2, 3].

Due to the high computational demands (in our equipment each network simulation ran for approximately 15 min and each iteration of the PSO algorithm required 500 simulations, for near 50 iterations in total), we executed all simulations in parallel in a computational facility with 200 available CPUs and with a modified parallelised code (Grid-SuperScalar technology at Barcelona Supercomputing Center—Centro Nacional de Supercomputacion for which we adapted the code).

To sample among the FPS we ran PSO 31 times, in each run we considered 50 particles, with a maximum of 50 iterations. An earlier stop was considered if the method was not finding a better-evaluated solution for more than 15 iterations. To define the working FPS we (1) first selected for each PSO run the best evaluated parameter set; then (2) we discarded 2 parameter sets because they were poorly evaluated (PSO did not find good solutions). Finally (3) non-filtered best solutions were used to define thresholds for the fitness functions (considering the maximum for each fitness function). Those thresholds (defined as 110 % of the maximum values) were used to define the selection criteria that defined FPS from all the parameter sets evaluated in all PSO runs; by doing this we were able to recover more than one high-quality parameter set per PSO run. In the analysis, **FPS** denotes the original 29 parameter sets while **rFPS** denotes FPS extended with those high-quality recovered parameter sets.

16.3.3 Increasing the Set of Solutions

Once we had FPS we sought to explore the relation between them by simulating networks (parameter sets) that interpolated between them. We designed an exploratory greedy search of the parameter space. For each pair of the best solutions of the 31 PSO we investigated 9 equidistant points in the linear path between them in parameter space. We call these linearly interpolated parameter sets **LIps**. Figure 16.5 shows an example of the methodology, C1, C2 and C3 are the best solutions, the lines represent the shortest path (line) in 16 dimensions between each pair; each cross represent a new parameter set to be evaluated, a (un)filled cross denotes a (un)feasible model. By using this procedure we can evaluate if different solutions form part of a large continuous region of solutions or else if they are separated by regions without solutions, suggesting that they could constitute qualitatively different solutions.

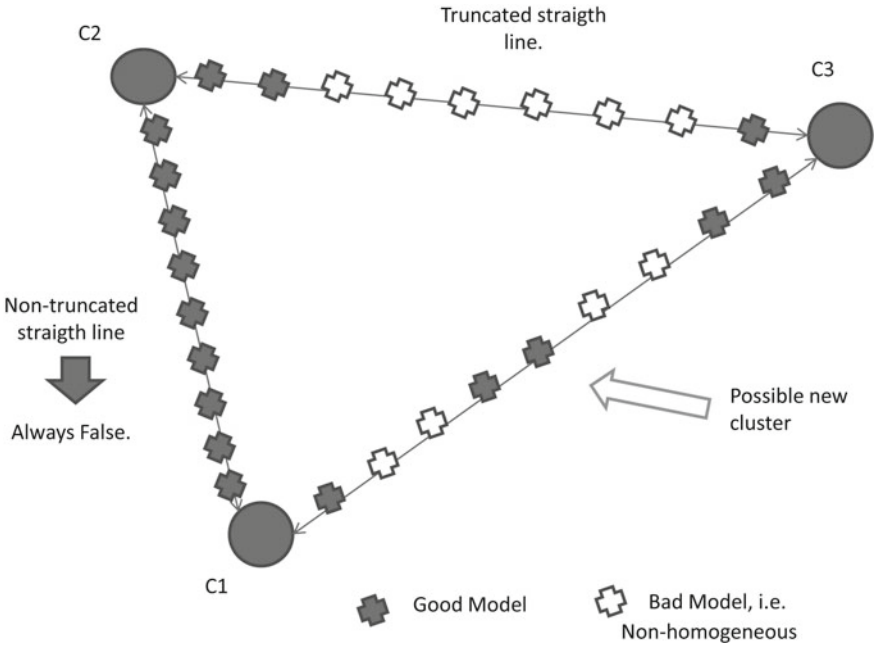

Fig. 16.5 Extending the solution space. All the linear paths between every pair of PSO's best solutions (FPS) are evaluated. For each linear path 9 equidistant points are evaluated. We observed that no single path contained all parameter sets evaluation as good (see "*Truncated straight line*"). In addition, a few possible new clusters of solutions were identified in some paths (see "*Possible new cluster*")

16.3.4 Testing the Allowed Range of Parameters

One concern in this whole procedure is that one needs to define beforehand an allowed range for each of the parameters explored (see Fig. 16.2a), and this may limit the capacity of the search algorithm to find the best solutions. We performed one analysis to test this, and we iterated the procedure if we found that some of the ranges needed to be expanded. We plotted the density functions of the FPS (Fig. 16.6, black continuous), FPS extended with new solutions by linear paths (**LIps**, Fig. 16.6 black discontinuous), and the set of all parameter sets explored (Fig. 16.6, grey). By analysing FPS density plots we can confirm that no parameter had a large number of FPSs clustered at one of the imposed range limits, so that parameter boundaries did not seem to be a limiting factor to find FPS. Interestingly, densities are in many cases centred in the point equidistant between the lower and upper bound; we consider two explanations for this observation: (a) it may reflect a bias generated by the range definition (centred values are explored more often in random trajectories) and (b) this may denote parameters that have no large effects on the network described by the

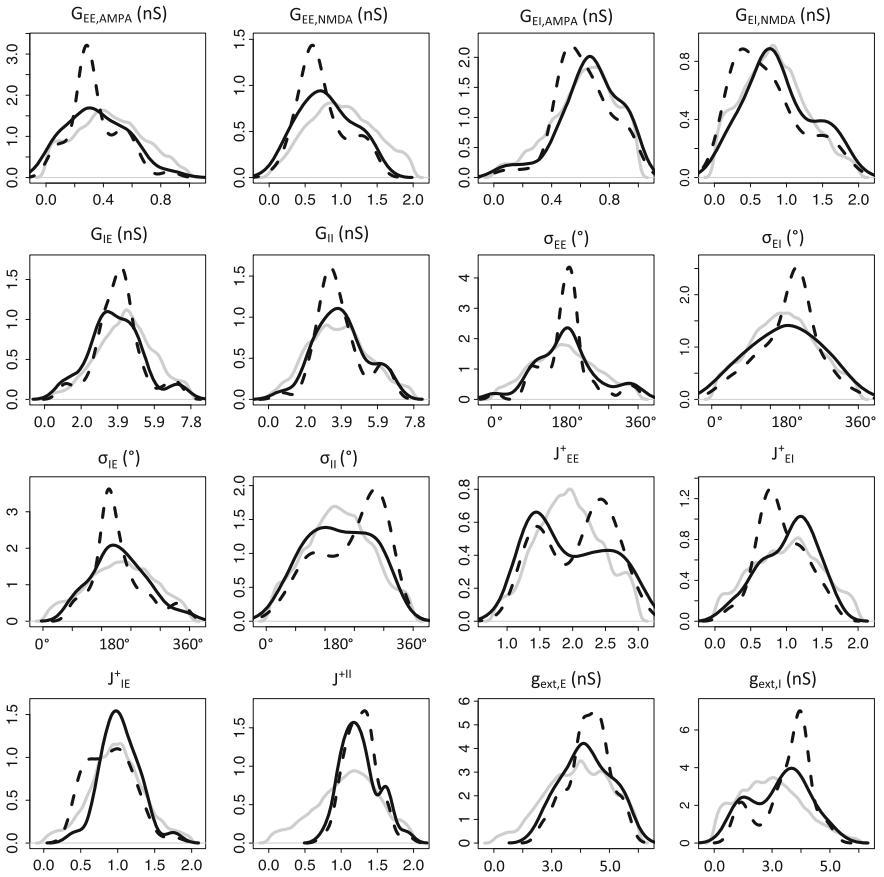

Fig. 16.6 Ranges of parameter in good solutions. For each parameter we evaluated the density of the values observed in all parameter sets evaluated (*grey*), FPS (*black*, continuous line), and rFPS (*black*, discontinuous line)

parameter set. Importantly, despite the benefits of extending FPS by linear paths, we consider that **LIps** reflect biases and have to be considered carefully when statistical analysis in parameter space are conducted.

16.4 Results

We applied the procedure described above and we found 93 different networks that could produce spatial working memory function as described in Fig. 16.1; 29 were identified in FPS and the rest were included through linear path extension (LIps, Fig. 16.5). We then analysed these solutions with the aim of identifying relevant characteristics of the solutions and parameters associated to them.

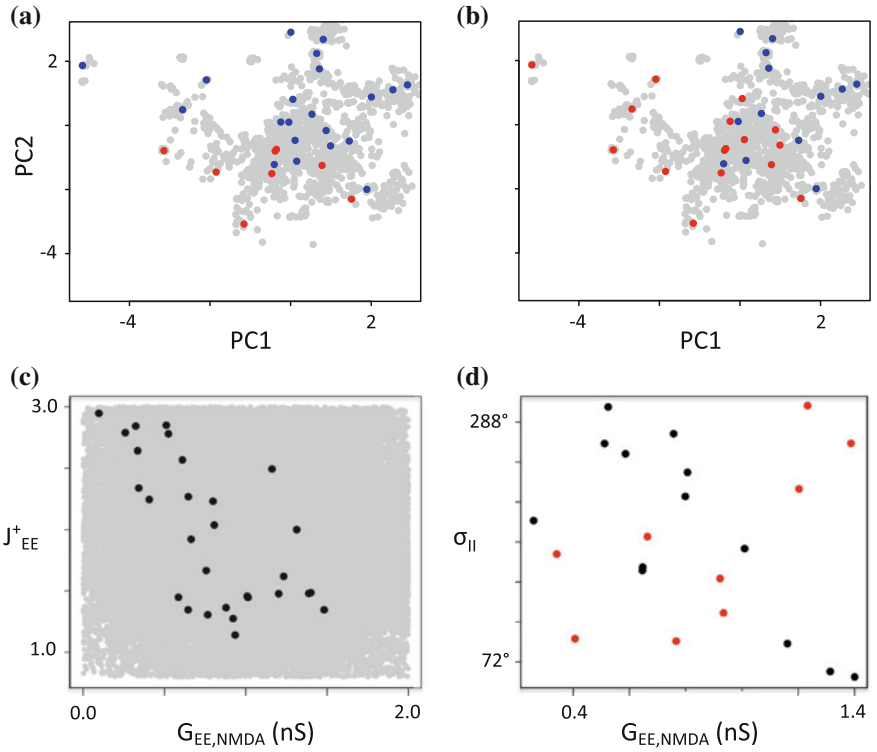

Fig. 16.7 Identification of solution grouping and compensatory mechanisms. **a** Shows the FPS's PCA with all parameter sets evaluated plotted in *gray* in the plane of the 2 first principal components. *Red dots* are FPS with two peaks in the firing rates of inhibitory neurons (see Fig. 16.7a, b) and *blue dots* are FPS with one single peak in inhibitory neuron activity (*blue*, see Fig. 16.7c, d) . **b** Similar to (**a**) but *blue dots* mark FPS where inhibitory neurons had activity centred around the stimulus during persistent activity (Fig. 16.7a, d) and *red dots* mark FPS for peak inhibitory activity opposite to stimulus location (Fig. 16.7b, c). **c** Identified Compensatory mechanism in FPS between $G_{EE,NMDA}$ and J^{+}_{EE}. *Black dots* represent solutions in FPS, while *grey dots* represented discarded parameter sets evaluated during the parameter search. **d** Opposite compensatory mechanism between $G_{EE,NMDA}$ and σ_{I1}: solutions with inhibitory activity centred around the stimulus had a negative correlation between $G_{EE,NMDA}$ and σ_{II} ($\rho = -0.66$, *black dots*), while this correlation was positive for solutions with inhibitory activity maximal 180° away from the stimulus ($\rho = 0.69$, *red dots*)

16.4.1 2-D Representation of FPS to Detect Structure in the Solutions

We applied Principal Component Analysis (PCA) to the parameter sets explored by our search algorithm (Fig. 16.7a). PCA analysis (see [4]) allowed dimensionality reduction from the 16 parameters to the 2 first principal components (which accounted for 33 % of the variance). This bi-dimensional representation allowed us to identify that many of the explored parameter sets clustered around the best solutions (FPS,

Table 16.1 Principal component analysis of FPS

Parameter	PC1	PC2	PC3
$G_{EE,AMPA}$	0.31	−0.27	0.33
$G_{EE,NMDA}$	0.35	−0.25	−0.35
$G_{EI,AMPA}$	0.31	−0.02	−0.15
$G_{EI,NMDA}$	0.03	0.16	−0.56
G_{IE}	0.26	−0.37	0.02
G_{II}	−0.23	−0.06	−0.40
σ_{EE}	−0.23	−0.05	−0.23
σ_{EI}	−0.02	0.20	0.05
σ_{IE}	0.07	0.10	−0.10
σ_{II}	0.08	−0.08	0.27
J^{+}_{EE}	−0.23	0.35	0.29
J^{+}_{EI}	0.42	0.16	0.04
J^{+}_{IE}	0.32	0.33	−0.21
J^{+II}	−0.19	−0.44	−0.02
$g_{ext,E}$	0.16	−0.27	−0.01
$g_{ext,I}$	−0.33	−0.35	−0.08

shown in red or blue). We may consider this initial clustering to be a consequence of the searching algorithm; however, from the observations during LIps computation (Fig. 16.5) we observed that many linear combinations of FPSs were not considered with quality enough to be part of the FPS; therefore we may conclude that FPS is non-convex (but we cannot conclude anything about parameter connectivity in the topological sense [18]).

The loadings of the 3 first PCA components (which accounted for 46 % of the variance) can be found in Table 16.1. Careful inspection of these loadings could help identify what parameters of our simulation were most informative in distinguishing between the different FPS. We found that connectivity strengths, especially among excitatory neurons, and the tuning strengths of all neuron connectivities were discriminating factors between solutions, while the width of the connectivities did not seem to differentiate them significantly.

16.4.2 Clustering Solutions in the Parameter Space

Next we aimed to identify if solutions that are close in parameter space were sharing specific characteristics in their network activity. As a first exploratory approach we considered the solutions as they were mapped on the 2-D PCA reduction. We investigated the rastergrams and firing rates of excitatory and inhibitory neurons. FPS are clearly separated into subgroups if we inspect the activity of inhibitory

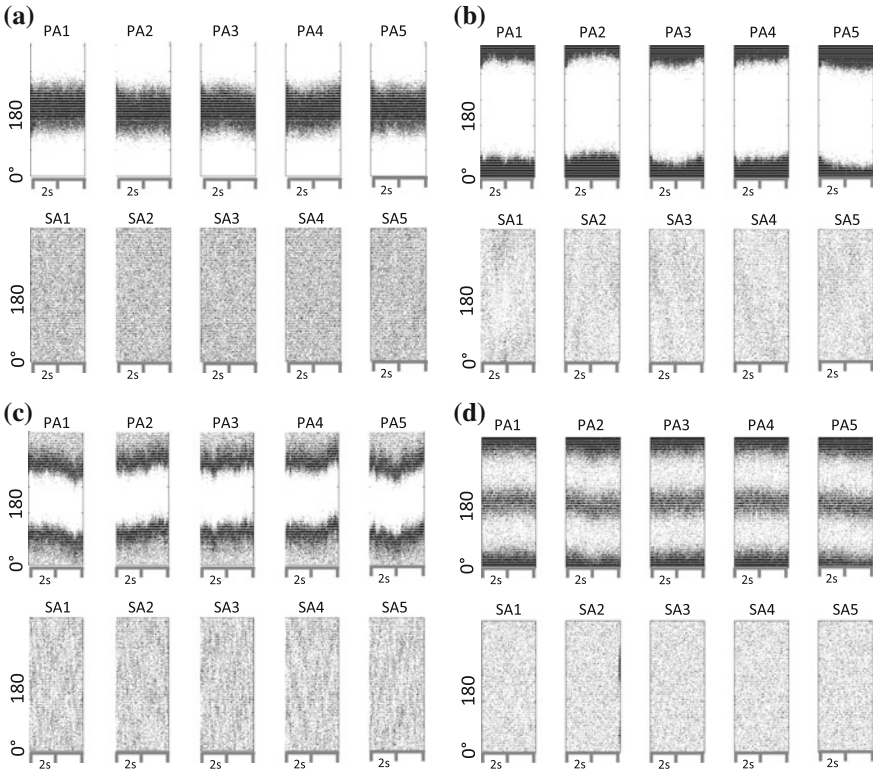

Fig. 16.8 Characterization of inhibitory rasters. Four types of inhibitory rastergrams are shown. In each case, 5 simulations initiated with different random seeds are shown, both for tuned persistent activity (PA 1–5 in the *upper panels*) and for spontaneous activity trials (SA 1–5 in the *lower panels*). These four examples illustrate: **a** single peak of activity centred at 180°; **b** single peak of activity not centred at 180°; **c** two peaks of activity, not centred at 180°; and **d** two peaks of activity centred at 180°

neurons. We considered two characterizations of inhibitory rasters: (1) CvsNC: centred around the presented stimulus (180°, Fig. 16.7a, d) or not centred (Fig. 16.8b, c); and (2) 1vs2: 1 peak (Fig. 16.8a, b) or 2 peaks (Fig. 16.8c, d). In both characterizations there is a separation between groups in the PCA, see Fig. 16.7a, b.

We identified with a Kruskal-Wallis test that J_{EI}^{+} was significantly different (after multiple testing correction) in comparison CvsNC, while J_{II}^{+} changed significantly in the 1vs2 comparison. Interestingly, J_{EI}^{+} and J_{II}^{+} were the most relevant parameters in the first and second component of the PCA, respectively.

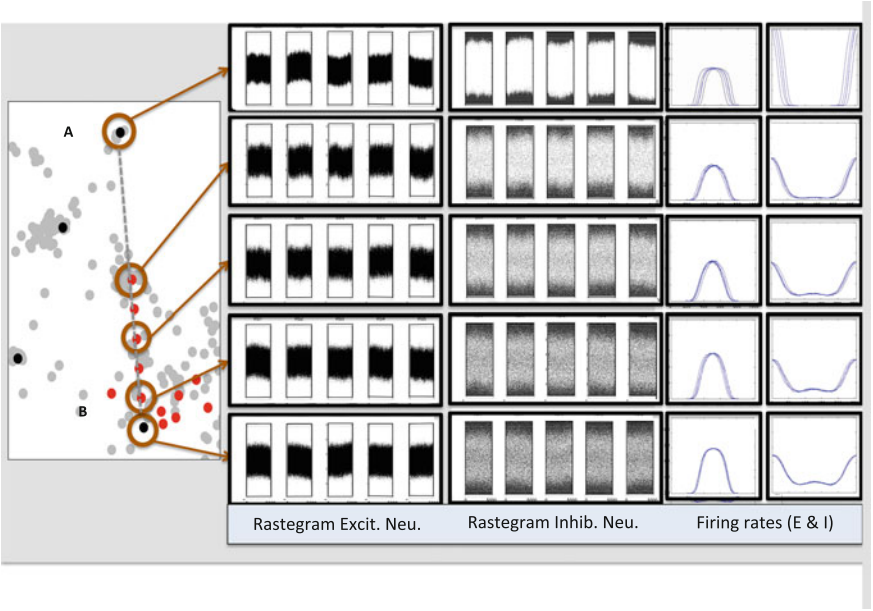

Fig. 16.9 PFC model solutions are linearly unconnected, but new solutions are Inferred: example of a novel cluster in the Linear Path between two FPS-B. The linear path (*grey dashed line*) between two best solutions (**A** and **B**) from PSO is inspected at 9 interpolating points. The *red points* in the *line* denote new feasible solutions. For **A**, **B** and 3 new solutions the rastergram and firing rates of excitatory (*left*) and inhibitory (*right*) are shown. Rastergrams are provided for 5 simulations with different seeds. Average firing rates over the delay period for excitatory and inhibtiory neurons are plotted in the *rightmost* panels, each simulation with a different line to show diffusion of activity in different trials

16.4.3 LIps Identifies Transition Between Solution Types

In LIps we extended FPS by investigating linear combinations of FPS parameter sets (Fig. 16.5). Through this methodology we observed that a straight line of feasible solutions connected no pair of solutions; therefore we validate the non-convexity of FPS.

In LIps we identified novel clusters of high-quality parameter sets; importantly, we found that in all cases the novel cluster firing rates are similar to those of one of the pair of solutions that define the linear combination, possibly suggesting a non-linear connection between them. Figure 16.9 shows an example of a novel cluster. We observe that in this case the characteristics of the behaviour of the novel cluster are similar to those found in Solution B. However, we observe a gradual transformation that reduces the high bump diffusion observed in solution A to the non-existent diffusion of solution B. We conclude that the study of the linear paths helps in the

identification of those parameters associated to specific characteristics. Interestingly, the major differences between Solution A and Solution B are not occurring in pyramidal neurons but in inhibitory neurons.

16.4.4 Studying Compensatory Mechanisms

We were interested in observing whether there were compensatory mechanisms between the parameters that may help to alleviate the fine-tuning problem biologically. We measured this by computing Spearman non-parametric correlation between the values of the parameter sets. We considered only parameter sets in FPS, without including the linear combinations of those, to prevent introducing an artificial bias. Table 16.2 shows all correlations. Among them, we selected those that were statistically significant, strong absolute correlations ($|\rho| > 0.5$). We found three compensatory mechanisms of interest: (1) $G_{EE,NMDA}$ versus J^{+}_{EE} ($\rho = -0.63$, Fig. 16.8c); (2) J^{+}_{EI} versus J^{+}_{IE} ($\rho = 0.62$, expected positive correlation); and (3) J^{+}_{EI} versus $g_{ext,I}$($\rho = -0.53$). These compensatory mechanisms can be interpreted as addressing critical principles of network operation that sustain the required working memory function. Thus, the first of these compensations, relating negatively the strength of $G_{EE,NMDA}$ and the tuning of excitatory connections J^{+}_{EE}, seems to work to keep a sufficient level of excitation locally: if the strength of excitatory synapses is strong enough, local potentiation relative to other synapses is not required and may actually lead to unreasonably high firing rates. The second compensatory mechanism (positive relation between the tuning of excitatory-to-inhibitory and inhibitory-to-excitatory connections) suggests a control of the spatial specificity in the local excitation-inhibition loop. Finally, the third compensation addresses the local excitation-inhibition balance: if inhibitory cells have high spontaneous activity (as a result of increased $g_{ext,I}$), local excitation from excitatory cells should be reduced. These particular requirements for the operation of working memory networks have not been scrutinized in previous studies and suggest new avenues to study the conditions for working memory operation in such networks.

Finally, we also searched for compensatory mechanisms that were opposite in different sub-groups as defined by 1vs2 and/or CvsNC. We found one such opposite compensatory mechanism relating $G_{EE,NMDA}$ and σ in the CvsNC classification (Fig. 16.7d). Thus, for solutions in which inhibitory neurons presented a stable firing pattern with maximal activity centred around the stimulus location (i.e. congruent with maximal activity in excitatory neurons), these two parameters correlated negatively, while they correlated positively for those solutions where inhibitory neurons had activity profiles peaking at the same location as excitatory neurons (Fig. 16.7d). This relationship is intriguing and suggests a different mechanistic link between within population interactions (inhibition to inhibitory neurons, and excitation to excitatory neurons) for each of these two network function organization (CvsNC). This remains to be explored in depth in a future study.

Table 16.2 Correlation structure between parameters in FPS

	$G_{EE,AMPA}$	$G_{EE,NMDA}$	$G_{EI,AMPA}$	$G_{EI,NMDA}$	G_{IE}	G_{II}	σ_{EE}	σ_{EI}	σ_{IE}	σ_{II}	J^{+}_{EE}	J^{+}_{EI}	J^{+}_{IE}	J^{+II}	$g_{ext,E}$	$g_{ext,I}$
$G_{EE,AMPA}$	1.00	0.21	0.18	−0.42	0.45	−0.40	−0.23	−0.06	0.06	0.02	−0.36	0.29	0.03	0.12	0.14	−0.13
$G_{EE,NMDA}$	0.21	1.00	0.19	0.41	0.45	−0.03	0.01	−0.02	−0.10	0.02	−0.63	0.40	0.39	0.13	0.23	−0.11
$G_{EI,AMPA}$	0.18	0.19	1.00	0.20	0.21	0.12	−0.12	−0.17	0.20	−0.05	−0.09	0.05	0.21	−0.34	0.46	−0.19
$G_{EI,NMDA}$	−0.42	0.41	0.20	1.00	−0.33	0.37	0.04	−0.09	0.07	−0.14	−0.15	0.04	0.22	−0.08	−0.02	0.04
G_{IE}	0.45	0.45	0.21	−0.33	1.00	0.18	−0.17	0.22	−0.12	0.02	−0.44	0.15	0.01	0.15	0.27	0.02
G_{II}	−0.40	−0.03	0.12	0.37	0.18	1.00	0.03	0.07	−0.10	−0.31	−0.25	−0.40	−0.21	−0.02	−0.26	0.22
σ_{EE}	−0.23	0.01	−0.12	0.04	−0.17	0.03	1.00	0.08	−0.09	−0.24	−0.16	−0.32	0.00	0.27	−0.12	0.16
σ_{EI}	−0.06	−0.02	−0.17	−0.09	0.22	0.07	0.08	1.00	−0.01	−0.08	0.01	0.08	0.06	−0.20	−0.17	−0.07
σ_{IE}	0.06	−0.10	0.20	0.07	−0.12	−0.10	−0.09	−0.01	1.00	−0.32	0.07	−0.04	0.19	−0.15	0.35	−0.16
σ_{II}	0.02	0.02	−0.05	−0.14	0.02	−0.31	−0.24	−0.08	−0.32	1.00	0.07	0.03	−0.10	−0.10	0.19	0.25
J^{+}_{EE}	−0.36	−0.63	−0.09	−0.15	−0.44	−0.25	−0.16	0.01	0.07	0.07	1.00	−0.05	−0.11	−0.13	0.12	0.00
J^{+}_{EI}	0.29	0.40	0.05	0.04	0.15	−0.40	−0.32	0.08	−0.04	0.03	−0.05	1.00	0.62	−0.11	0.03	−0.53
J^{+}IE	0.03	0.39	0.21	0.22	0.01	−0.21	0.00	0.06	0.19	−0.10	−0.11	0.62	1.00	−0.43	−0.04	−0.44
J^{+II}	0.12	0.13	−0.34	−0.08	0.15	−0.02	0.27	−0.20	−0.15	−0.10	−0.13	−0.11	−0.43	1.00	0.23	0.45
$g_{ext,E}$	0.14	0.23	0.46	−0.02	0.27	−0.26	−0.12	−0.17	0.35	0.19	0.12	0.03	−0.04	0.23	1.00	0.17
$g_{ext,I}$	−0.13	−0.11	−0.19	0.04	0.02	0.22	0.16	−0.07	−0.16	0.25	0.00	−0.53	−0.44	0.45	0.17	1.00

16.5 Conclusions

We show here that large-scale biological simulations of neural networks for specific cognitive functions can be evaluated for generality using an optimization procedure in their high-dimensional parameter space. Typically, these simulations are very unconstrained and generality has been tested using mathematical simplifications in mean field formulations. While this approach is indeed very general, the initial assumptions on the simplifications to use may constrain the validity of the results to a subset of possible solutions. Our computational workflow approach can classify what parameters are more critical for the identified behaviour, and what compensatory or synergistic associations between parameters are imposed by the required behaviours. These relationships can guide simplifications for further mathematical analysis.

Our study provides two major conclusions arguing for the exploration of multiple high-quality parameter sets (or solutions), which support and extend those shown in Marder et al. [17]. First, to consider a single solution (such as a single set of parameters fitting the expected data) provides limited insights on a given model: *are we sure that the conclusions observed in a single solution (parameter set) are true for all feasible parameter set*s? Secondly, the characterization of the set of feasible parameter sets provides a deeper understanding of the model because it can (a) characterize and enumerate the set of hypotheses that cannot be rejected based on the present experimental and theoretical understanding of the phenomenon; (b) identify specific experiments that can be most informative in distinguishing between these pending alternative scenarios; and (c) provide insights about what parameters of the models are critical, and could be used as targets for specific manipulations in subsequent simulations.

We have made of use of the approach proposed in the present chapter in other computational works [2, 3]; where this parameter exploration procedure was used to confirm that a specific property observed in one model was general to the class of possible models constrained by experimental and behavioural results. See also [5] for other parameter exploration-based approaches to generate hypothesis.

References

1. Achard, P., De Schutter, E.: Complex parameter landscape for a complex neuron model. PLoS Comput. Biol. **2**(7), e94 (2006)
2. Ardid, S., Wang, X.J., Gomez-Cabrero, D., Compte, A.: Reconciling coherent oscillation with modulation of irregular spiking activity in selective attention: gamma-range synchronization between sensory and executive cortical areas. J. Neurosci. Off. J. Soc. Neurosci. **30**(8), 2856–2870 (2010)
3. Cano-Colino, M., Almeida, R., Gomez-Cabrero, D., Artigas, F. Compte, A.: Serotonin regulates performance nonmonotonically in a spatial working memory network. Cereb. Cortex (New York, N.Y.:1991) **24**(9), 2449–2463 (2014). doi:10.1093/cercor/bht096
4. Cedersund, G., Samuelsson, O., Ball, G., Tegnér, J., Gomez-Cabrero, D.: Optimization in biology parameter estimation and the associated optimization problem. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)

5. Cedersund, G.: Prediction uncertainty estimation despite unidentifiability: an overview of recent developments. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
6. Compte, A., Brunel, N., Goldman-Rakic, P.S., Wang, X.J.: Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb. Cortex **10**(9), 910–923 (2000)
7. Funahashi, S., Bruce, C.J., Goldman-Rakic, P.S.: Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J. Neurophysiol. **61**(2), 331–349 (1989)
8. Gomez-Cabrero, D., Compte, A., Tegnér, J.: Workflow for generating competing hypothesis from models with parameter uncertainty. Interf. Focus **1**(3), 438–449 (2011)
9. Hodgkin, A.L.: Chance and Design in Electrophysiology: An informal account of certain experiments on nerve carried out between 1934 and 1952. J. Physiol. **263**(I), 1–21 (1976)
10. Hodgkin, A.L., Huxley, A.F.: Currents carried by sodium and potassium ions through the membrane of the giant axon of Loligo. J. Physiol. **116**, 449–472 (1952)
11. Hug, S., Schmidl, D., Bo Li, W., Greiter, M.B., Theis, F.J.: Bayesian model selection methods and their application to biological ODE systems. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
12. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: IEEE International Conference on Neural Networks Proceedings, Vols. 1–6, 1942–1948 (1995)
13. Kirk, P., Silk, D., Stumpf, M.P.H.: Reverse engineering under uncertainty. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
14. Lejon, A., Samaey, G.: Stochastic modeling and simulation methods for biological processes: overview. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
15. Mannakee, B.K., Ragsdale, A.P., Transtrum, M.K., Gutenkunst, R.N.: Sloppiness and the geometry of parameter space. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
16. Marder, E., Taylor, A.L.: Multiple models to capture the variability in biological neurons and networks. Nat. Neurosci. **14**(2), 133–138 (2011)
17. Marder, E., Tobin, A.E., Grashow, R.: How tightly tuned are network parameters? Insight from computational and experimental studies in small rhythmic motor networks. Progr. Brain Res. **165**, 193–200 (2007)
18. Munkres, J.R.: Topology, 2nd edn. Prentice Hall, New Jersey (2000)
19. Poli, R.: Analysis of the publications on the applications of particle swarm optimisation. J. Artif. Evol. Appl. **2008**, 1–10 (2008). doi:10.1155/2008/685175
20. Prinz, A.A., Billimoria, C.P., Marder, E.: Alternative to hand-tuning conductance based models: construction and analysis of databases of model neurons. J. Neurophysiol. **90**(6), 3998–4015 (2003)
21. Prinz, A.A., Bucher, D., Marder, E.: Similar network activity from disparate circuit parameters. Nat. Neurosci. **7**(12), 1345–1352 (2004)
22. Sunnåker, M., Stelling, J.: Model extension and model selection. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
23. Tucker, W.: Interval methods. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)
24. Tuckwell, H.: Introduction to Theoretical Neurobiology (2 vols.). Cambridge Studies in Mathematical Biology. Cambrige University Press, Cambrige (1988)
25. Van Geit, W., Achard, P., De Schutter, E.: Neurofitter: a parameter tuning package for a wide range of electrophysiological neuron models. Front. Neuroinform. **1**, 1 (2007)
26. Van Geit, W., De Schutter, E., Achard, P.: Automated neuron model optimization techniques: a review. Biol. Cybern. **99**(4–5), 241–251 (2008)
27. Van Schepdael, A., Carlier, A., Geris, L.: Sensitivity analysis by design of experiments. In: Uncertainty in Biology, A Computational Modeling Approach. Springer, Chem (2016, this volume)