Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 00 (2020) 000–000

Procedia Computer Science

www.elsevier.com/locate/procedia

24th International Conference on Knowledge-Based and Intelligent Information & Engineering Systems

Decision making model based on attractor network with binary neurons

Marcin Penconek*

Quantitative Psychology and Economics
Faculty of Psychology, University of Warsaw, Stawki 5/7, 00-183 Warsaw, Poland

Abstract

Recent neuroscience research provides evidence on how decision making systems can be implemented in cortical networks of human brains. In this paper we investigate the possibility of developing a biophysically-realistic decision making model based on the asynchronous recurrent network with attractors and binary neurons. Success of our approach suggests that decision making is a network phenomenon and biologically plausible performance of the model does not rely on neuron activation dynamics of the leaky integrate-and-fire neurons. Results presented here provide supplementary evidence that the postulates formulated by recent empirical research are indeed essential for the implementation of decision making system into neuronal networks.

Keywords: decision making models; recurrent neural networks; attractor networks; binary neurons; accuracy; reaction time

1. Introduction

Decision making models have been of interest to psychology and neuroscience since the development of the drift diffusion model (Ratcliff, 1978) and the race model (LaBerge, 1962, Pike, 1966, and Vickers, 1970). Recent advances in neuroscience provide insights related to biophysically-realistic decision making models and how they

* Corresponding author. Tel.: +48-661-960-033
E-mail address: m.penconek2@uw.edu.pl

can be implemented in the brain. In his book on emotions and decision making, Rolls presents evidence in favor of the neuronal model based on an attractor network with recurrent connections, and formulates four postulates related to its features:

- strong recurrent connections in attractor pools,
- global feedback inhibition facilitating winner-take-all competition,
- stochasticity due to finite-size effects and random Poisson inputs,
- long synaptic time-constant (NMDA) in the leaky integrate-and-fire neurons (Rolls, 2014).

A model consistent with the above principles based on the leaky integrate-and-fire neurons (Tuckwell, 1988) was developed by Wang (2002, Wong and Wang, 2006). The model consist of a recurrent network composed of two populations of excitatory neurons selective to stimulus, a third pool of excitatory interneurons not selective to the stimulus and the pool of inhibitory neurons. This architecture was previously used by Amit and Brunel (1997).

The Wang model heavily relies on the leaky integrate-and-fire neurons and the incorporation of slow NMDA-mediated synapses in the network connections which are thought to play role in integrating neuronal activity during decision formation. Thus, it is unclear whether representing neuron activation dynamics in the model is essential to its performance. Our aim is to investigate this question by constructing the model with McCulloch-Pitts (1943) binary neurons. Phenomenological approach presented here suggests that decision making is a network phenomenon depended on the existence of attractors sets in the network rather than any specific properties of neurons themselves.

Apart from using binary neurons, we also reduce the network architecture of the model to excitatory neurons only. The underlying assumption is that the role of inhibitory neurons relies on providing the global feedback inhibition mechanism. The mechanism that we introduce in our model ensures persistent chaotic behavior of the network, and as a result, creates stochastic inputs into the attractor pools. It also facilitates winner-take-all competition postulated by Rolls. Our model shows biophysically-realistic performance and its predictions are directionally consistent with empirical evidence on accuracy and reaction time.

This article is organized as follows. First, we describe the way our model is designed and outline underlying phenomenological assumptions. Then, we discuss model implementation and its parameters. In the results section, we demonstrate that the model predictions are directionally consistent with empirical evidence. Appendix shows a sample of model simulations.

2. Model design

Our model is based on the classical ideas of the asynchronous recurrent artificial neural network (Hopfield, 1982) with McCulloch-Pitts (1943) binary neurons. Its specific construction is inspired by the Wang model and the postulates of Rolls.

The network is formally described as directed graph $G = \langle V, E \rangle$ where V is the set of binary neurons and E is the set of connections between them. Let N be the number of neurons in V, i.e., $|V| = N$ ($N = 1000$). The set of edges E is defined by the function on pairs of neurons $w: V \times V \rightarrow \{0,1\}$. The connectivity weights w_{ij} between the pair of neurons v_i and v_j are assumed to be either 0 or 1, i.e., $w_{ij} \in \{0,1\}$ for all v_i and v_j in V. Thus, unlike in the Wang model which uses weighted connections, in our model the neurons are either connected (if $w_{ij} = 1$) or not (if $w_{ij} = 0$). We will follow the convention that v_i is the postsynaptic neuron and v_j is the presynaptic one (if they are connected). Unlike in the Hopfield network, no restrictions on the connections are imposed. In particular, the network is not assumed to be symmetric (i.e., w_{ij} might differ from w_{ji}), as network symmetry is not biophysically realistic. We shall also allow $w_{ii} = 1$ for some v_i in V. In our model network connections are randomly generated. No learning mechanism is introduced, i.e., the set of connections is fixed and stays the same during network evolution.

Evolution of the network is described by the state function $s_t: V \rightarrow \{0,1\}$ which attributes values 0 or 1 to all neurons. We say that the neuron is active (or spiking) at the time t, if $s_t(v_i) = 1$. For simplicity we shall use the abbreviation $s_t(i)$ for all $i \in N$ when referring to the value of the state function for neuron v_i. All neurons are assumed to be excitatory neurons. The role of inhibitory neurons is captured by the global inhibitory mechanism.

Updates of the state function are governed by the following rule:

$$s_t(\mathrm{i}) = 1 \text{ iff } \frac{\sum_{j=1}^{N} s_t(\mathrm{j})\, w_{ij}}{\sum_{j=1}^{N} w_{ij}} > \frac{1}{\Theta} \cdot \left(\frac{\sum_{j=1}^{N} s_t(\mathrm{j})}{N} \right)^2$$

where the threshold activation level is the same for all neurons and depends on the current rate of active neurons in the system and the inhibitory constant Θ (Θ = 0.13). The threshold rule forms the global inhibitory mechanism as postulated by Rolls. Please note that our rule does not include any random noise effect. Stochasticity as postulated by Rolls is produced by the network itself.

The network evolves continuously over time (t > 0) with discrete-time asynchronous updates of the state function. Model simulations are run with I = 100000 updates. The order of the updates is governed by time delays assigned to neurons. When the neuron v_i is updated at time t, a random time delay of t(i) is assigned to the neuron. The time delay t(i) is assumed to have the exponential distribution, i.e., t(i) ~ Exp(λ), where the rate constant λ (λ > 0) depends on the state s_t(i) of the neuron v_i after the update at time t:

$$t(\mathrm{i}) \sim \begin{cases} \mathrm{Exp}(\lambda_1), & s_t(\mathrm{i}) = 1 \\ \mathrm{Exp}(\lambda_2), & s_t(\mathrm{i}) = 0 \end{cases}$$

The next update of neuron v_i will happen at time t_i = t + t(i). Thus, the state of the neuron stays the same over the interval period of [t, t_i). In case of active neurons, this period is supposed to cover the full cycle of a spiking neuron including depolarization, repolarization and the refractory stage. The rate for a spiking neuron λ_1 = 0.07 is assumed to be higher than the rate for an inactive one (λ_2 = 0.005), This assumption implies that active neurons are updated faster. The implementation of the continuous evolution of the network over time (t > 0) relies on identifying the neuron v_i with the lowest value of t_i and initiating the update of this neuron. The network is initiated with random state s_0(i) and corresponding random delays t_0(i) assuming the rate of active neurons is equal to Θ.

Now, let us define the specific architecture of the network which forms the decision making system. Decisions made by the system are represented by attractor sets A and B. A and B are two disjoint subsets of V, i.e., A ∩ B = Ø, of size n (n=100). Attractor sets exhibit the higher density of connections in comparison to the other parts of the network. Higher density of connections could be a result of hebbian learning (Hebb, 1949), however, no learning mechanism is introduced in our model. Two distinct density constants d_1 = 0.55 and d_2 = 0.36 ($d_1 > d_2$) determine the probability of connection between two neurons depending on whether they both belong to the defined decision set or not.

$$P(w_{ij} = 1) = \begin{cases} \mathrm{d}_1, & if\ v_i, v_j \in \mathrm{A} \text{ or } if\ v_i, v_j \in \mathrm{B} \\ \mathrm{d}_2, & \text{otherwise} \end{cases}$$

Please note that we allow sets A and B to be directly connected as well with the probability of connecting a neuron from A with a neuron from B at the level of d_2.

Stimulus (input) to the attractor sets is provided by two separate pools of neurons S_A and S_B. The density of the connections between neurons from the attractor set A (or B) and the neurons from S_A (or S_B) is set to be d_1, the same as the density of connections within attractor sets. The resulting connections are added to the set of connections governing updates for neurons in A and B, however they do not affect the threshold activation level. The underlying assumption is that the inhibitory mechanism is applicable to a specific spatial location only and the neurons from S_A and S_B are located in a different brain area where reward value (or evidence in favor of decisions) is computed. Stimulus is provided during a predefined period of time by setting all neurons in the selected pool (S_A or S_B or both) active. The level of stimulus is controlled by the sizes of pools S_A and S_B, and the duration of time during which the stimulus is provided. S_A and S_B can be thought of as evidence towards decisions A and B, respectively.

We would say the system made a decision in favor of A (or converged to A), if the rate of active neurons in A in comparison to the rate of active neurons in B exceed a predefined threshold level Ψ (Ψ = 0.75) for a defined period of time T (T = 500ms). Formally, A wins at t_0, iff

$$\sum_{i \in A} s_t(v_i) - \sum_{j \in B} s_t(v_j) > \Psi$$

$$\int_{t_0}^{h} \left(\sum_{i \in A} s_t(v_i) - \sum_{j \in B} s_t(v_j) \right) dt > h\Psi \ \ for\ all\ h \in (t_0, t_0 + T)$$

and t_0 is the minimal t such that the above conditions are satisfied.

The way our model is designed reflects the following phenomenological assumptions about how the cortical network of human (and animal) brain works:

- Single neurons do not perform decision making computations. Computation is a network phenomenon and the result of its architecture.
- Excitatory neurons take part in the computation, while the role of inhibitory neurons relies on providing feedback inhibition which can be simulated through a global mechanism. The mechanism proposed in our model requires that inhibitory neurons are capable of gathering information about the rate of currently active excitatory neurons in spatial area where decision making system is located.
- Brain computes reward values for (evidence in favor of) alternative decisions separately (in other brain area) and decision making is a separate system capable of comparing them. This assumption is consistent with the neurophysiological evidence (Rolls, 2014). Reward values are encoded in the number of active neurons.
- Stochasticity is a result of the proposed inhibitory mechanism that ensures persistently chaotic behavior in spontaneous state (not presented here) and creates random inputs to the attractor sets. In our model we purposely withdraw from implementing on-going random effects other than those resulting from governing updates with the exponential distribution.
- Neurons perform their operations in parallel and their updating time is independent from updating time of other neurons. This is reflected in the model design based on asynchronous network.

Note that with many regards our model is different than the Wang model. Apart of using binary neurons, our model is also based on binary connections (weighted connections in the Wang model). The core of network architecture is the same, however, we abandon the pool of inhibitory neurons. Its role is replaced by the global inhibition mechanism. Unlike the Wang model, our model does not require a baseline stimulation, as the network creates continuous stochastic behavior. Stimulus is provided by additional pools of active neurons (while Wang stimulates the system with frequency inputs). Finally, we reduce on-going random effects to setting random time for updates.

3. Model parameters and implementation

Some model parameters were set arbitrarily, including the size of the network N= 1000 and sizes of the attractor sets n=100, as well as the level of inhibitory constant Θ = 0.13. Note that in our model Θ should be in the range from 0.1 to 0.2. The lower limit ensures that (at least) one attractor can win, the upper limit protects the system to converge to both attractors at the same time. In model simulations, $\lambda_1 = 0.07$. This translates to measuring time in milliseconds (ms) assuming the frequency of spikes of an active neuron is 70Hz. The assumption is consistent with the firing frequency rate of LIP neurons in the human brain (Rolls, 2014). In model simulations λ_2 is equal to 0.005. Note that this does not predetermine the frequency of spontaneous spikes of neurons, as after inactivity period a neuron can become inactive again. The parameter λ_2 was set experimentally. Lower values of λ_2 in relation to λ_1 facilitate slower integration of neuronal activity during decision formation. Density constants d_1 and d_2 are equal to 0.55 and 0.36, respectively, and they were set experimentally. Although high density is not biologically plausible

(Lefort et al., 2009), we decided to simulate our model with these parameters to account for small network size. Relationship between density constants heavily affects model performance. We shall discuss it later.

Despite limited number of parameters, the model shows complex dynamics. Understanding the impact of model parameters on these dynamics requires further research and is beyond the scope of our study.

The model was implemented in R version 3.6.1 (2019-07-05). Program code is available upon request. Computational complexity of the algorithm governing the network evolution is $O(N \cdot I)$. Network connections are regenerated for each simulation in order to ensure validity of results in a general case. Regenerating network connections requires $O(N^2)$ operations.

4. Results

Throughout this section we shall present evidence that the model can provide a seemingly realistic performance in terms of accuracy of decision making and reaction time. In order to test the model, we have run 300 simulations with randomly chosen conflicting inputs provided for a period of 500ms. The results are shown in Fig 3. Red dots represent model decision in favor of A and blue dots show simulations where the model converged to decision B. Dots with no filling indicate model runs with no decision taken.

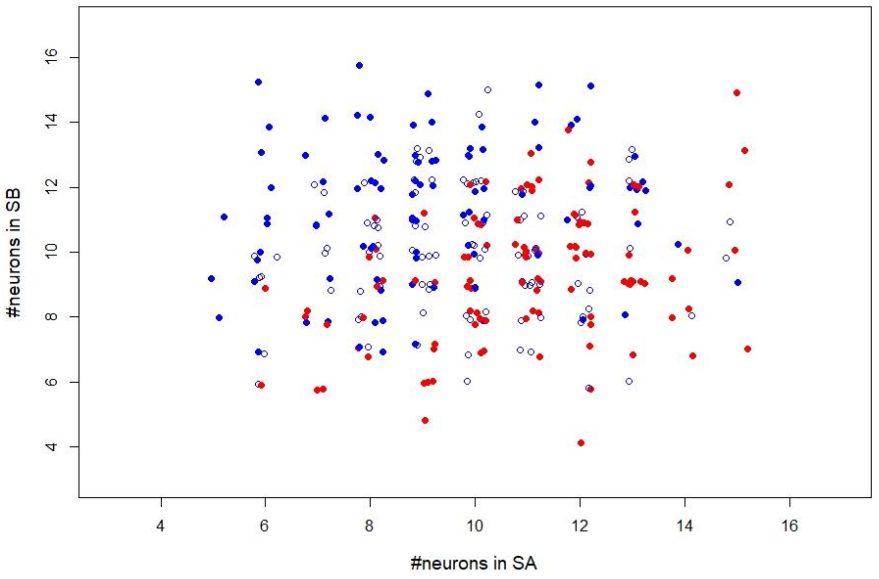

Fig. 1. Decisions is favor of A (red dots) and B (blue dots) in regenerated networks based on 300 models runs with randomly selected conflicting input levels (binomial distribution 0 to 20 active neurons); stimulus duration 500ms. Dots with no filling indicate cases when no decision was made. Model inputs consisting of natural numbers are presented as randomized numbers (in a close neighborhood) in order to show multiple decisions for the same levels of stimulus.

High accuracy level is a desired property of any decision making system. Empirical research shows positive relationship between stronger evidence and decision accuracy. In our model the rate of correct decisions (i.e., in favor of stronger inputs) increases with the difference between conflicting stimulus (one-way fixed effects ANOVA, $F = 43.86$, $df = 1$, $p < 2e\text{-}10$). Difference of one active neuron (between sets S_A and S_B) results in 58% of decisions in favor of the stronger stimulus (33 vs. 24), differences of two and three neurons translate to 74% (31 vs. 11) and 92% (23 vs. 2) correct decisions, respectively.

Model runs show that the same stimulus can result in different decisions (Fig. 1), even in cases where the difference in stimulus level is substantial. A compelling argument in favor of evolutionary benefits of probabilistic decision making is provided by Glimcher. First, decision randomness is essential in uncertain environments. Moreover, optimal behavioral strategies often rely on the implementation of mixed-strategies with probabilities other than zero-one in games with other players (Glimcher, 2004).

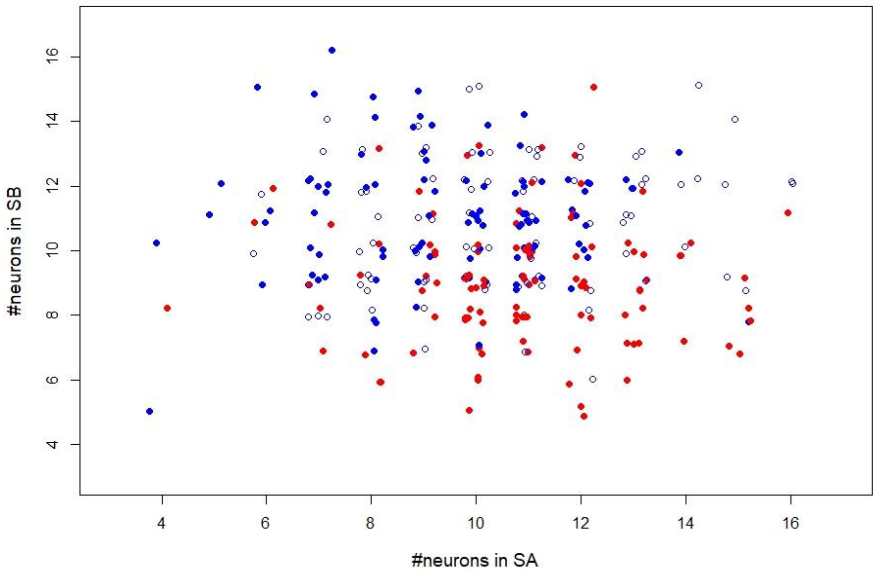

Fig. 2. Decisions is favor of A (red dots) and B (blue dots) in the same randomly generated network based on 300 models runs with randomly selected conflicting input levels (binomial distribution 0 to 20 active neurons); stimulus duration 500ms. Dots with no filling indicate cases when no decision was made. Model inputs consisting of natural numbers are presented as randomized numbers (in a close neighborhood) in order to show multiple decisions for the same levels of stimulus.

Network connections were regenerated for each simulation, so one can suspect that different decisions are a result of model computations in different networks. In order to check it, we conducted Monte Carlo analysis of decisions while running all simulations in the same randomly generated network. This time we have also fixed the connections between decision sets and sets S_A and S_B which provide inputs to the system. The same network still produces different decisions as response to the same non-random stimulus (Fig. 2) suggesting that network evolution over time is the main source of stochastic decisions.

We shall now investigate model predictions for reaction time subject to level and duration of inputs. Reaction time (RT) of the system is measured as the delay between the time when the stimulus started till the time the decision was made. It is expected that higher stimulus level and its longer duration would produce shorter reaction time for correct decisions.

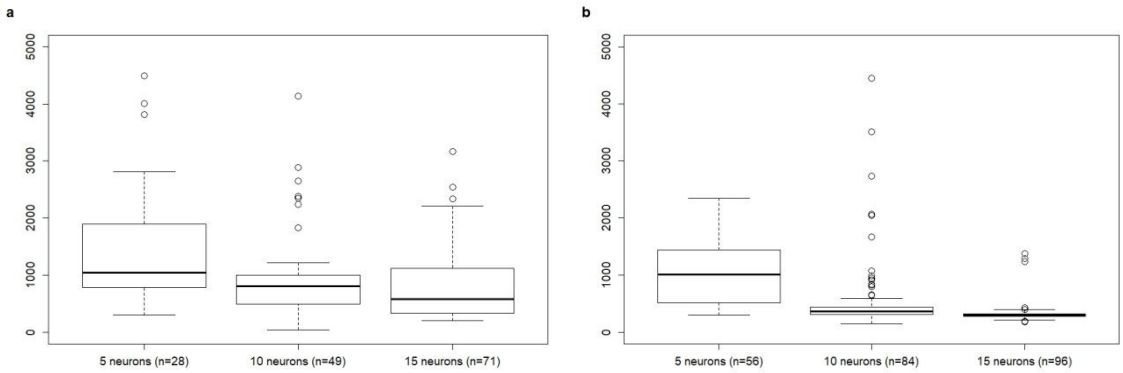

Fig. 3. Reaction time (ms) for model runs with three input levels of 5, 10 and 15 active neurons; a) stimulus duration 300ms; b) stimulus duration 500ms; simulations with correct decisions were selected from 100 runs per group.

In order to test whether the model correctly predicts differences in RT, we conducted a sample of 100 simulations for each condition defined by three stimulus levels of 5, 10 and 15 active neurons and two duration times 300ms and 500ms. Our results confirm the hypothesis (Fig. 3). Reaction time for correct decisions depends on stimulus level (two-way fixed effects ANOVA, F = 52.873, df = 1, p < 3e-12) and stimulus duration (F = 36.407, df = 1, p < 4e-09); interactions between factors are not significant (F = 0.987, df = 1, p = 0.321).

Empirical research provides evidence for longer reaction time on error trials (Rolls, 2014). It also suggests that reaction time encodes the level of confidence in a decision and this information can be used by other brain systems to evaluate the decision accuracy at a later cognitive stage. Thus, differences in RT on correct and incorrect decisions are an important component of the decision making process. We test the model predictions by analyzing reaction time in model runs with conflicting stimulus. Two runs consisting of 300 simulations each were conducted for inputs provided by 15 and 10 active neurons lasting 300 and 500ms (Fig. 4). For both duration times, the difference in reaction time between correct and error trials is significant on the confidence level of 0.95 (one-sided t-test, t = 2.456, df = 24.44, p < 0.0108 and t = 2.555, df = 14.27, p < 0.0114 for 300ms and 500ms, respectively).

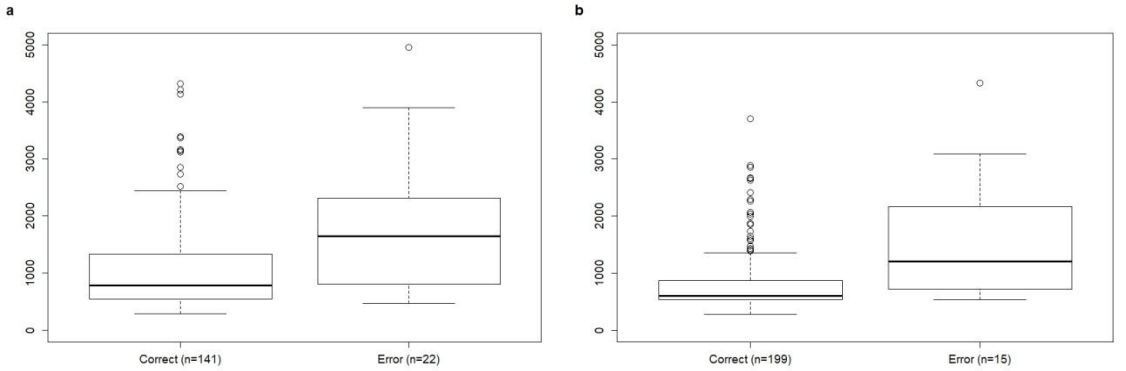

Fig. 4. Difference in reaction time (RT) on correct vs. error trials for 200 model runs with conflicting stimulus of 15 and 10 active neurons; a) stimulus duration 300ms, n=141 correct trials vs. n=22 incorrect; b) stimulus duration 500ms, n= 199 correct trial vs. n=15 incorrect.

5. Discussion

Model simulations were run for density constants $d_1 = 0.55$ (attractor sets) and $d_2 = 0.36$ (elsewhere). These parameters were set based on Monte Carlo analysis consisting of 500 model simulations for various randomly selected pairs of density constants. All runs were with no inputs. Results are shown in Fig. 5a.

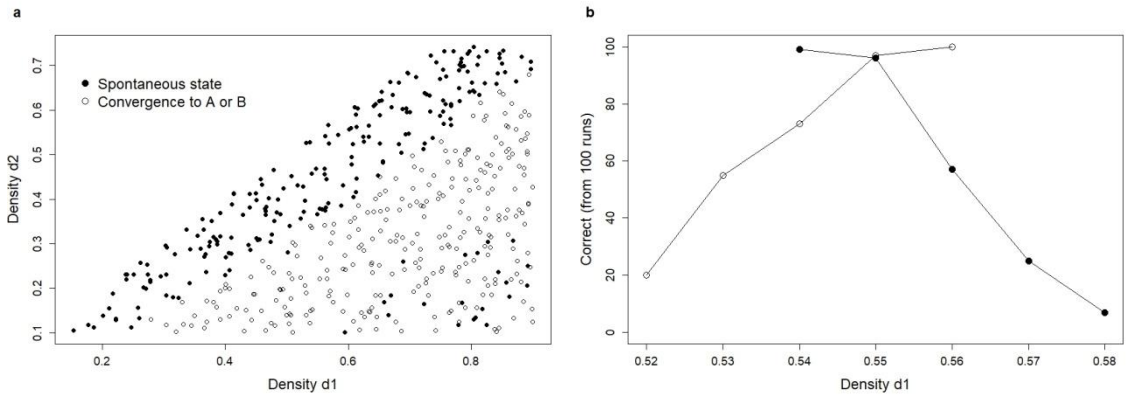

Fig. 5. Decision making in relation to density parameters; a) Monte Carlo analysis based on 500 model runs with no inputs; b) correct responses for varying density parameter d_1 ($d_2 = 0.36$ in all simulations), black dots: spontaneous state with no input, white dots: convergence with input.

Black dots indicate simulations that maintained the spontaneous state (correctly) and white dots represent simulations when the model converged to decision A or B.

The results clearly indicate that the only valid combination of density constants lies on the tiny boarder between the areas of white dots where the model is unstable and black dots where the model quickly becomes insensitive to stimulus. In order to observe how tiny this space is, we analyzed model runs with no inputs and runs with the input provided by 15 active neurons for 500ms. This time we keep density $d_2 = 0.36$ constant and vary the values of d_1 in the range of 0.52 to 0.58 (100 runs per condition). Model response was considered correct, if the system maintained spontaneous state for runs with no input and converged to (any) decision for runs with stimulus (Fig 5b). Both, stability and sensitivity to inputs is only achieved for $d_1 = 0.55$. Other values of the parameter either produce high levels of instability or make the model irresponsive.

6. Conclusions

It might come as a surprise that a decision making model with very limited assumptions can provide a seemingly biologically plausible performance. Further validation of the model in comparison with experimental data is necessary, but results presented in this paper are promising. Our approach of constructing the model with binary neurons suggests that decision making is a network phenomenon resulting from the existence of attractors sets in the network rather than a result of the activation dynamics described by the leaky integrate-and-fire model. It also provides supplementary evidence that the network postulates formulated by Rolls (2014) are indeed essential for the implementation of decision making system into neuronal networks. We expect that our model will help developing new hypotheses on the role of structural network components in cortical computation of decisions. Thanks to its simplicity, the model can also be applied in psychology and neuroeconomics for investigating neurophysiological determinants of human decision making.

Apart from further validation, our study suggests various directions of potential future research. The proposed model works only under very specific choice of the density of connections in attractor sets and the rest of the network. This does not seem to be biologically plausible feature. Hence, it could be interesting to understand what network mechanisms can potentially enhance the range of density parameters under which the network performs. A broader understanding how different assumptions and model parameters impact the model itself and the patterns of decision making could be useful. In this paper, we have decided to present the model with two attractor sets. However, we realize that the model works also in a general case with multiple attractors. It might be interesting to investigate whether the model shows linear relationship between reaction times and the logarithm of the number of alternatives as predicted by the Hick's law (Hick 1952).

Acknowledgements

The study was conducted under the supervision of prof. Dominika Maison. Author would like to express his gratitude for her comments and suggestions. Previous versions of this article were reviewed by prof. Katarzyna Kopczewska and reviewers representing KES International. Author thanks for their feedback and comments.

Appendix A. Model simulations

Examples of model runs are presented in Fig. 6. Charts show the proportions of active neurons in each pool over time. Red line represents the attractor set A, blue line stands for set B and the black line shows the proportion of active neurons in the rest of the network. Stimulus, if applied, is shown through vertical dotted lines indicating the start and end of providing inputs. If a decision was made, solid vertical line represents decision time consistent with the decision rule described above. Time is measured in milliseconds (ms).

The model shows integration of neuronal activity during decision formation after stimulus expires and keeps the decision (in a working memory) as a semi-stable state. When converging to a decision, a strong inhibitory reaction is observed for neurons in the dominated pool. This effect is facilitated by the proposed inhibitory mechanism and the value of inhibitory constant $\Theta = 0.13$ which does not permit both attractor sets to converge at the same time. A weaker level of similar inhibitory reaction is also observed in the rest of the network (Fig. 6c and d).

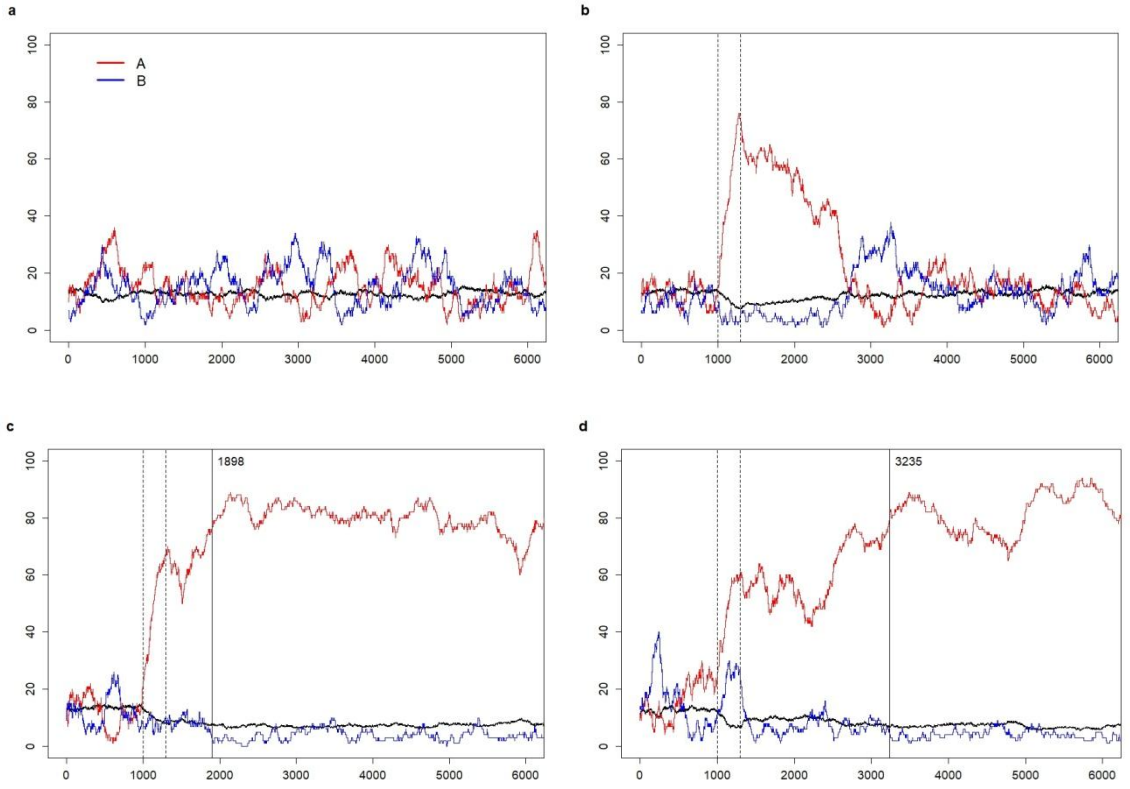

Fig. 6. Model runs under different conditions: a) no inputs, model maintained the spontaneous state; b) stimulus provided to pool A (red line) by 10 active neurons for 300ms (in between dotted lines), model maintained the spontaneous state; c) stimulus provided to pool A (red line) by 10 active neurons for 300ms, model converged to decision A, the solid vertical line indicates decision time; d) conflicting stimulus provided by 10 and 5 active neurons to pools A and B, respectively, model converged to decision A.

Frequency data can be recovered from our model by calculating the number of neuron updates which resulted in its activation in a given period of time. The analysis was conducted slicing time by 10ms periods.

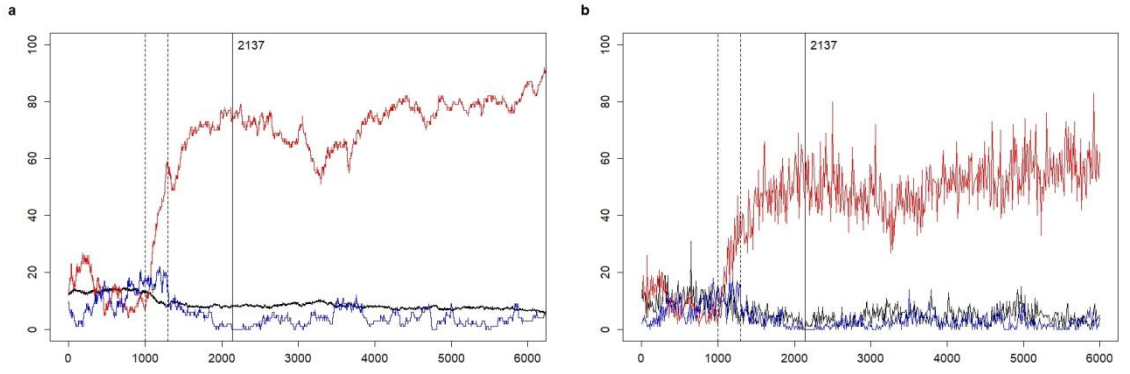

Fig. 7. The same model simulation with stimulus provided by 10 active neurons for 300ms represented by: a) proportion of active neurons in each pool; b) average frequency of spiking neurons in pools A, B and 100 neurons from the rest of the network.

The same model run represented by population of active neurons and by average frequency of spiking neurons per pool (A, B and selected 100 neurons from the rest of the network) is illustrated in Fig. 7a and b, respectively. In the attractor state active neurons fire with the frequency rate in the range of 40-80Hz on average.

References

[1] Amit, D.J., and Brunel, N. (1997) "Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. " *Cereb. Cortex* 7: 237–252.

[2] Hebb, D.O. (1949) "The Organization of Behavior: A Neuropsychological Theory." *New York, NY: John Wiley & Sons.*

[3] Hick, W.E. (1952) "On the rate of gain of information." *Quarterly Journal of Experimental Psychology* **4**: 11–26.

[4] Hopfield, J.J. (1982) "Neural networks and physical systems with emergent collective computational abilities." *Proceedings of the National Academy of Sciences of the USA* **79** (**8**): 2554–2558.

[5] Glimcher, P. W. (2004) "Decisions, Uncertainty, and the Brain. The Science of Neuroeconomics." *Cambridge: MIT Press*

[6] LaBerge, D. A. (1962) "A recruitment theory of simple behavior." *Psychometrika* **27**: 375–396.

[7] Lefort, S., Tomm, C., Sarria, J.C.F. and Petersen, C.C.H. (2009) "The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex." *Neuron* **61**: 301–316.

[8] Luce, R.D. (1986) "Response Time: Their Role in Inferring Elementary Mental Organization" *New York: Oxford University Press.*

[9] McCulloch, W.S.; Pitts, W.H. (1943) "A logical calculus of the ideas immanent in nervous activity." *Bulletin of Mathematical Biophysics* **5** (**4**): 115–133.

[10] Pike, A. R. (1966) "Stochastic models of choice behavior: Response probabilities and latencies of finite Markov chain systems." *British Journal of Mathematical and Statistical Psychology* **21**: 161–182.

[11] Ratcliff, R. (1978) "A theory of memory retrieval." *Psychol. Rev.* **85**: 59–108.

[12] Rolls, E. T. (2014) "Emotion and Decision-Making Explained." *Oxford University Press.*

[13] Tuckwell, H.C. (1988) "Introduction to Theoretical Neurobiology" *Cambridge: Cambridge University Press*.

[14] Vickers, D. (1970) "Evidence for an accumulator model of psychophysical discrimination." *Ergonomics* **13**: 37–58.

[15] Wang, X.-J. (2002) "Probabilistic Decision Making by Slow Reverberation in Cortical Circuits." *Neuron* **36** (**5**): 955-968.

[16] Wong, K.-F. and Wang, X.-J. (2006) "A recurrent network mechanism of time integration in perceptual decisions." *Journal of Neuroscience* **26**: 1314–1328.