Spectra of the Spike Flow Graphs of Recurrent Neural Networks

Filip Piekniewski

Faculty of Mathematics and Computer Science,
Nicolaus Copernicus University, 87-100 Torun, Poland
philip@mat.umk.pl
http://www.mat.umk.pl/~philip

Abstract. Recently the notion of power law networks in the context of neural networks has gathered considerable attention. Some empirical results show that functional correlation networks in human subjects solving certain tasks form power law graphs with exponent approaching ≈ 2. The mechanisms leading to such a connectivity are still obscure, nevertheless there are sizable efforts to provide theoretical models that would include neural specific properties. One such model is the so called *spike flow model* in which every unit may contain arbitrary amount of charge, which can later be exchanged under stochastic dynamics. It has been shown that under certain natural assumptions about the Hamiltonian the large-scale behavior of the *spike flow model* admits an accurate description in terms of a *winner-take-all* type dynamics. This can be used to show that the resulting graph of charge transfers, referred to as the *spike flow graph* in the sequel, has scale-free properties with power law exponent $\gamma = 2$. In this paper we analyze the spectra of the spike flow graphs with respect to previous theoretical results based on the simplified *winner-take-all* model. We have found numerical support for certain theoretical predictions and also discovered other spectral properties which require further theoretical investigation.

Key words: power law network; spike flow model; graphs spectrum

1 Introduction

Power law networks (often referred to as scale-free networks, which sometimes causes confusion [1]) are now an established field of study in random graph theory. Diverse empirical evidence have shown that power law connectivity emerges spontaneously in miscellaneous systems ranging from the World Wide Web [2], science collaboration networks [3], citation networks [4], ecological networks [5], linguistic networks [6], cellular metabolic networks [7, 8] to telephone call network [9, 10] and many others. In many cases networks featuring power law degree distributions also include certain structural properties which enhance tolerance against attacks or bandwidth (the correspondence between power law degree distribution and structural properties is not straightforward and has been discussed in [1]). It is quite natural to ask whether neural systems could benefit from such an architecture, and if so, whether there are any mechanisms inherent to neural activity that might lead to a power law connectivity. Early studies of C.elegans worm nervous system showed exponential decay of degree distribution [11, 12], however the network of C. elegans is very small (the whole organism has only about 1000 cells, and the total number of neurons is just above 300) whereas the mechanisms of self-organization leading to a power law structure might emerge in larger populations of neurons with significant feedback wirings. One strong empirical evidence that this might be the case are the results of [13, 14], which show that the network composed of centers of activity observed in human brain by FMRI, connected whenever their activity is correlated[1] above a certain threshold is scale-free with power law exponent ≈ 2.

It is worth noting that power law graphs are usually sparse (in the sense that the number of edges depends linearly on the number of vertices) and yet well connected (power law graphs are more likely to form a giant component than corresponding – in terms of edge density – Erdős- Rényi random graphs - see chapter 6 in [15] for related study). These features seem to be advantageous for recurrent neural networks, and indeed some studies [16, 17] have proved that power law architectures are useful for artificial NN. In that case however, the connectivity was not a result of neural activity but was rather imposed as a background for already existing models.

Many of the existing models describing the development of power law networks are stemming from the model of Barabási and Albert [18] based on growth and preferential attachment. This model however does not describe well the situation considered in [13] since growth in this case is very limited. Another reason why Barabási-Albert model is inadequate to the situation is that in its most natural setup it leads[2] to power law exponent $\gamma = 3$ while empirical studies of [13] strongly suggest $\gamma = 2$. In our attempt to provide a more adequate theoretical description we have developed the so called *spike flow model* [19] which essentially resembles a typical Boltzmann machine but has more capacitive space of states and a bit tweaked Hamiltonian (the details are supplied in the next section) . Quite unexpectedly the *spike flow model* turned out to be mathematically tractable (at low enough temperatures), which allowed to establish explicit results [20] on the asymptotic properties of the dynamics and the emergence of a power law charge transfer graph (referred to as the *spike flow graph* in the sequel). Further theoretical research allowed to characterize the spectra of the *spike flow graph* in the asymptotic regime [21]. The study of spectral properties is particularly important to determine, whether the *spike flow model* is an adequate description of the mechanisms leading to power law connectivity in nervous system. The results from [21] impose that a certain kind of power law-like dependence should also be present in the distribution of graph eigenvalues (in section 4 below there is a brief discussion concerning the details). In this paper we provide numerical simulations which support claims of [21] which can be

[1] The patient was asked to perform certain simple tasks during the measurement.

[2] There are ways of reaching exponent 2 with variants of Barabási-Albert model, but they are even less suitable for the phenomena discussed.

regarded as a theoretical foundations of the results presented here. The theory in [21] however, is based on the simplified asymptotic version of the model (described below as well) whereas the presented material is based on the full-blown version of the *spike flow model*. Nevertheless the results show the existence of a spectral regime in which the predicted dependency is present. There are also other features of the spectra which require further theoretical study.

The rest of the paper is organized as follows. In section 2 we briefly describe the spike flow model, its basic properties and theoretical results (subsection 2.1) and motivations for studying spectral characteristics (subsection 2.2). In section 3 we describe the numerical setup of the simulation. In further two sections we provide results of the simulation and conclusions.

2 The Spike flow model

2.1 Basic properties

The model consists of nodes $\sigma_i,\ i = 1 \dots N$. Each node's state is described by a natural number from some fixed interval $[0, M_i]$. In the scope of this paper we assume $M_i = \infty$, that is the state space is unbounded (when $M_i = 1$ on the other hand the model much resembles Hopfield network). The network is built on a complete graph in that there is a connection between each pair of neurons $\sigma_i, \sigma_j,\ i \neq j$, carrying a real-valued weight $w_{ij} \in \mathbb{R}$ satisfying the usual symmetry condition $w_{ij} = w_{ji}$, moreover $w_{ii} := 0$. The values of w_{ij} are drawn independently from the standard Gaussian distribution $\mathcal{N}(0, 1)$ and are assumed to remain fixed in the course of the network dynamics. The model is equipped with the Hamiltonian of the form:

$$\mathcal{H}(\bar{\sigma}) := \frac{1}{2} \sum_{i \neq j} w_{ij} |\sigma_i - \sigma_j| \qquad (1)$$

if $0 \leq \sigma_i \leq M_i,\ i = 1, \dots, N$, and $\mathcal{H}(\bar{\sigma}) = +\infty$ in the other case. Here $\bar{\sigma}$ denotes of the state of the whole system. The dynamics of the network is defined as follows: at each step we randomly choose a pair of neurons (units) $(\sigma_i, \sigma_j),\ i \neq j$, and denote by $\bar{\sigma}^*$ the network configuration resulting from the original configuration $\bar{\sigma}$ by decreasing σ_i by one and increasing σ_j by one, that is to say by *letting a unit charge transfer from* σ_i *to* σ_j, whenever $\sigma_i > 0$ and $\sigma_j < M_j$. Next, if $\mathcal{H}(\bar{\sigma}^*) \leq \mathcal{H}(\bar{\sigma})$ we accept $\bar{\sigma}^*$ as the new configuration of the network whereas if $\mathcal{H}(\bar{\sigma}*) > \mathcal{H}(\bar{\sigma})$ we accept the new configuration $\bar{\sigma}^*$ with probability $\exp(-\beta[\mathcal{H}(\bar{\sigma}^*) - \mathcal{H}(\bar{\sigma})]),\ \beta > 0$, and reject it keeping the original configuration $\bar{\sigma}$ otherwise, with $\beta > 0$ standing for an extra parameter of the dynamics, in the sequel referred to as the inverse temperature conforming to the usual language of statistical mechanics. In the present paper we will assume β fixed and large, that is the system is in low temperature regime and so such "stochastic" jumps are rare.

Note that in this setup positive weights $w_{i,j}$ favor agreement of states σ_i and σ_j, while negative weight favor disagreement. Whenever a unit of charge is exchanged between two nodes that fact is recorded by increasing the counter associated with a corresponding edge. The edges (and nodes) being frequently visited by units of charge are in the focus of our interest. We refer to the resulting weighted[3] graph as to the *spike flow graph.*

In [20] a number of results related to the *spike flow model* have been established:

- In contrast to a seemingly complex dynamics, with high probability there is a unique ground state of the system, in which all the charge is gathered in a unit that maximizes
$$S_i := -\sum_{j\neq i} w_{ij}. \tag{2}$$
referred to as *support* in the sequel. The proof goes by a mixture of rigorous and semi rigorous calculations and has a rather asymptotic character, but is in full agreement with numerical simulations for systems containing between a couple hundreds to a couple of thousands of nodes.
- The system's behavior eventually admits a particularly simple approximation in terms of a kind of winner-take-all dynamics: almost all transfers converge to units of higher support (referred to as the *elite*, while the others referred to as the *bulk*), which then compete in draining charge from each other. That is to say, whenever a pair of units is chosen, the transfer occurs from the unit of lower support to the unit of higher support. Ultimately the unit of maximal support gathers all of the charge and the system freezes in a ground state. This approximation was used in [21] to establish explicit theoretical results on the properties of the spectra of the *spike flow graph.*
- The node degree distribution (where by degree we mean the sum of counters of edges adjacent to a given node[4]) obeys a power law with exponent $\gamma = 2$. The proof is based on the elite/bulk approximation and properties of ordering sequences. Again there is a strong agreement with numerical results

2.2 Spectral properties

The graph's spectrum (in this paper by *graph's spectrum* we mean the set of eigenvalues of the adjacency matrix, not the eigenvalues of combinatorial laplacian which are also studied in the literature, see [22] for comprehensive introduction) is among the most basic characteristic features, yet it provides insights into various properties which are usually faint in the typical analysis. In [21] some basic properties of the spectra of simplified *spike flow model* were investigated. By the simplified version of the model here we mean a model equipped with the asymptotic winner-take-all version of the dynamics, that is each charge transfer occurs according the direction of increasing support (however for the

[3] Weighted by edge counters that are not directly related to $w_{i,j}$ which remain fixed as a background to the energy function.
[4] Since charge transfers are directed, we distinguish in and out degrees, but asymptotically these two are equal in terms of distributions.

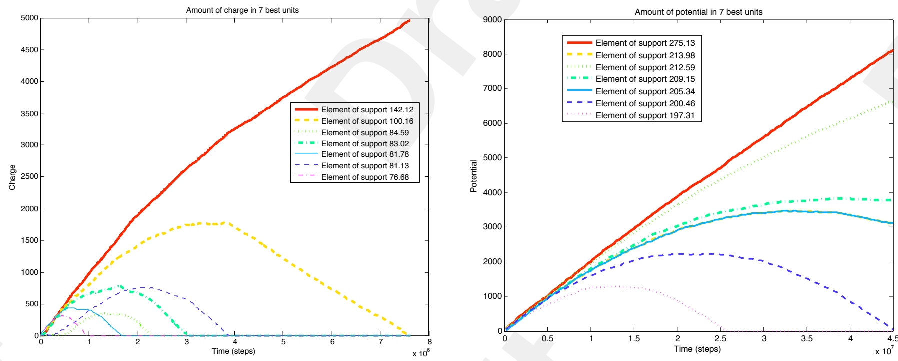

Fig. 1. Typical evolution of the amount of charge in seven units of highest support in the spike flow simulation. The left figure shows the system of 1000 vertices, while the right one of 5000 vertices. At the late stage of simulation when only a few units of highest support contain any charge the winner-take-all dynamics is certainly valid. At this stage the simulation can be simplified according to the restricted dynamics for efficiency. It is not obvious however, at which stage the winner-take-all approximation becomes acceptable. The study of spectral properties might shed some light into such issues.

validity of spectral analysis edge directions were dropped). It is worth noting, that even though such a dynamics becomes reasonably valid in large instances of the *spike flow model* after some number of steps, it is not valid in the early stages of the simulation when there is still a lot of charge in the bulk units. Correspondingly, the resulting *spike flow graph* may be noisy and contain various distortions. Nevertheless the shape of the spectrum depends on global features and should to some extent exhibit the predicted properties. The results of [21] imply that when sorted descending, the k-th eigenvalue behaves like $\frac{C}{k^2}$ for some constant C (note, this paper [21] provides a theoretical background of the results presented here). This result, established by investigating the spectrum of appropriate Hilbert-Schmidt type operators associated to the random evolution in the asymptotic regime, is valid for the simplified *spike flow graph* truncated at both ends by some δ_1 and δ_2 i.e. the nodes of degree less than δ_1 and more than δ_2 are removed from the graph. Recall that removing single vertices from the graph is not easily expressible in terms of graph's eigenvalues and in particular it does not correspond to removal of any particular eigenvalues from the spectrum[5].

3 Numerical setup

The simulations were carried out with the spike flow model consisting of 5000 units in the low temperature regime ($\beta = 100$ which results in extremely rare

[5] The resulting graph has less vertices and consequently there are less eigenvalues in the spectrum, but these remaining eigenvalues correspond to a different adjacency matrix.

transitions against the energy factor). At the beginning of the simulation each unit received 5 units of charge. The simulation was run until all of the available charge ended up in a small fraction of units of maximal support (5 in the case of simulations of 5000 vertices). By then, the winner-take-all approximation becomes perfectly valid (see figure 1), and consequently at the final stage the remaining simulation can be executed with simpler and faster version of the dynamics (winner-take-all) without affecting the resulting *spike flow graph*. The obtained weighted adjacency matrix was symmetrized by adding matrix to its transpose (consequently edge weights in both directions were summed).

4 Results

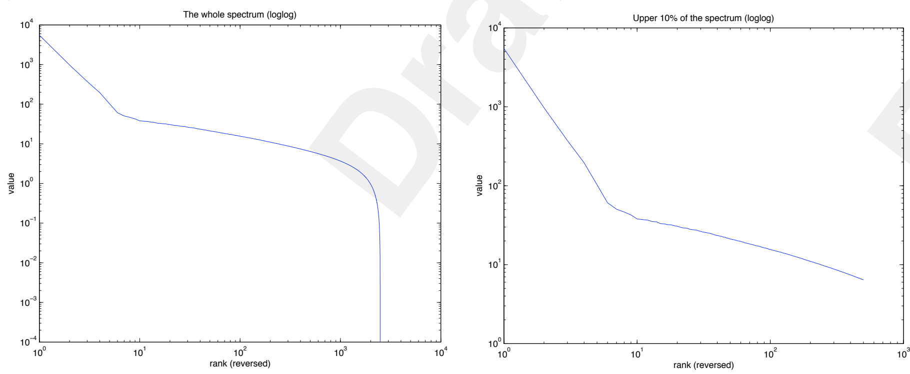

Fig. 2. An example spectrum of a graph consisting of 5000 vertices. The eigenvalues were sorted descending and plotted on a log-log plot. The left figure shows the whole spectrum, the right one only 10% largest eigenvalues. Clearly, the left linear part of the plot behaves like $y = -2x$ which implies that k-th eigenvalue is proportional to k^{-2}. This relation is only visible for the top ten eigenvalues.

In the present study we simulated the setup described in subsection 2.2 by truncating empirical *spike flow graph* at various levels, having in mind however, that the winner-take-all approximation is itself valid for the upper part of the range of vertex degrees. That is, we expect that significant cutoff of low degree vertices should not alter (interesting part of) the spectrum significantly whereas even minor cutoff of high degree vertices might have a devastating effects on the shape of the spectrum[6] (at least at the part which is in the focus of our interest, that is the set of largest eigenvalues). .

The above considerations proved to be true for the investigated model. To make things clearly visible we plotted the spectrum in the descending order on

[6] Consequently, finding out at which point the predicted spectral properties give up provides an insight into how accurate the winner-take-all approximation actually is.

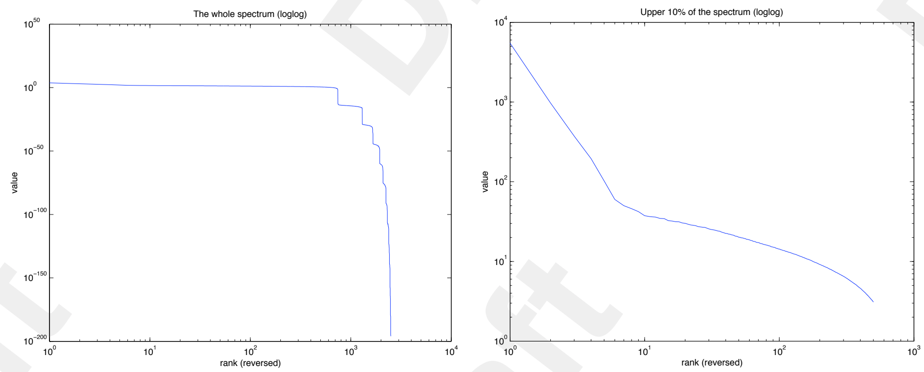

Fig. 3. An example spectrum of a graph consisting of 5000 vertices. The eigenvalues were sorted descending and plotted on a log-log plot. The left figure shows the whole spectrum, the right one only 10% largest eigenvalues. In this case however, 70% of the nodes of lowest degrees were cut off. Nevertheless the left linear part of the plot still behaves like $y = -2x$. The rightmost part of the spectrum (which corresponds to small eigenvalues) exhibits interesting "stair regime", which requires further investigation.

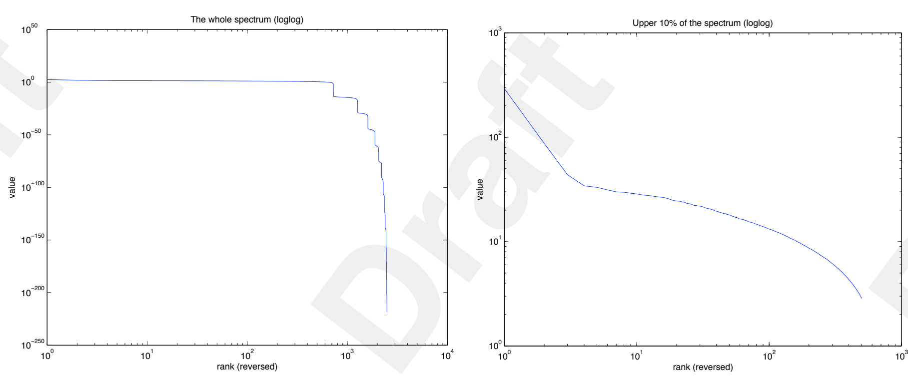

Fig. 4. An example spectrum of a graph consisting of 5000 vertices. The eigenvalues were sorted descending and plotted on a log-log plot. The left figure shows the whole spectrum, the right one only 10% largest eigenvalues. In this case however, 70% of the nodes of lowest degrees and 1% (50) of the nodes of highest degrees were cut off. A couple of top eigenvalues behave like $y = -2x$. The "stair regime" is clearly visible at right of the spectrum.

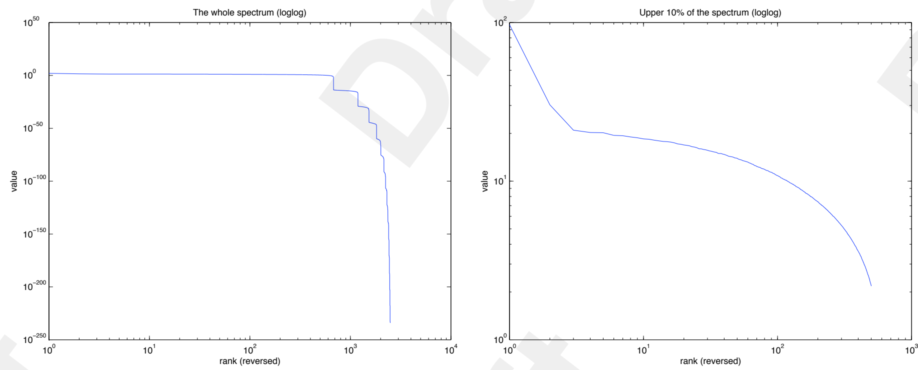

Fig. 5. An example spectrum of a graph consisting of 5000 vertices. The eigenvalues were sorted descending and plotted on a log-log plot. The left figure shows the whole spectrum, the right one only 10% largest eigenvalues. In this case however, 70% of the nodes of lowest degrees and 3% (150) of the nodes of highest degrees were cut off. Here the $y = -2x$ becomes irrelevant, which suggests that the property is related to the elite part of the graph which has been fairly cut off.

a log-log plots (figures 2,3,4,5) with various cutoffs. We expected that the initial (leftmost) part of the spectrum would form a straight line on the log-log plot with slope of ≈ -2. Clearly such a straight line is visible on figure 2 where the spetrum of the full (not truncated) graph is presented. This regime is valid for about 10 largest eigenvalues (it might not seem that significant, but note that these 10 largest eigenvalues contain notable part of the total mass of the spectrum), for smaller eigenvalues the approximation breaks down due to distortions related to more complex dynamics of the full blown model[7]. As expected, the investigated part of the spetrum remains nearly intact (figure 3) after significant cutoff of the low degree vertices (70% of the low degree vertices were removed). Interestingly, the other part of the spectrum (i.e., small eigenvalues) started to exhibit somewhat discrete decay resembling a stairway (there are groups of eigenvalues having nearly same value). The reasons for such a spectral characteristic are yet unclear and require further theoretical explaination. Figure 4 shows the spectrum of the graph, whose 70% of low degree vertices and 1% of high degree vertices were removed. Clearly the initial part of the spetrum in which the straight line approximation is valid had shrunk to about 4 eigenvalues. The removal of 1% of high degree vertices (50 vertices) knocks down fair amount of the elite and consequently the winner-take-all approximation becomes inaccurate. This effect is even more visible in figure 5 where 3% (150) of high degree vertices were removed and the straight line regime is nearly absent, although still the first two eigenvalues seem to follow the expected relation. As mentioned earlier, the largest eigenvalues are not directly related to largest degree vertices

[7] This is not very surprising since the winner-take-all dynamics is certainly not valid for low degree vertices.

and the above result should rather be interpreted in terms of validity of the winner-take-all approximation of the resulting truncated graph.

5 Conclusions

It seems that the theoretical predictions of [21] are to some extent observable in the full blown *spike flow model* equipped with significantly more complex dynamics that the winner-take-all simplification. The results of [21] are themselves of rather asymptotic character, and consequently it was not obvious whether any of the predicted properties would be visible in the simulation of 5000 vertices, where the winner-take-all approximation is only valid for some fraction of the units of high support. Empirical data give insight into more complex spectral properties of the *spike flow model*, notably the *stairway regime* which might be related to the fact, that the truncated version of the graph can become disconnected, and consequently the particular flat regions in the spectrum could be attributed to various connected components, but this requires further investigation.

Acknowledgments. The author gratefully acknowledges the support from the Polish Minister of Scientific Research and Higher Education grant N N201 385234 (2008-2010). The author also appreciates fruitful collaboration with Dr Tomasz Schreiber.

References

1. Li, L., Alderson, D., Tanaka, R., Doyle, J.C., Willinger, W.: Towards a theory of scale-free graphs: Definition, properties, and implications (extended version) (Oct 2005)
2. Albert, R., Jeong, H., Barabási, A.L.: Diameter of the world-wide web. Science **401** (Septmeber 1999) 130–131
3. Barabási, A.L., Jeong, H., Néda, Z., Ravasz, E., Schubert, A., Vicsek, T.: Evolution of the social network of scientific collaborations. Physica A **311**(4) (2002) 590–614
4. Redner, S.: How popular is your paper? an empirical study of the citation distribution. European Physical Journal B **4**(2) (1998) 131–134
5. Montoya, J.M., V., R.V.S.: Small world patterns in food webs. Journal of Theoretical Biology **214**(3) (February 7 2002) 405–412
6. i Cancho, R.F., Solé, R.V.: The small-world of human language. Proceedings of the Royal Society of London B **268**(1482) (November 7 2001) 2261–2265
7. Bhalla, U.S., Iyengar, R.: Emergent properties of networks of biological signaling pathways. Science **283**(5400) (January 1999) 381–387
8. Jeong, H., Tombor, B., Albert, R., Oltvai, Z.N., Barabási, A.L.: The large-scale organization of metabolic networks. Nature **407**(6804) (October 5 2000) 651–653
9. Abello, J., Buchsbaum, A., Westbrook, J.: A functional approach to external graph algorithms. ESA '98: Proceedings of the 6th Annual European Symposium on Algorithms (1998) 332–343

10. Aiello, W., Chung, F., Lu, L.: A random graph model for massive graphs. In: STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing, New York, NY, USA, ACM (2000) 171–180
11. Amaral, L.A., Scala, A., Barthelemy, M., Stanley, H.E.: Classes of small-world networks. Proc Natl Acad Sci U S A **97**(21) (October 2000) 11149–11152
12. Koch, C., Laurent, G.: Complexity and the Nervous System. Science **284**(5411) (1999) 96–98
13. Eguíluz, V.M., Chialvo, D.R., Cecchi, G.A., Baliki, M., Apkarian, A.V.: Scale-free brain functional networks. Phys Rev Lett **94**(1) (January 2005)
14. Sporns, O., Chialvo, D.R., Kaiser, M., Hilgetag, C.C.: Organization, development and function of complex brain networks. Trends Cogn Sci **8**(9) (September 2004) 418–425
15. Chung, F., Lu, L.: Complex Graphs and Networks (Cbms Regional Conference Series in Mathematics). American Mathematical Society, Boston, MA, USA (2006)
16. Perotti, J.I., Tamarit, F.A., Cannas, S.A.: A scale-free neural network for modelling neurogenesis. Physica A Statistical Mechanics and its Applications **371** (November 2006) 71–75
17. Stauffer, D., Aharony, A., da Fontoura Costa, L., Adler, J.: Efficient hopfield pattern recognition on a scale-free neural network. The european physical journal B (32) (2003) 395–399
18. Barabási, A.L., Albert, R.: Emergence of scaling in random networks. Science (286) (October 1999) 509–512
19. Piękniewski, F., Schreiber, T.: Emergence of scale-free spike flow graphs in recurrent neural networks. In: Proc. IEEE Symposium Series in Computational Intelligence - Foundations of Computational Intelligence, Honolulu, Hawai USA (April 2007) 357–362
20. Piękniewski, F., Schreiber, T.: Spontaneous scale-free structure of spike flow graphs in recurrent neural networks. Neural Networks **21**(10) (2008) 1530 – 1536
21. Schreiber, T.: Spectra of winner-take-all stochastic neural networks. In review, available at arXiv, http://arxiv.org/abs/0810.3193 (2008)
22. Chung, F.R.K.: Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92) (Cbms Regional Conference Series in Mathematics). American Mathematical Society (February 1997)