The impact of connectivity on the memory capacity and the retrieval dynamics of Hopfield-type networks

Christophe L. Labiouse[*], Albert A. Salah[**], Irina Starikova[***]

[*]Department of Cognitive Sciences, B33
University of Liege, 4000 Liege, Belgium
clabiouse@ulg.ac.be

[**]Perceptual Intelligence Laboratory
Boğazici University, 80815 Istanbul, Turkey
salah@boun.edu.tr

[***]Department of Philosophy
Novosibirsk State University, 630090 Novosibirsk, Russia
shi@sibmail.ru

Most models of neural associative memory have used networks with broad connectivity. However, this seems unrealistic from a neuroanatomical perspective. A simple model of associative memory with emergent properties was introduced by Hopfield [5]. We choose this widely known model to investigate the impact of connectivity on the storage capacity and the retrieval dynamics in artificial associative networks. In this paper, we use sparse topologies where only a small fraction of possible connections were actually active. For these diluted topologies, we test different kinds of architecture; namely, randomly connected networks, regularly connected networks, small-world networks, and modular networks. Computer experiments reveal that, among these diluted topologies, the modular architectures, that exhibit a relatively high clustering coefficient and whose characteristic path lengths and total physical wiring lengths are short, produce the far better results in terms of memory storage and generalization abilities. Theoretical implications and extensions to this work are discussed.

Keywords: Neural Networks, Associative Memories, Graph Theory, Small-World Networks

1. Introduction

Neural networks of associative memory were intensively studied for the past twenty years [4, 5], and a lot of computer experiments, theoretical results and formal analyses were performed on them [1, 2, 6, 8-11]. However, these models used to be fully connected which is not a realistic structure of real biological associative memories in the brain. Available evidence suggests a sparse connectivity in the neocortex and in the hippocampus [12]. Moreover, the patterns of connectivity seem to be better captured by structured topologies and not by randomly connected graphs. In fact, some studies show that both the neocortex and hippocampal regions present modular architectures or, at least, a structured connectivity [7, 12]. In this paper, we want to investigate the impact of different diluted topologies on classical measures of networks' performance; namely, the storage and the generalization capacities, the resistance to noise. In particular, we look at randomly connected, regularly connected, 'small-world', and modular networks. Each of these topologies will be described in details in the network model section. Our basic results show that the modular architectures yield the best results on the tested criteria and exhibit a relatively high clustering coefficient, a short characteristic path length and a short total physical wiring length. First, we present the background ideas underlying this work. Then we detail the different architectures we create and the methods to generate them. We explain the computer experiments performed and expose the results. We conclude by a discussion on the theoretical implications and on the potential future directions to follow.

2. Background and Motivation

The Hopfield model was designed to solve an important computational problem: the associative memory problem [4, 5]. The aim is to store a set of patterns in such a way that when presented with

a new pattern, the network responds by producing one of the stored patterns that is the most similar to this new pattern. In other words, the goal is to develop a computational device that could act as a content-addressable memory, to complete patterns, to generalize correctly, and to be quite resistant to noise. For twenty years now, an amazing number of theoretical and empirical studies have dissected this model. Although this model offers an unrealistic and poor model of human memory, its simplicity offer insights to build new and more efficient models. However, the impact of structured connectivity has never been systematically investigated with the help of graph theory concepts. In the fifties, two Hungarian mathematicians, Erdös and Renyi have written seminal papers on random graphs [3]. Since then, random graphs have grown as a nearly independent field and have generated a massive literature. However, most works considered that graphs were either perfectly regular or totally random. Recently, some researchers [13] have shed lights on the nature of artificial and natural networks, revealing a structure far different from the supposed regular or random graphs. Then, many complex systems were inspected with this brand new flexible and expressive graph formalism, assuming that structure and organization might help to understand function, process, emergence, and eventually, give insightful views on the nature's workings. The graph structure models entities and their relations to each other, properties of those relations, and allows higher-level properties to emerge as functions of the whole graph. Duncan and Watts showed that a particular set of high-level properties exist in many real systems that can be modeled with specific random graphs (small-worlds). Capturing these properties allows the models to come closer to the system under study, and some of the unrealistic assumptions that were historically attributed to the problem can be discarded. Our purpose is to discard one assumption in the Hopfield-type networks; the assumption that the underlying network topology is fully connected, and regular. The associative models are claimed to be biologically plausible, and it is desirable to reduce the artificial assumptions made to achieve good results. Our hypothesis is that connectivities similar to the actual brain connectivity work just as well, and we investigate it in a comparative way.

Small-world networks are characterized by two mathematical values: the characteristic path length L and the clustering coefficient C. The former is the typical average measure of separation between every pair of nodes in a graph G. L is defined as:

$$L(G) = \frac{1}{N(N-1)} \sum_{i \neq j \in G} d_{ij}$$

where d_{ij} is the shortest path length between i and j, i.e. the minimum number of edges traversed to get from a vertex i to a vertex j. The latter is a local property of G measuring the average cliquishness of a node. If G_i defines the subgraph of the neighbors of a vertex i, C is defined as:

$$C(G) = \frac{1}{N} \sum_{i \in G} C_i$$

where C_i is the cliquishness of vertex i, and is defined as:

$$C_i = \frac{\#\text{ of edges in } G_i}{\frac{k_i(k_i-1)}{2}}$$

Small-world networks are in fact highly clustered, as regular lattices, yet having small characteristic path lengths, like random graphs.

3. Network Models

In order to investigate the effect of connectivity profiles on the properties of associative memories, we create networks of N neurons which accept bipolar states (either +1 or –1). The different neural

interconnectivities we define are specified by an undirected graph G on the vertices $[N]\times[N]$ where a connection from neuron i to neuron j exists if $\{i,j\} \in G$. The training is made on P randomly generated unbiased patterns of length N and the connection strength between node i and node j is computed by a Hebbian learning rule:

$$w_{ij} = \frac{1}{N}\sum_{\mu=1}^{P} \zeta_i^{\mu}\zeta_j^{\mu} c_{ij}$$

where

$$c_{ij} = \begin{cases} 0 \text{ if } \{i,j\} \notin G \\ 1 \text{ if } \{i,j\} \in G \end{cases}$$

The state of a neuron i is updated in discrete steps by the following rule:

$$S_i(t+1) = \operatorname{sgn}\left(\sum_{j\neq i}^{N} w_{ij} S_j(t)\right)$$

where sgn( ) is the sign function and $S_j(t)$ is the output activation of neuron j at time t. Updating the activations is done asynchronously by randomly picking a single neuron and evaluating whether its current state has to change or not. This guarantees the convergence of the network into a fixed point attractor (either the correct pattern or spurious attractors) after a sufficient number of iterations.

The basic network topology we use is a one-dimensional ring lattice on which we place the neurons. This allows us to compute the Euclidean inter-neuron distances, and eventually the total wiring length of a given network. We define the degree $k = \alpha N$, where k is the number of neurons which are connected to a given neuron, and α is the overall connectivity of the network. For each type of the diluted topologies, we construct different graphs G. We start from a regular ring lattice where each neuron is connected to its k neighbors. This constitutes the regular networks. Then, by keeping constant the overall connectivity α, we create increasingly random graphs using the procedure of Watts and Strogatz [13]. For each vertex, we rewire its edges, with probability p, to a randomly chosen vertex of the network. This progressively introduces long-range connections as p increases. The graphs change from regular connectivity at $p = 0$ to random connectivity at $p = 1$, by steps of 0.1. As noted by Watts and Strogatz, relatively low probabilities (between $p = 0.001$ and $p = 0.1$) are sufficient to produce small-world graphs. Besides the single procedure to produce regular, small-world, and random graphs, we also create modular networks of equal overall connectivity. These networks consist of M modules. Each module consists of n neurons, each of which is connected to all n-1 neurons within the same module. Besides these intra-module connections, inter-module connections are randomly generated to keep constant the overall connectivity α, and to allow some fraction of shortcut connections. These are randomly distributed throughout the entire network. This kind of structure is supposed to mimic neocortical columns and some hippocampal regions (refs). It is important to note that all the sparse topologies we model in this paper do not have power-law degree distributions. Either the degree is constant for each node (in regular lattices) or the degree distribution is Poisson-like (in random, small-world, and modular topologies).

4. Computer Experiments and Results

The networks we create have $N = 100$ neurons, and $P = 5$ memories. The ratio is chosen largely below the theoretical storage limitation of $0.138N$ for fully connected Hopfield networks [4, 10] in order to avoid any catastrophic forgetting of memories. In fact, it was shown that, when a Hopfield network is trained on a proportion of memories above this criterion value, the entire set of memories is not learned and the memory model becomes useless. We choose different amounts of overall connectivity from $\alpha = 0.06$ to $\alpha = 0.50$ by steps of 0.04. For the modular architectures, we

build networks with either $M = 10$, $n = 10$ or $M = 20$, $n = 5$ or $M = 25$, $n = 4$. We compare the performances of fully connected nets and those of our sparse topologies.

To test the networks, we probe them with patterns containing increasing amounts of corrupted bits in order to evaluate the generalization capacities of the networks and to roughly check the sizes of the basins of attraction. The amount of corruption is either 0 % (this tests how perfectly a network is able to recall the patterns, and is called the 'memory' condition) or 20 % (this tests how tolerant to noise a network is, and is called the 'generalization' condition) of the number of neurons. To test the quality of recall, we compute the Hamming distance between the expected attractor (the stored pattern) and the actual attractor reached by the network. A low average Hamming distance indicates a good recall and/or an accurate generalization. On the other hand, a high average distance is a sign of an inability to store patterns and/or a poor capacity to generalize (this might indicate the small size of the basins of attraction). We also record the retrieval dynamics by storing the change of activation states through time. However, we do not report any data related to these dynamics here. Results reported here are averaged on 50 runs performed on each type of simulation. Given the 12 different amounts of overall connectivity, the 15 various architectures (one fully connected, one regular, one random, nine due to the random rewiring procedure, and three modular architectures), and the 2 different amounts of noise present in the probes, we therefore run more than 300 types of simulations. All the procedure (graph creation, learning, and recall) is done in a C++ program we specifically wrote for this project.

In the 'generalization' condition, the Hamming distances as a function of average degree are shown in Figs. 1 and 2. The major results are: First, given the small number of patterns used in these simulations, fully connected networks give, on average, the best results, being able to recover most of the time the right attractors. That is why we do not report their performance on the figures. Second, independently of the structure, the higher is the average degree, the lower is the Hamming distance. In other words, the more connections present in a network, the better is its capacity to generalize. This relationship seems to be nonlinear: even a relatively small average degree, so sparsely connected networks, suffices to produce accurate generalizations. Third, random networks perform better than regular lattices. Fourth, the performance seems to be a nonlinear function of the probability of random rewiring. In fact, rewiring 10 % of the edges gives results that are already different from those exhibited by regular networks. With 30 % of rewiring, the performance is nearly at the same level than the one of random networks. Fifth, modular architectures generalize as good as random networks. The generalization performance as a function of physical wiring length is shown in Fig. 3 and 4. There seems to be no clear differences between all the simulated topologies. Actually, for a same level of overall connectivity, modular and regular architectures have physical wiring lengths far shorter than those of random nets. However, random nets with very short wiring lengths perform far worse than regular and modular nets of equivalent wiring lengths. This is due to the fact that random nets with short wiring lengths are too sparse and are not able to transmit accurately information for an accurate retrieval. So, too much sparseness prevents correct recalls. To sum up, modular architectures seem to offer the better trade-off for the generalization task. Using short wiring lengths, they generalize quite accurately, even with relatively sparse connectivities.

In the 'memory' condition, the Hamming distances as a function of average degree are shown in Fig. 5. The important results are: First, the memory performance is better than the generalization performance for all the topologies. Second, as in the generalization task, independently of the topology, the higher is the average degree; the lower is the Hamming distance. In other words, too sparse connections impede a correct learning of the right attractors. Anyway, there is no need to use a full connectivity to achieve such a correct memorization; a relatively small average degree is sufficient to yield an accurate memorization. Third, random nets are better than regular lattices. Fourth, a small amount of random rewiring is sufficient to give results as good as those of totally

random nets. Fifth, and this is the most important result, the modular architectures outperform random, regular, and small-world nets, especially with very small average degrees. Random, regular, and small-world networks need to have at least an average degree of 25-30 % of the number of neurons to give a reasonable performance. Modular architectures only need average degrees of 5-10 % of the number of neurons. The memory performance as a function of physical wiring length is shown in Fig. 6. This supports the preceding analysis. The modular architectures outperform all the other topologies. Even with incredibly short wiring lengths, they are able to correctly memorize patterns.

5. Discussion

The computer experiments seem to indicate that a relatively small amount of randomness in the rewiring process produces a performance qualitatively similar to the one of random networks. However, contrary to the Bohland et al.'s claims [1], the probability of random rewiring (between 0.25 and 0.30) that gives a performance as good as random nets is far higher than the very low probabilities (between 0.001 and 0.1) needed to generate small-world memories. Actually, with a probability of 0.25, the characteristic path length is obviously decreased but the clustering coefficient dramatically dropped, too. This is not a feature of small-world networks, as stated by Watts and Strogatz. Most importantly, the results clearly indicate the advantage of modular architectures. Their performances are impressive, especially with very low levels of connectivity, and better than random networks, yet exhibiting much shorter physical wiring lengths and a clustered structure virtually absent in random nets. Finally, among the topologies we test in this paper, the modular architectures seem to nicely fit with available neuroanatomical evidence. Four potentially interesting issues seem to be open to pursue this project.

First, instead of examining, by trail and error, some specific topologies, it would be potentially fruitful to see the issue as an optimization problem. We could examine how a network is growing by itself through its interactions with the learning environment. Genetic algorithms seem to offer a reasonable alternative to see whether topologies emerge from the scratch. Moreover, GA techniques could support the results presented here if modular architectures were the most fitted topologies for associative problems.

Second, time might be come to unite trends interesting in those issues. Neuroanatomical research on cortical connectivity might expand its knowledge by borrowing recent concepts from random graph theories and the study of network dynamics. Cognitive science could also take advantage of these different approaches to understand cognitive processes such as concept learning and memory with more biologically plausible models.

Third, modular architectures we model here have high clustering coefficients and short characteristic path lengths as the Watts and Strogatz's small-world networks do. That indicates that the Watts and Strogatz's model might not be the only way to generate small-world nets. Building a classification of different small-world graphs might be a worthy project for the future. In the same vein, modeling architectures with a power-law distribution of degrees might yield insightful ideas about how associative memories actually behave.

Finally, the development of spiking neuron models and integrate-and-fire models to test the effects of connectivities might give insightful results on the way real brains process information. Moreover, given the focus on temporal synchronies in these models, the real physical distances between neurons certainly play a key role in the temporal coding of information. This would support the hypothesis of a bias towards short connections given their reliability and economy to transmit signals.

Acknowledgments

Christophe L. Labiouse is a Research Fellow of the Belgian National Fund of Scientific Research. The authors are grateful to the Santa Fe Institute, the Central European University, and the Eötvös Loránd University in Budapest.

References

[1] J.W. Bohland, A.A. Minai, Efficient associative memory using small-world architecture, *Neurocomp. 38-40* (2001) 489-496.
[2] Canning, E. Gardner, Partially connected models of neural networks, *J. Phys. A 21* (1988) 3275-3284.
[3] P. Erdös, A. Renyi, On random graphs I, *Publ. Math. Debrecen 6* (1959) 290-291.
[4] J. Hertz, A. Krogh, R.G. Palmer, *Introduction to the Theory of Neural Computation*, Addison-Wesley, Redwood City, CA, 1991.
[5] J.J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA 79* (1982) 2445-2458.
[6] R. Huerta, F. Corbacho, L.F. Fago-Fernández, J.A. Sigüenza, Fast response and temporal coherent oscillations in small-world networks, *Phys. Rev. Lett. 84* (2000) 2758-2761.
[7] R.A. Jacobs, M.I. Jordan, Computational consequences of a bias towards short connections, *J. of Cog. Neurosci. 4* (1992) 331-344.
[8] J. Komlós, Effect of connectivity in an associative memory model, *J. Comput. System Sciences 47* (1993) 350-373.
[9] N. Levy, E. Ruppin, Associative memory in a multi-modular network, *Neural Comput. 11* (1999) 1717-1737.
[10] R.J. McEliece, E.C. Posner, E.R. Rodemich, S.S. Venkatesh, The capacity of the Hopfield associative memory, *IEEE Trans. on Information Theory 33* (1987) 461-482.
[11] D. O'Kane, A. Treves, Short and long-range connections in autoassociative memory, *J. Phys. A 25* (1992) 5055-5069.
[12] K.E. Stephan, C.C. Hilgetag, G.A.P.C. Burns, M.A. O'Neill, M.P. Young, R. Kotter, Computational analysis of functional connectivity between areas of primate cerebral cortex, *Philos. Trans. R. Soc. London B 355* (2000) 111-126.
[13] D. Watts, S. Strogatz, Collective dynamics of 'small-world' networks, *Nature 393* (1998) 440-442.

Figures[1]

Fig.1. Hamming distances as a function of average degree with 20 % corrupted bits (generalization task)

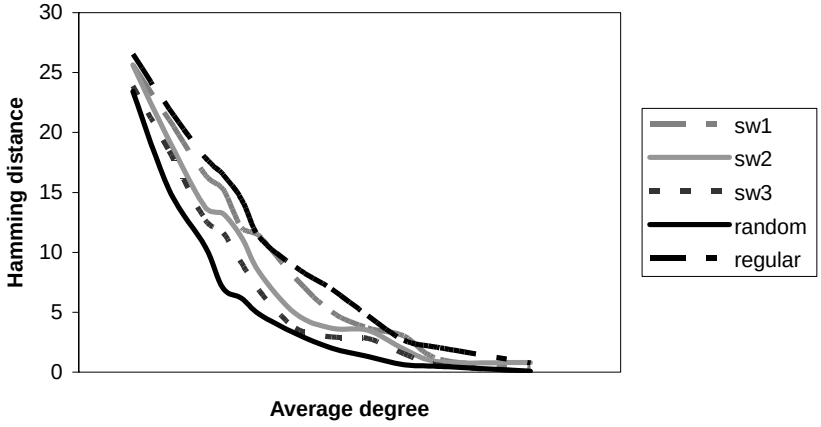

Fig.2. Hamming distances as a function of average degree with 20 % corrupted bits (generalization task)

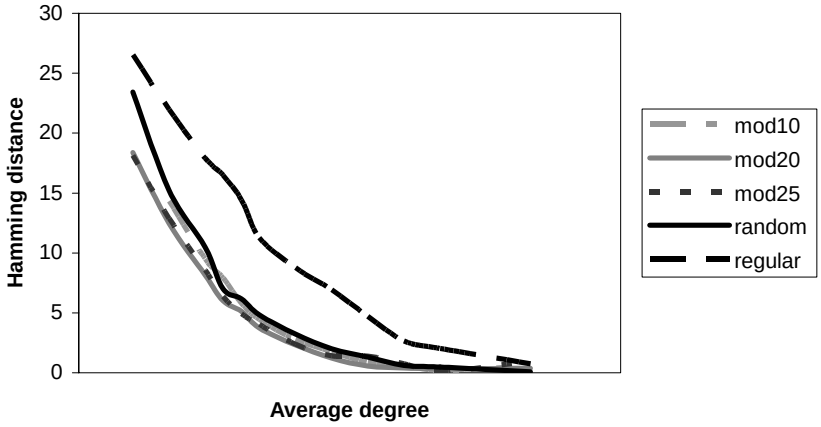

Fig.3. Hamming distances as a function of total wiring length with 20 % corrupted bits (generalization task)

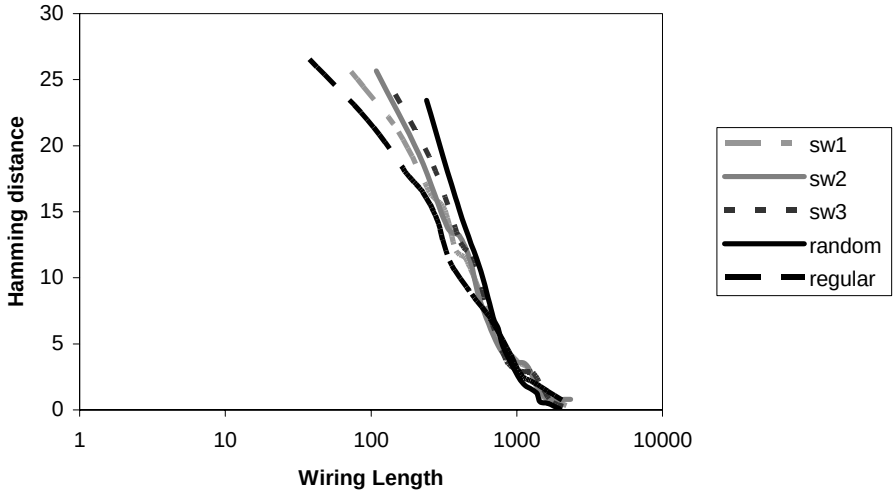

[1] In the figures, "sw1" means 10 % of random rewiring, "sw2" means 20 % of random rewiring, "sw3" means 30 % of random rewiring, "mod10" means 10 modules of 10 neurons, "mod20" means 20 modules of 5 neurons, and "mod25" means 25 modules of 4 neurons.

Fig.4. Hamming distances as a function of total wiring length with 20 % corrupted bits (generalization task)

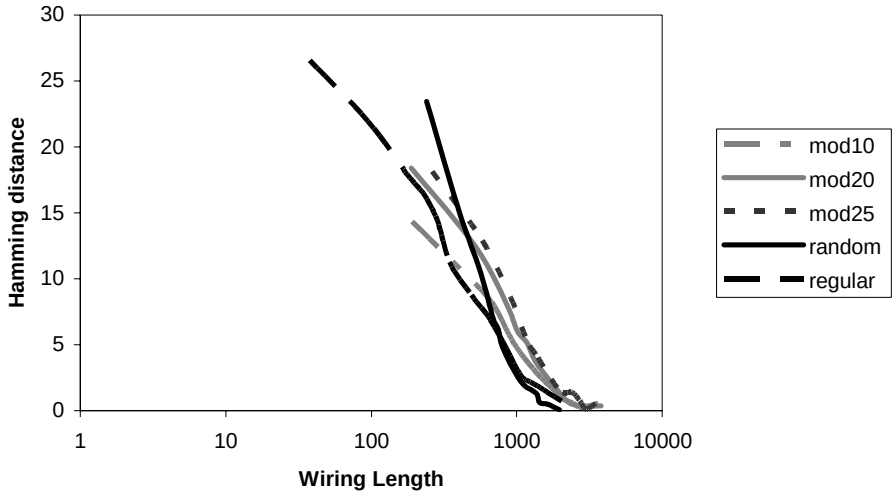

Fig.5. Hamming distances as a function of average degree with no corrupted bits (memory task)

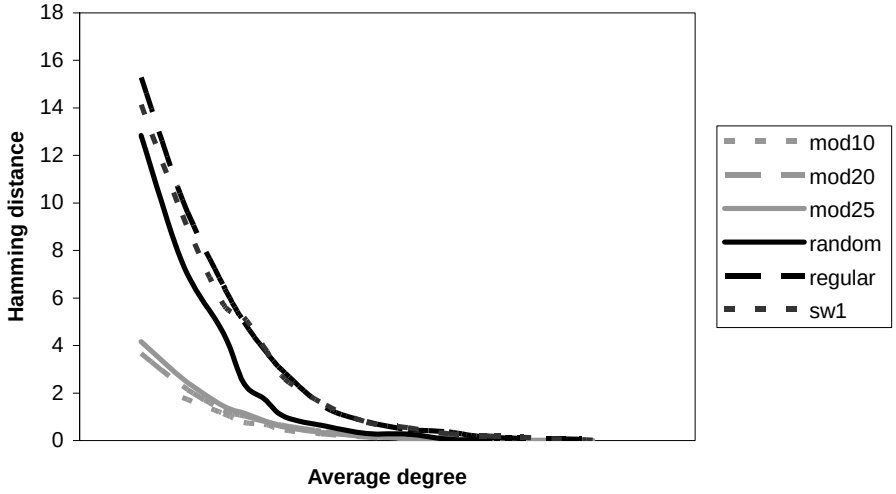

Fig.6. Hamming distances as a function of total wiring length with no corrupted bits (memory task)