Effects of simulated brain damage on small-world neural networks

Kristen Fortney[1*], Juergen Pahle[2], Jose Delgado[3], Gregor Obernosterer[4], Vikas Shah[5], Michael Wojnowicz[6], and Natasha Qaisar[7]

[1]Dept. of Physiology, University of Toronto
[2]Bioinformatics and Computational Biochemistry, EML Research, Heidelberg
[3]Dept. of Psychology, University of Arizona
[4]Institute of Molecular Biotechnology, Austrian Academy of Sciences
[5]Dept. of Internal Medicine, Banner Good Samaritan Medical Center
[6]Dept. of Psychology, Cornell University
[7]Centre of Biotechnology, University of Peshawar

*Send correspondence to k.fortney@utoronto.ca

In the language of graph theory, the brain is a small-world network, as are many other complex systems such as social networks and the Internet. In rough terms, this means that its web of connections lies midway between randomness and the regular pattern seen, for example, in layered backprop networks. Here we demonstrate one advantage of small-world connectivity for systems like the brain: it makes learning and memory more robust to damage. We show with simulations that small-world neural networks are more functionally robust to random removal of synapses than regular networks (and as robust as random networks), in the sense that they remember more of the patterns that they were trained on prior to damage. This advantage of small-world structure increases with the number of layers and synapses of the network.

INTRODUCTION

Small-world networks have become the focus of intense examination over the past decade since the formal description of these networks by Watts and Strogatz (1998). "Small-world" networks can be visualized as regular networks (such as standard feedforward neural networks, where each interior vertex connects to all vertices in directly adjacent layers, see Fig. 1) that have been randomly rewired by moving edges to create a small number of long-range shortcuts. More specifically, small-world networks are typically characterized through two network properties: short average path length L and large clustering coefficient C. The average path length L describes the average, over all pairs of nodes in the graph, of the minimum number of edges that need to be traversed to travel between those two nodes. The clustering coefficient C, (dense local connectivity) reflecting network transitivity, is the probability that nodes B and C are connected to each other if both are connected to A, for any triplet of nodes (A,B,C) in the network. In regular networks, path lengths are long and the clustering coefficient is high - local interactions dominate and very few long range interactions are seen. Watts and

Strogatz reported that small-world transformations caused by rewiring shortcuts leads to a drop in path length but a slower drop in clustering. While such rearrangements ultimately result in a random network (which exhibit short path lengths and low clustering coefficients), at some point in the rearrangement clustering remains relatively high but path lengths are substantially reduced as compared to a regular network. Watts and Strogatz termed these small-world networks; they have turned out to be ubiquitous in biological and other complex systems.

In particular, networks in the brain exhibit small-world topology (Sporns *et al*. 2004, Bassett and Bullmore 2006). Anatomically, local clustering is maintained within the cortical gray matter, where neurons in the same vertical column, 1 mm in diameter, are 10,000 times more likely to have a connection than distant cortical neurons, yet short path lengths are maintained through the sparse long-range connectivity of the white matter which is wrapped around the prominent folds of the cortex (Laughlin and Sejnowski 2003). Functionally, these same small-world properties have shown up in temporal correlations in fMRI data (Sporns *et al*. 2004) and signal synchronization in high-resolution EEG, even though these data don't always depend in any simple way on anatomical connectivity. According to this work, functional networks exhibit high intra-system connectivity (high clustering) yet the number of steps needed to go from one functional area to another is relatively short (short path length) (Eguiluz *et al.* 2005, Sporns *et al.* 2004).

Though the brain shows small-world structure in several ways, the advantage to the organism remains unclear. A recent paper suggests that learning speed may provide at least part of the explanation: Simard *et al.* (2005) found that small-world neural networks can learn substantially faster than similar-sized regular or random networks. Here we demonstrate another advantage of small-world topology: it makes learning and memory more robust to damage, and so may help to preserve brain function in the face of aging, injury or disease.

METHODS

I.   Defining 'small-world' for a (possibly) disconnected graph

As originally defined by Watts and Strogatz (1998), small-world networks are graphs that combine the dense local connectivity of regular networks with the short average path length of random networks, measured respectively by the quantities C and L. Recently the definition was extended by Latora and Marchiori (2003) to apply to also to disconnected graphs. Their new measures are more appropriate for our study because our procedure for creating small-world and random networks can produce disconnected graphs. For a graph G, Latora and Marchiori define the efficiency $E(G)$ and the local efficiency $E_{loc}(G)$ as follows:

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

where N is the number of nodes in the network and d_{ij} is the length of the shortest path between nodes i and j.

$$E_{loc}(G) = \frac{1}{N}\sum_{i \in G} E(G_i)$$

where G_i is the subgraph of G consisting of the neighbors of i.

Efficiency E is the inverse of the harmonic mean of $\{d_{ij}\}$ and local efficiency E_{loc} is the average efficiency of these subgraphs. Because E_{loc} is proportional to C and E is proportional to $1/L$, small worlds have high values of both E and E_{loc}, *i.e.* they are both globally and locally efficient. The anatomical and functional networks of the brain are also small-world in the sense of Latora and Marchiori (Latora and Marchiori 2004, Achard 2007).

II.   *Neural network generation*

Following Simard *et al.* (2005) we generated three classes of feedforward neural networks: regular, small-world, and random (Fig. 1). Neural networks with regular topology were constructed in the standard way: neurons in layer i were connected to every neuron in layer i -1 and every neuron in layer i +1. From the regular network we created random & small-world networks by this rewiring procedure:

1.   Choose a random connection in the network and delete it.
2.   Randomly choose two unconnected nodes (not in the same layer) and add a forward connection.

Repeating these steps over and over eventually results in a random network, but if instead we stop after some appropriate, smaller number of rewirings then E and E_{loc} will both be high, and the network will have a small-world structure (Fig. 1).

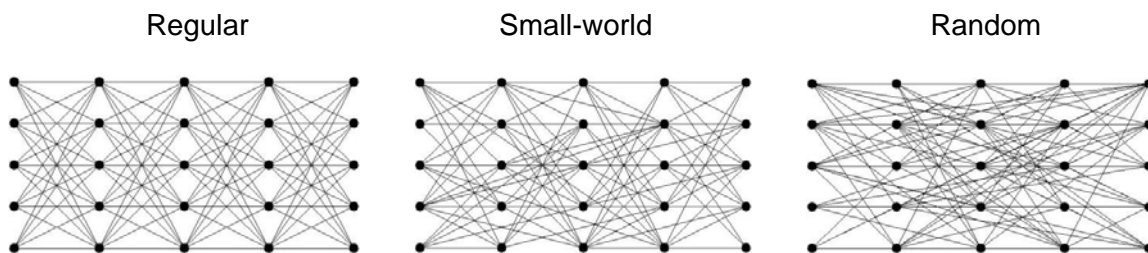Regular   Small-world   Random

Fig. 1. Randomly rewiring a regular network produces a small-world network, and eventually a random network.

For our experiments we used the same five network architectures as Simard *et al.* (2005). The table below, summarizing these networks, is based on their paper.

Layer L	Neurons per layer N/L	Synaptic Connections	Random rewirings needed to generate a small world
5	5	25	19 +/- 5
5	8	40	28 +/- 5
10	10	100	400 +/- 100
15	8	120	830 +/- 50
15	15	225	2300 +/- 300

Table 1. We performed experiments on five different neural network architectures.

III.   Experiments

For each of the five architectures in Table 1, we trained networks using standard learning algorithms — backprop or resilient backprop (RPROP) — to recognize a set of random binary input-output patterns. We used two learning algorithms because backprop frequently became stuck in local minima and never achieved the accuracy we wanted, particularly on networks of more than 10 layers. With RPROP, we were able to include deeper networks.

We varied the *pattern-set size* — number of patterns the networks had to learn — to see whether the difficulty of the task affected the results. For the bigger networks — architectures 3–5 — we used pattern-set sizes of 8, 16, 32, 64, or 128, and for each of these sizes we tested10 regular, 10 random, and 10 small-world nets. For the smaller networks — architectures 1 and 2 — we used pattern-set sizes of 8, 16, or 32, and for each we again tested 10 regular, 10 random, and 10 small-world nets (since these two architectures have only 5 input neurons, they can receive only $2^5 = 32$ distinct binary input vectors). The networks were trained until they achieved the same low rate of error (MSE = 1.0 for architectures 1-3, whether trained by backprop or RPROP; MSE = 1.2 for architectures 4 or 5, which were always trained by RPROP).

We then simulated brain damage by randomly eliminating some percentage of the network's edges ("synapses") — either 5%, 10%, 25%, or 50%. After damage, we tested the networks on the pattern set they had previously been trained on, and recorded their MSE. The damaged networks were not given a chance to relearn before testing. So in total we performed 20 different tests on architectures 3, 4, and 5 (4 damage levels × 5 pattern set sizes), and 12 tests on architectures 1 and 2 (4 damage levels × 3 pattern set sizes). All tests were carried out using SNNS, the Stuttgart Neural Network Simulator, developed in Stuttgart, and JavaNNS, a Java based neural network simulator, developed at the University of Tuebingen.
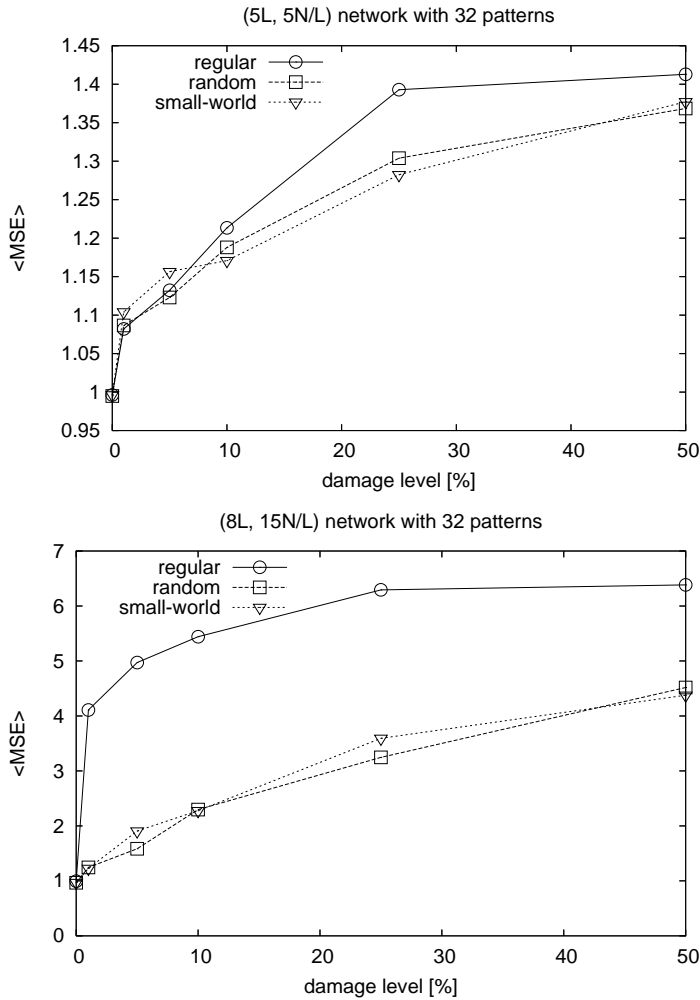
Fig. 2. The beneficial effect of a small-world topology is much more pronounced for deep network architectures (bottom) than for shallow ones (top).

RESULTS

For each network architecture, for each pattern set size (8, 16, 32, 64, or 128), and for each level of damage (5%, 10%, 25%, 50%), we performed two statistical tests, comparing our 10 small-world networks to 10 regular networks, and the same small-world networks to 10 random networks. For each of these two comparisons we used a two-tailed Student t-test (allowing for the possibility of different variances). Then, for each network architecture, we performed a Bonferroni correction. In all our tests, the criterion for significance was a P value < 0.05 after this Bonferroni correction.

I.   *RPROP results*

<u>Small-world nets vs. Regular nets</u>

Our simulations showed that small-world networks were significantly more robust to damage than regular networks.

- In all 20 tests (*i.e.* for every size pattern set and for every damage level) performed on the deep networks — architectures 3–5 — small-world networks were significantly more robust to damage than regular networks.
- In 5 of the 12 tests on second-shallowest networks — architecture 2 — small-world networks were significantly more robust to damage than regular ones.
- For the shallowest networks — architecture 1 — none of the 12 tests showed a significant difference between the regular and small-world.

For any given architecture and level of damage, the difference between the small-world and regular nets grew as difficulty of the learning task increased (Fig. 3).

        These results indicate that small-world networks are more robust to damage than regular networks. The difference becomes more pronounced for networks with more layers or more neurons per layer.
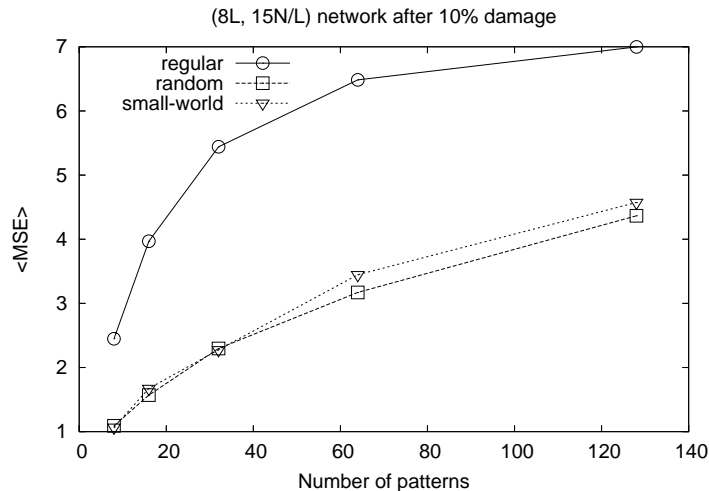

Fig. 3. In a deep network, the difference between the small-world and regular networks grows as the difficulty of the learning task (*i.e.* number of patterns to be learned) increases.

<u>Small-world nets vs. Random nets</u>

We found no consistent difference between the small-world and random networks:

- In all 12 tests performed on the shallow architectures, 1 and 2, and all 20 tests on architecture 4, there was no significant difference between small-world and random networks.
- On 1 of the 20 tests for each of architectures 3 and 5, there was a significant difference, with the random networks outperforming the small-world networks.

These results suggest that there is usually no difference between the functional robustness of small-world versus and random networks, except perhaps for deep nets with many synapses, where random networks may occasionally surpass small worlds.

II.   Backprop results

Results of our backpropagation experiments closely matched those with RPROP, though with backprop we have no satisfactory data for architectures 4 and 5, or architecture 3 with the largest pattern set, because the regular networks never learned to the desired level of accuracy (MSE = 1).

Again, we found no difference between small-world and regular networks, or small-world and random networks, for the shallowest networks (architecture 1). In the (8L, 5N/L) architecture, the small-world nets outperformed the regular nets on 2 of the 20 tests, and the random nets outperformed the small-world ones on 1 test. Finally, in the (8L, 15N/L) architecture, the small worlds outperformed the regular nets on 15 of the 16 tests, and there was no difference between the small-world nets and the random ones.

These results again support the conclusion that small-world networks are more robust to damage than regular ones, and that there is little difference in this regard between random nets and small worlds.

DISCUSSION

Our simulations demonstrate that small-world neural networks are more functionally robust to random removal of synapses than regular networks (and as robust as random networks). This difference in performance seems to be more pronounced for larger and deeper network architectures; presumably, as the number of layers increases, the shortcuts created by the rewiring procedure are able to have a more drastic effect on communication in the network. Another factor that affects performance is the difficulty of the learning task: Fig 3 shows that, as the number of patterns to be learned increases, so does the difference between the small-world and regular networks.

Our results suggest an explanation for why the brain is wired up as a small-world network: it makes our learning less vulnerable to the effects of aging, injury, or disease. Our study is relevant to the process of normal aging and the field of dementia research (Alzheimer's disease, frontotemporal dementia, dementia with Lewy bodies, multi-infarct dementia). By age 90, 35-40% of the neurons in some areas of the brain are lost (Dani et al. 1997). Interestingly, Alzheimer's disease seems to involve a loss of small-world properties (Stam et al. 2007). Our results suggest that small-world properties are

beneficial for healthy aging, and may provide a new perspective from which to approach treatments for dementia.

Having established that small-world neural networks are more robust to damage than regular ones, we would like to follow up with a more detailed ANOVA analysis to determine which factors contribute most to this effect. We plan to test a greater variety of neural network architectures to isolate the different effects of: network depth, total number of neurons, total number of synapses, difficulty of the learning task, and degree of damage. We will also investigate the effect of randomly deleting neurons in our networks, instead of synapses, and model the specific pattern of damage seen in brain diseases like Alzheimer's.

REFERENCES

Achard S, Bullmore E (2007). Efficiency and Cost of Economical Brain Functional Networks. PLoS *Computational Biology* **3:2**, e17 doi:10.1371/journal.pcbi.0030017

Albert R, Jeong H, and Barabasi AL (2000). Error and attack tolerance of complex networks. *Nature* **406**, 378-382.

Bassett DS, Bullmore E (2006). Small-world brain networks. *The Neuroscientist* **12: 6**, 512-523

Dani SU, Hori A, Walter GF (1997). *Principles of neural aging.* Amsterdam: Elsevier Science, p. 264.

Eguiluz VM, Chialvo DR, Cecchi GA, Baliki M, and Apkarian AV (2005). Scale-free brain functional networks. *Physical Review Letters* **94**, 018102.

Latora V, Marchiori M (2003). Economic small-world behavior in weighted networks. *European Physical Journal B* **32:2**, 249-263.

Laughlin SB and Sejnowski TJ (2003). Communication in neural networks. *Science* **5641**, 1870-1874.

Newman M (2003). The structure and function of complex networks. *SIAM Review* **45**, 167-256.

Simard D, Nadeau L, Kröger H (2005). Fastest learning in small-world neural networks. *Physics Letters A* **336:1**, 8-15.

Sporns O, Chialvo DR, Kaiser M, Hilgetag CC (2004). Organization, development and function of complex brain networks. *Trends in Cognitive Sciences* **8:9**,418-425.

Stam C, Jones B, Nolte G, Breakspear M, and Scheltens P (2007). Small-world networks and functional connectivity in Alzheimer¨s Disease. *Cerebral Cortex* **17**, 92-99.

Watts DJ, Strogatz SH (1998). Collective dynamics of 'small-world' networks. *Nature* **393**, 409-10.

SNNS - Stuttgart Neural Network Simulator. http://www-ra.informatik.uni-tuebingen.de/SNNS/ [accessed Aug 15th, 2007]

JavaNNS - Java Neural Network Simulator. http://www-ra.informatik.uni-tuebingen.de/software/JavaNNS/welcome_e.html [accessed Aug 15th, 2007]