

Simulating highly diluted neural networks

Jeferson J Arenzon† and Ney Lemke‡

Instituto de Física, Universidade Federal do Rio Grande do Sul, CP 15051, 91501-970-Porto Alegre-RS, Brazil

Received 14 March 1994

Abstract. The asymmetric diluted Hopfield model is studied through numerical simulations and reliable results are obtained even when the network sizes are not exponentially large. The mean final overlap, the mean convergence time and a quantity that gives information on the relative importance of spurious states are measured as a function of the load capacity α .

Neural network models analogous to spin magnetic systems have been studied in the last few years using several powerful frameworks. Among them, the technique developed by Derrida, Gardner and Zippelius [1] (hereafter DGZ) to study strongly diluted systems has been used in a large variety of models, mainly focusing in the ability of recalling stored patterns that are attractors (stable fixed points) of the network dynamics. The obtained results are exact in the limit where the mean connectivity (C) satisfies $C \ll \ln N$, where N is the size of the network. At first, this condition led to the conclusion that numerical simulations should use network sizes scaling as $\exp(C)$ to satisfy the above condition. However, in this paper we demonstrate that the magnetization can be measured and reliable results may be obtained when large though finite networks, with a finite (but not very small) C , are considered. We evaluate the mean convergence time and also show that spurious states play a role in the dynamics if the network size is finite.

The network, built of N binary spins, $S_i = \pm 1$, is updated in parallel according to the following $T = 0$ dynamics

$$S_i(t+1) = \text{sgn}\left(\sum_{j \neq i} J_{ij} S_j(t)\right) \quad (1)$$

where the connections J_{ij} are given by the rule

$$J_{ij} = C_{ij} \sum_{\mu=1}^P \xi_i^{\mu} \xi_j^{\mu} \quad (2)$$

and $C_{ij} = 1$ for C randomly chosen synapses out of the possible N for each neuron and zero otherwise. During the temporal evolution of the system we monitored the overlap m_{μ} between the state S and the memory ξ^{μ} ,

$$m_{\mu}(t) = \frac{1}{N} \sum_i \xi_i^{\mu} S_i(t). \quad (3)$$

† E-mail address: arenzon@ifl.ufrgs.br

‡ E-mail address: lemke@ifl.ufrgs.br




Figure 1. The asymptotic overlap m versus α for several values of N . The analytical results [1] are shown for $C = 20$ (full circle) and in the limit $C \rightarrow \infty$ (full curve), while the results from numerical simulation are the open symbols. The errors bars are for the $N = 16000$ case.

The parallel dynamics of the diluted Hopfield model when considering initial macroscopic overlap, $m(t = 0)$, with only one of the embedded patterns is ruled by [1]:

$$m(t + 1) = \sum_{n=0}^C \sum_{s=0}^{C(P-1)} \frac{(1 + m)^{C-n} (1 - m)^n}{2^{CP}} \binom{C}{n} \binom{C(P-1)}{s} \text{sgn}(CP - 2n - 2s) . \tag{4}$$

Actually, equation (4) is slightly different from DGZ: it does not present the average on the Poisson-distributed synapses since we are interested in a fixed connectivity. In the limit of $C \rightarrow \infty$ it obeys [1]

$$m(t + 1) = \text{erf} \left[\frac{m(t)}{2\alpha} \right] . \tag{5}$$

In either case the network is able to store P memories ξ_i^μ ($i = 1, \dots, N; \mu = 1, \dots, P$) up to $\alpha_c = P_c/C = 2/\pi$ [1] and the transition from the retrieval phase ($m \neq 0$) to a disordered one ($m = 0$) is continuous as $\alpha \rightarrow \alpha_c$.

In order to perform the numerical simulation, one does not need to store the whole $N \times N$ matrix of connections since most of the elements would be zero. We use two matrices $N \times C$: one to store the neighbours of a given neuron i and the other for the values of each *a priori* non-zero synaptic connection. The averages were taken over 20 sets of memories and P initial states.

The mean final overlap provides information that allows a comparison between results from numerical simulation with those obtained through analytical techniques: the maximum load capacity α_c , the kind of transition that occurs at α_c , the sizes of the basins of attraction of the patterns (or equivalently, the robustness against noise), etc.

In figure 1 we show m versus α for N ranging from 4000 to 16 000 neurons. Both the full curve and the full circles are the solution for $C \rightarrow \infty$ (equation (5)) and for $C = 20$ (equation (4)), respectively. The error bars are shown just for the $N = 16\,000$ case and they are significant only in the region near the transition, as expected. The initial state is chosen as one of the embedded memories, that is, $m(t = 0) = 1$. The predicted α_c is




Figure 2. The overlap $m(t)$ versus the time t for $C = 20$ and several values of N and α . The continuous curves are the analytical results of (4) for $C = 20$.

verified: for increasing values of N the simulation curves tend to attain a non-zero limit if $\alpha < \alpha_c = 2/\pi$, while going down to zero if $\alpha > \alpha_c$. It must be remarked that the results for $C = 20$ and $C \rightarrow \infty$ are quite the same and there is no need to further increase the connectivity. It is also interesting to notice that, different to the simulation of the fully connected case, here there is no remanent magnetization for $\alpha > \alpha_c$: m decreases to zero as N increases. Although the final overlap obtained with equations (4) and (5) agrees with the one obtained in the simulation, the dynamics is re-obtained with reasonable accuracy, as can be seen in figure 2.

Information about the sizes of the basins of attraction may be obtained using initial states with $m(t = 0) < 1$ [2]. It was verified that for any positive initial overlap and large values of N , the system converges towards the specified memory, as predicted analytically.

The mean convergence time [3], that is, the number of parallel updatings needed in order to achieve equilibrium, was measured (see figure 3). For $\alpha > \alpha_c$ the simulation yields an almost constant value for $\langle T \rangle$ that grows as the network increases, which does not happen for low values of α . As $\alpha \rightarrow \alpha_c$ there is a slowing down in the relaxation time. This is different from the fully connected model with sequential updating where the convergence time diverges exponentially [4]; here the data seems to fit quite well with a linear behaviour near α_c , although larger networks and connectivities would be required to rule out other possible behaviours.

Since, during the dynamical evolution, the system may be trapped in states that are not memories (spurious states), it is important to measure their influence in the temporal evolution. In the case of first-order transitions and when the final overlap has a value near unity, the mean convergence time and the dispersion around it provide information on the smoothness of the phase space around the embedded memories [3]. Since here the transition is continuous, another tool must be used (see [5] for a comparison of both cases).

A useful quantity that may provide such information is [6]

$$\mathcal{M} = \langle\langle m_{m_0=1} \rangle\rangle - \langle\langle m_{\text{ris}} \rangle\rangle \tag{6}$$

where $m_{m_0=1}$ is the maximum absolute value of the final overlap, $\max\{|m_\mu|, 1 < \mu < P\}$, when the initial state is equal to one of the embedded memories and m_{ris} when it is chosen at random. When the only attractors in the phase space are the embedded memories, both




Figure 3. The mean convergence time $\langle T \rangle$ versus α obtained through numerical simulation for several values of N and $C = 20$. The error bars are just shown for the $N = 16000$ case.




Figure 4. The value of \mathcal{M} versus α obtained through numerical simulation for several values of N and $C = 20$. For low values of α , $\mathcal{M} = 0$, and the curve's area decreases as N increases.

$m_{m_0=1}$ and m_{ris} will be the same and hence $\mathcal{M} = 0$. On the other hand, when there are attractors other than the patterns with extensive basins around it, $\mathcal{M} \neq 0$. Thus, \mathcal{M} is a quantity proportional to the fraction of initial states that lead to spurious states. For small α one obtains $\mathcal{M} = 0$ (see figure 4), signalling that no spurious states with considerable basins of attraction exist. This is also confirmed by the zero dispersion around the mean convergence time: the phase space is smooth around the attractors (for small α the retrieval overlap with the memories is still near unity and the dispersion around the mean is related with spurious states). As α increases, \mathcal{M} grows almost linearly with α , attains a maximum and decreases again to zero. The location of the maximum depends on N and converges to α_c as the sizes get larger. Also for increasing values of N , the area under the curve decreases: the presence of spurious states seems to be just an artifact of the finite network, disappearing when infinite size is achieved. The fluctuations result from the fact that the

quantities $\langle\langle m_{m_0=1} \rangle\rangle$ and $\langle\langle m_{\text{ris}} \rangle\rangle$ are of the same order and large samples must be used to obtain reasonable averages (the averages were taken over 100 sets of patterns, preventing us from using larger sizes). Thus, for finite N there is an intermediate region in the α -axis where the spurious states play a role in the dynamics of the network. Nevertheless, the basins of attraction of the memories are large and, unless the initial overlap is too close to zero, the system will recall the desired pattern. The spurious states found here are mainly superpositions (symmetric or not) of memories.

To summarize, we showed that reliable simulations may be performed in strongly diluted systems and compared with analytical results even without the use of exponentially large systems. With these results the complete range of values of α and C is covered by simulation techniques and the whole set of analytical results provided by rather distinct methods (replicas, strong dilution, etc) may be tested. Besides this, we evaluate the importance of spurious states in the system evolution showing that for small networks they are prominent.

Acknowledgments

We acknowledge fruitful discussions with R M C de Almeida and F A Tamarit. This work was partially supported by the Brazilian agencies CNPq, FINEP and FAPERGS.

References

- [1] Derrida B, Gardner E and Zippelius A 1987 *Europhys. Lett.* **4** 167
- [2] Forrest B M 1988 *J. Phys. A: Math. Gen.* **21** 245
- [3] Kanter I 1990 *Phys. Rev. A* **40** 2611
Kohring G 1990 *J. Phys. A: Math. Gen.* **23** 2237
- [4] Gosh M, Sen A K, Chakrabarti B K and Kohring G A 1990 *J. Stat. Phys.* **61** 501
Gosh M 1992 *Physica* **181A** 23
- [5] Arenzon J J, de Almeida R M C, Iglesias J R, Penna T J P and de Oliveira P M C 1993 *Physica* **167A** 1
- [6] Arenzon J J and de Almeida R M C 1993 *Phys. Rev. E* **48** 4060