On the Storage Capacity of the Hopfield Model with Biased Patterns

Matthias Löwe
Universität Bielefeld
Fakultät für Mathematik
Postfach 100131
33501 Bielefeld
email: loewe@mathematik.uni-bielefeld.de

Keywords: Hopfield model, neural networks, storage capacity, Markov chains, large deviations

AMS Subject Classification: 82C32, 82B44, 60K35

Abstract

We introduce a form of the Hopfield model that is able to store an increasing number of biased i.i.d. patterns (it is well known that the standard Hopfield model fails to work properly in this context). We show that this new form of the Hopfield model with N neurons can store $\frac{N}{\gamma \log N}$ or αN biased patterns (depending on which notion of storage is used). The quantity γ increases with an increasing bias of the patterns, while α decreases when the bias gets large.

1 Introduction

Neural networks, especially the so–called Hopfield model, have been in the centre of interest in the recent probabilistic and physical literature. Originally introduced by Pastur and Figotin [FP77] as a simplified model of a spin glass the Hopfield model connects the ideas of neural computing with those of statistical mechanics. However, it was its re-interpretation as a neural network by Hopfield [Ho82] that created the interest in the model.

The basic idea behind this model is to choose the information to be stored – which usually is referred to as patterns – as the local minima of an appropriate Hamiltonian (this is going to be made more formal in the next section) on the set $\{1, \ldots, N\}$ which is thought of as the set of neurons. Hence the retrieval dynamics (usually a Monte–Carlo dynamics, possibly at zero temperature) will eventually converge to the stored patterns. Since, according to the other rules of neural networks, this is possible only up to a certain accuracy depending on the number of patterns to be stored, one may introduce the notion of a storage capacity of the Hopfield model as the critical number of patterns (depending on N, of course) up to which the model can successfully reconstruct these patterns.

To be more precise, at zero temperature there are two distinct definitions of storage capacity: The first one introduces the gradient (or zero temperature Monte–Carlo) dynamics on the energy landscape on $\{1, \ldots, N\}$ associated with the Hamiltonian of the system and requires that the patterns are fixed points of this dynamics (and thus minima of the Hamiltonian). The second one going back to a seminal paper of Newman [N88] defines the notion of storage capacity more liberally requiring that a pattern is surrounded by an energy barrier of height proportional to N implying that it is close to a local minimum of the Hamiltonian.

It is well known and first was announced by McEliece et al. [MPRV87] that the first definition leads to a storage capacity of at least $\frac{N}{\gamma \log N}$ (where the possible values of γ range between 2 and 6 and basically depend on whether we are interested in results with high probability or in results for almost all patterns) for the Hopfield model with N neurons and randomly and independently chosen unbiased patterns (also see [M92], [B94], and the recent survey paper by Petritis [P95]). On the other hand the essential progress of Newman's work ([N88]) was that he was able to give a proof for the non-rigorous result of Amit, Gutfreund and Sompolinsky [AGS87] showing that the Hopfield model allows to store even αN patterns (for an $\alpha > 0$ small enough), if small errors are tolerated. The value for α obtained by Newman was $\alpha = 0.056$, which has been recently improved by Loukianova [L94] to $\alpha = 0.071$ by a refined large deviations analysis. The latest result concerning the critical storage capacity is due to Talagrand [T95b],[T96] who has given a bound of $\alpha = 0.08$, obtained via an improvement of Loukianova's idea. Nevertheless, it is still an open question, whether the prediction of $\alpha = 0.14$ in [AGS87] on the basis of the non-rigorous replica method and computer simulations can be mathematically justified.

The other important source of interest in the Hopfield model is that it may serve as a simplified spin glass model, where the complexity of the random environment can be controlled by the number of stored patterns. As a matter of fact, if this