On the Storage Capacity of Hopfield Models with Weakly Correlated 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 analyze the storage capacity of different forms of the Hopfield model with semantically correlated patterns ξ_i^ν (i.e. the patterns are correlated in ν but independent in i). We show that the standard Hopfield model of neural networks with N neurons can store $\frac{N}{\gamma \log N}$ or αN semantically correlated patterns (depending on which notion of storage is used), provided that the correlation comes from a homogeneous Markov chain. The quantity γ increases with an increasing correlation of the patterns, while α decreases with large correlations. This answers the open question whether the standard Hopfield can store any increasing number of correlated patterns at all in the affirmative. Moreover, we show that there is a version of the Hopfield model that can store a number of semantically correlated data that is even larger than the bound for the number of independent patterns that can be stored in the standard Hopfield model provided we require that the patterns are exactly reproduced. Finally we show that this model can store αN patterns (provided Newman's notion of storage capacity is used) with an α not depending on p.

1 Introduction

Neural networks are intensively studied in the recent probabilistic and physical literature. The so–called Hopfield model has an exposed position, since it connects the ideas of neural computing with those of statistical mechanics.

Originally the Hopfield model has been introduced by Pastur and Figotin [FP77] as a simplified model of a spin glass. Five years later it was rediscovered as a model of neural networks by Hopfield [Ho82]. 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. 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, there are two definitions of storage capacity: The first one introduces the gradient 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 less strictly requiring that a pattern is surrounded by an energy barrier of height proportional to N.

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 $\frac{N}{6 \log N}$ for the Hopfield model with N neurons and randomly and independently chosen patterns (also see the recent survey paper by Petritis [P95]). On the other hand Newman [N88] gave 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 found 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 [T95] who has announced a bound of $\alpha = 0.085$, obtained via isoperimetric inequalities. 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.

But its abilities as an auto-associative memory only partly motivated the recent interest in the Hopfield model. The other important feature is that it may serve as a sort of a simplified spin glass model, where the complexity of the interactions between two sites, hence the complexity of the random environment, can be tuned by the number of stored patterns. As a matter of fact, if this number equals one the standard Hopfield model (as compared to the generalized Hopfield models to be introduced in the next section) is (by a simple change of variables) equivalent to the classical Curie-Weiss model of a ferromagnet and thus the interactions are purely deterministic. In contrast to that in the region where the number of patterns