Explorations of Fitness Landscapes of a Hopfield Associative Memory with Random and Evolutionary Walks

Akira Imada
Graduate School of Information Science
Nara Institute of Science and Technology
8916-5 Takayama, Ikoma, Nara, 630-01 Japan
akira-i@is.aist-nara.ac.jp

Keijiro Araki
Graduate School of Information Science
and Electrical Engineering, Kyusyu University
6-1 Kasuga-koen, Kasuga, Fukuoka, 816 Japan
araki@c.csce.kyusyu-u.ac.jp

Abstract

We apply evolutionary computations to Hopfield's neural network model of associative memory. In the model, some of the appropriate configurations of the synaptic weights give the network a function of associative memory. One of our goals is to obtain the distribution of these configurations in the synaptic weight space, together with the information of local optima. That is to say, our aim is to know a geography of a fitness landscape defined on the space. As a step toward this goal, we concentrate on the local optima, in this paper. Hence, we use a walk by the Gaussian mutation to explore the fitness landscape, rather than more effective evolutionary walks, expecting its high probability to be trapped at local optima.

1 Introduction

Associative memory is a dynamical system which has a number of stable states with a domain of attraction around them [1]. If the system starts at any state in the domain, it will converge to the stable state (content addressability). In 1982, Hopfield [2] proposed a fully connected neural network model of associative memory in which information is stored by being distributed among neurons and retrieved by being relaxed dynamically. The dynamical behavior of its neuron states strongly depends on synaptic strength between neurons. The synaptic strength from neuron j to neuron i is called a *weight*, and denoted as w_{ij} in this paper. Hopfield used the Hebbian rule [3] to determine the synaptic weights, while we specify them with using evolutionary computations (see, e.g., [4, 5]).

For the fully connected neural networks with N neurons, all the possible configurations of w_{ij} construct the N^2-dimensional Euclidean space. Hereafter, we call this a (*synaptic*) *weight space*. Every configuration of w_{ij} is represented as a point in the weight space. Fitness values of the points are estimated according to the capability of the networks that are determined by the configurations.

Gardner [6] discussed solutions of the weight configurations in the space, in terms of the volume of the solutions. However, the number and distribution of these solutions in the space are still open problem. Hence, the task of our evolutionary computations is to locate the global and local optima in the space to obtain the information as to the number and distribution of these points.

Our previous results show that as the number of patterns to be stored increases, the task to locate one of the optimal solutions becomes difficult. We conjecture that this phenomenon is due to the increased ruggedness of the fitness surface. As will be described below, the fitness surface constructed on a configuration space is called a *fitness landscape*. In our problem, a capability of a network to store patterns as an associative memory defines the fitness landscape on the synaptic weight space. As a first step to study the ruggedness of the fitness landscape of the Hopfield model, we concentrate, in this paper, on the case that only a few patterns are stored in the network, where the number of local optima is not massively large. And the landscapes are explored by a walk using the Gaussian mutation rather than more effective evolutionary walks, expecting it to be trapped more easily at one of the local optima.

Although there have been a fair amount of analyses of the fitness landscape of the Hopfield

model [7, 8, 9], they are all defined on Boolean hypercube of the configuration of *binary neuron states*. The analysis of the landscape defined on the *continuous weight space* instead has not been explored, as far as we know.

2 Fitness Landscape & Two Evolutionary Walks

The concept of the fitness landscape was first introduced by Wright [10] to study biological evolutionary processes. Since then, this concept has been used not only in evolutionary biology but also in chemistry, physics, computer science and so on. In chemistry, for example, a molecule can be represented as a string with N genes and their k alleles, or equivalently, a sequence of N letters with each letter being chosen from an alphabet of size k [11]. Twenty amino acids ($k = 20$) for proteins or four nucleotides ($k = 4$) for nucleic acids can be considered as examples of the alphabet. The k^N possible combinations of the letters construct a configuration space of the string. Then, for example, the free energy of RNA folding into secondary structures [12], or the ability of peptides to bind to a particular substrate to catalyze a specific reaction [13] is assigned as a fitness value to each point in the configuration space of the strings. We can also find the other example in physics. The Hamiltonian energy of Ising spins defines a fitness landscape on the configuration space of N spins. In this case, each spin takes the value either 1 or -1 ($k = 2$). Bray et al. [7, 8] argued the number and distribution of meta-stable states (local optima) of the Hamiltonian energies.

To explore these fitness landscapes, we need a rule with which a point in the space moves to one of its neighbors. Then consecutive movements of a point to the neighbors form a *walk* on the landscape. Macken et al. [11] used *random point mutation*, that changes a single letter in the string, to specify neighbors of the string. They studied the statistical properties of the landscape of the chemical affinity of antibody for antigen in immune response by sampling points along *"an evolutionary walk"* in which point moves to the *firstly found fitter* neighbor. Weinberger [14] used two different walks: "gradient walk" in which the walker steps to the *best* of its neighbors; and *"random adaptive walk"* in which the next step is chosen *at random* from the set of better neighbors, to investigate the Kaufman's NK landscape which is a model formulated in more general form [15].

The fitness landscapes described above are all constructed on the space of *discrete* configurations of string. What we are concerned, however, is a fitness landscape defined on the *continuous* weight space, in the sense of Gardner. Hence, we use a walk by the Gaussian random mutation where each coordinate of a point in N^2-dimensional weight space, x_i, is given small perturbation using the Gaussian random variables with mean 0 and standard deviation σ, i.e.,

$$x_i^{\text{new}} = x_i^{\text{old}} + \sigma \cdot N_i(1, 0).$$

This forms a "random walk" in the continuous weight space. Or if we want a "hill-climber", we give the mutation λ times to the current point. Then the point with the *best* fitness among these λ mutants is selected as the next point.

3 Hopfield Model & Fitness Evaluation

Before proceeding further, we summarize the Hopfield model briefly, and define our fitness evaluation.

The Hopfield Model

The Hopfield model consists of N neurons and N^2 synapses in which each neuron takes one of two states ± 1. Then patterns

$$\xi^\nu = (\xi_1^\nu, \cdots, \xi_N^\nu), \qquad \nu = 1, \cdots, p$$

where ξ_i^ν takes the value of either 1 or -1, are to be memorized as equilibrium states. After one of these patterns is given to the network, neuron states are updated asynchronously, i.e., at most

one neuron changes its state at a time, according to

$$s_i(t+1) = sgn\left(\sum_{j\neq i}^{N} w_{ij}s_j(t)\right),$$

where $s_i(t)$ is a state of i-th neuron at time t. The behaviors of the collective states of individual neurons are characterized by synaptic weights. Provided the synaptic weights w_{ij} are determined appropriately, networks store some number of patterns as associative memory. Hopfield specified w_{ij} by the rule which was at first advocated by Hebb [3] and later formalized by Cooper [16], i.e.,

$$w_{ij} = \frac{1}{N}\sum_{\nu=1}^{p} \xi_i^\nu \xi_j^\nu \ (i \neq j),\ w_{ii} = 0.$$

Then, giving one of the memorized patterns, possibly including a few errors, to the network as an initial state will result in the stable state after updating.

Fitness Evaluation

When one of stored patterns ξ^ν is given to the network as an initial state, the state of neurons varies from time to time afterwards (unless ξ^ν is stored as a fixed point). In order for the network to function as an associative memory, these two states must be similar. The similarity as a function of time is defined by,

$$m^\nu(t) = \frac{1}{N}\sum_{i=1}^{N} \xi_i^\nu s_i^\nu(t).$$

This is conventionally called an *overlap*. The $m^\mu(t)$ is calculated for each stored pattern and averaged, being denoted as $\langle m(t)\rangle$. Then we define fitness value as a temporal average of this $\langle m(t)\rangle$. That is, our fitness value f is

$$f = \frac{1}{t_0 \cdot p}\sum_{t=1}^{t_0}\sum_{\nu=1}^{p} m^\nu(t).$$

In this paper, t_0 is set to $2N$, twice the number of neurons. Note that the fitness 1 implies all the p patterns are stored as fixed points, while fitness less than 1 includes many other cases.

4 Experimental Results and Discussion

All the experiments in this paper were made on networks with 49 ($= N$) neurons. After storing p random bipolar patterns using the Hebbian learning rule, the fitness landscapes were investigated. The weight configuration determined by the Hebbian rule is naturally the global optimal point in the weight space, when p is small enough. However, the landscape also includes the other global optima as well as many local optima. When $p = 1$, the model corresponds to the Mattice model of Ising spin [18] where the landscape is smooth. On the other hand, when p is infinitely large, the model corresponds to the Sherringon-Kirkpatrick model [19], where the landscapes are highly rugged. So far, we have experimented with p in between these two extremes, with the aim being obtaining the maximum p, under which the network remains content addressability. In this paper, we concentrate on small p cases (p =1, 2, 3). The reason for choosing small p is that we expect the comparatively smooth landscape makes it easier for us to obtain the insight into its geometry than the larger p cases.

4.1 Fitness Distribution

To know the rough image of the fitness landscape, the simplest way is to pick up samples randomly from the whole weight space. So, we pick up, say, 240,000 samples randomly from the fitness landscape derived from the Hebb rule associative memory that had learned only one pattern ($p = 1$).

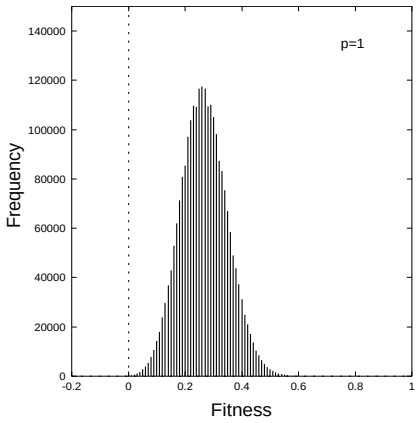

Figure 1: **Fitness along a down hill walk**

In Figure 1, we show the fitness distribution in this case. From this Figure, we can learn two things. First, the distribution of the fitness values is very similar to the Gaussian distribution. This suggests that the landscape is more or less like the one in which each point is assigned random fitness value, such as Kauffman's NK model [15]. Second, it is very difficult for the randomly chosen samples to statistically represent the whole space. As a matter of fact, we can see in the Figure that it is highly unlikely for a randomly chosen point to be near one of peaks. We can guess from this observation that the peaks are very sharp.

4.2 Down Hill Random Walk from a Peak

Then, we experiment a down hill walk from a peak which is one of the global optima found by other evolutionary technique, like Evolutionary Programming [17]. In this case, the walk is guaranteed to include the information of the points around at least one optimum.

Though the primary objective of this experiment is to find an appropriate step size of the walk formed by the Gaussian mutation, and in fact, it is from this experiment that we set the value of σ to 0.002, we can obtain further information about the peak. Different random walks from a specific point are expected to start in different directions. In addition, our Gaussian mutation random walk marches approximately radially from the starting point until it is trapped at a local/global optimum. Hence the observations of several random down hill walks will give us a geometry of the peak. Here, we showed three down hill walks from a global optimum peak in Figure 2. The Figure suggests that the geometry is not isotropic around the peak.

4.3 Gaussian Mutation Hill Climbing

In this subsection, we show the results of Hill climbing walks by the Gaussian mutation, exploring the fitness landscapes constructed for p =1, 2 and 3. All the walks here started at the origin at the weight space.

The walk length, by which we mean the number of steps of a walk needed from its starting point to the local/global optimum that captures the walk, is strongly related to both the density and basin size of the local optima in the landscape. We show this following after Macken et al. [11].

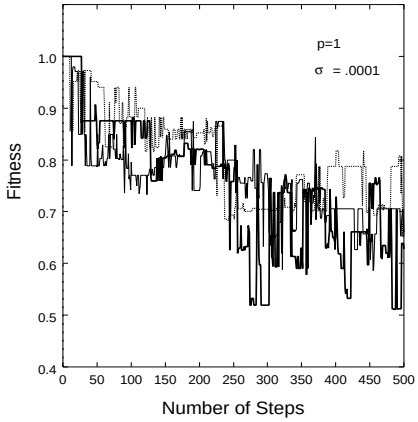

Figure 2: **Fitness along a down hill walk**

Now, let us denote the probability for a point to reach optimum in exactly k steps, starting at the fitness u_0, as $p_k(u_0)$. In order for this to occur, two independent events must happen. First, the point must reach a fitness in the range $(u, u+du)$ in k steps starting at the fitness u_0; Second, all the λ mutant neighbors of the point must have the fitness lower than u. Assume these probabilities are $f_k(u;u_0)du, (u_0 < u < 1)$ and $h(u)$, respectively. Then we obtain

$$p_k(u_0) = \int_{u_0}^{1} f_k(u;u_0)h(u)du.$$

Clearly the function $h(u)$ is roughly the *density of local optima in the space*, and $f_k(u;u_0)du$ should related to the size of basin of attraction of the local optima, though we have not known the explicit form of each function in the equation.

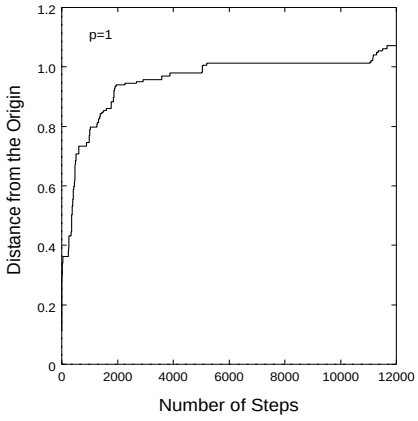

Figure 3: **Typical Fitness vs Generation**

We repeated the simulation 120 times for $p = 1$, and 30 times for both $p = 2$ and 3, with different random number seed at the start. The stop criterion that we set in this paper is 12,000 steps unless one of the global optima is found. For $p = 1$, 32 out of 120 runs located one of the global optima, while for $p = 2$ and 3, all the 30 runs are trapped at a local optimum. The statistics are shown in Table 1.

The probability for a walk to escape the local optimum is not zero. It is in contrast to the mutation used in standard genetic algorithm with binary chromosomes (flip a bit). In fact, we observed that some walks escaped the local optima after remaining there for a long period, as shown in Figure 3. This suggests that our stop criterion 12,000 steps is not long enough. Nevertheless, our observation of the tendency toward longer walk length to the local optima as p increases (see minimum and average walk length in Table 1), seems to tell us that the larger the number of patterns, the smaller the basin size.

Table 1: Statistics of the walks to the optimum

number	ultimate-fitness				walk length				number
of patterns	max	min	avg	std	max	min	avg	std	of runs
(succeeded in locating one of the global optima)									
1	1.000	—	1.000	—	10959	18	347	143.9	32
(trapped at one of the local optima)									
1	0.980	0.853	0.950	0.00298	11450	10	905	251.9	88
2	0.984	0.866	0.927	0.00645	11795	136	2768	577.2	30
3	0.926	0.789	0.879	0.00667	11674	1133	6185	674.3	30

5 CONCLUSION

We have described explorations of evolutionary walks on the fitness landscape of the Hopfield model of associative memory, in which a few patterns are stored by the Hebbian learning algorithm. The landscape is defined on the *continuous* synaptic weight space by assigning the ability of the network to store patterns as associative memory, to each point in the space. Since the space is astronomically huge, statistics of the point sampled at random or along a random walk do not work to represent the whole space. In this paper, the walks are formed by the Gaussian random mutation. Because of its indeterministic property, which is in contrast to the standard binary mutation, the statistical analysis, such as using *correlation* between two neighboring points, or *autocorrelation function* along an evolutionary walk are also very resistant to be used. Hence, as a first step to know the geometry of our landscape, we observed down hill random walk from a peak and Gaussian mutation hill climbing. We have conjectured that the shape of the local hills is not isotropic around the peak, and the sharpness of the hills (the size of basin of attraction) becomes smaller as the number of storing patterns increases.

The study has not been matured at all, and many things are waiting for our explorations. The next step will be to analytically derive the description of the walk length as a function of both the density and the basin size of local optima.

References

[1] J. Komlós, and R. Paturi (1988) *Convergence Results in an Associative Memory Model.* Neural Networks 1, pp239–250.

[2] J. J. Hopfield (1982) *Neural Networks and Physical Systems with Emergent Collective Computational Abilities.*, Proceedings of the National Academy of Sciences, USA, 79, pp2554-2558.

[3] D. O. Hebb (1949) *The Organization of Behavior.*, Wiley.

[4] A. Imada, and K. Araki (1995). *Genetic Algorithm Enlarges the Capacity of Associative Memory.* Proceedings of the 6th International Conference on Genetic Algorithms. pp413–420.

[5] A. Imada, and K. Araki (1997) *Evolution of Hopfield Model of Associative Memory by the Breeder Genetic Algorithm.* Proceedings of the 7th International Conference on Genetic Algorithms, pp784–791.

[6] E. Gardner (1988) *The Phase Space of Interactions in Neural Network Models.* J. Phys. 21A, pp257-270.

[7] A. Bray, and M. Moore (1980) *Metastable states in spin glasses.* Journal of Physics, C: Solid State Physics, 13, ppL469–L476.

[8] A. Bray, and M. Moor (1981) *Metastable states in spin glasses with short-ranged interactions.* Journal of Physics, C: Solid State Physics, 14, pp1313–1327.

[9] B. Derrida, and E. Gardner (1986) *Metastable states of a Spin Glass Chain at 0 Temperature.* Journal of Physique 47, pp959–965.

[10] S. Wright (1932) *The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution.* Proceedings of the 6th International Congress of Genetics, pp356–366.

[11] C. Macken, P. Hagan, and A. Perelson (1991) *Evolutionary walks on rugged landscapes.* SIAM Journal of Applied Mathematics, 51, pp799–827.

[12] W. Fontana, and P. Schuster (1987) *A Computer Model of Evolutionary Optimization.* Biophysical Chemistry, 26, pp123–147.

[13] J. M. Smith (1970) *Natural Selection and the Concept of a Protein Space.* Nature,225. pp563.

[14] E. D. Weinberger (1990) *Correlated and uncorrelated fitness landscapes and how to tell the difference.* Biological Cybernetics, 63, pp325–336.

[15] S. A. Kauffman, and S. Levin (1987) *Towards a general theory of adaptive walks on rugged landscapes.* Journal of Theoretical Biology, 128, pp11–45.

[16] L. N. Cooper (1973) *A Possible Organization of Animal Memory and Learning.* Proceedings of the 24th Nobel Symposium on Collective Properties of Physical Systems. pp252–264

[17] A. Imada, and K. Araki (1998) *How does the Landscape of a Hopfield Associative Memory Look Like?* (submitted.)

[18] D. C. Mattis (1976) *Solvable Spin Systems with Random Interactions.* Physics Letters 56A, pp421–422.

[19] D. Sherrington and S. Kirkpatrick (1975) *Solvable Model of a Spin Glass.* Physical Review Letters 35, pp1792–1796.