Domains of attraction and the density of static metastable states, in single-pattern iterated neural networks

1991 J. Phys. A: Math. Gen. 24 1083

(http://iopscience.iop.org/0305-4470/24/5/022)

J. Phys. A: Math. Gen. 24 (1991) 1083–1092. Printed in the UK

Domains of attraction and the density of static metastable states in single-pattern iterated neural networks

Thomas B Kepler†

Physics Department, Brandeis University, Waltham, MA 02254, USA

Received 27 July 1990, in final form 26 October 1990

Abstract. I calculate, for a single-pattern iterated neural network, the density of static metastable states, $\partial\mathcal{N}/\partial m$, as a function of the field distribution and the symmetry of the synaptic matrix. The features of this function strongly suggest that the boundary upon which this density vanishes gives the critical overlap for the memory state, i.e. gives a measure of the size of its domain of attraction. This heuristic interpretation is shown to agree with the exact result for vanishing symmetry, the only case where a direct calculation can be performed. I explicitly calculate critical overlaps as a function of symmetry and mean field strength when the field distribution is a delta-function and when it is a unit-width Gaussian.

1. Introduction

Spin-glass inspired neural networks [1,2] have been used as a paradigm for auto-associative memory, in which an initial condition (an *input* state) containing partial or noisy information on one of the memorized patterns evolves via the system dynamics to the completed pattern. In other words, the memorized patterns are attractors. The system is comprised of N Ising-valued ($+1$ or -1) neural elements interacting through the 'synaptic' coupling matrix $\mathbf{J}$. The state of the system at time t is represented by the vector $\boldsymbol{S}(t)$, and evolves according to the dynamics given by

$$S_i(t+1) = \mathrm{sign}\left(\sum_{j=1}^{N} J_{ij} S_j(t)\right). \tag{1}$$

The overlap of two states $\boldsymbol{S}^1$ and $\boldsymbol{S}^2$ gives a measure of their similarity, or proximity in phase space, and is simply their normalized dot-product

$$m(\boldsymbol{S}^1, \boldsymbol{S}^2) = \frac{1}{N}\sum_{i=1}^{N} S_i^1 S_i^2. \tag{2}$$

If the system has learned p memory patterns $\{\boldsymbol{\xi}^\mu\}_{\mu=1,2,\ldots,p}$, that is, the $\boldsymbol{\xi}^\mu$ have been made fixed points of equation (1), then a question of substantial importance is: how large an overlap with a given memory state $\boldsymbol{\xi}^1$ is required of an initial state

† Present address: Biology Department, Brandeis University, Waltham, MA 02254, USA.

0305-4470/91/051083+10$03.50

to ensure that it evolves to $\boldsymbol{\xi}^1$? What is the size of the domain of attraction? The domain of attraction for the fixed point $\boldsymbol{\xi}^1$ is defined as the set of all those initial states that flow to $\boldsymbol{\xi}^1$ after sufficient time. The work of Forrest [3] suggests that this question is sensible; there is a critical value, m_c, of the overlap such that almost all initial states with $m_0 > m_c$ will flow to the memory state, and almost all the rest will not. A direct determination of the domains of attraction would require a calculation of the sequence of functions $m_n(m_0)$ giving the overlap at the nth time step as a function of the initial overlap and then finding the limit as $n \longrightarrow \infty$. For the special case where $\mathbf{J}$ is separable [1] (the so-called Hebb matrix) Gardner *et al* [4] have laid out this calculation and have given explicit results for $n = 1$ and $n = 2$. Beyond this, the number of order parameters increases dramatically and immediately becomes prohibitive. So even in this restricted and relatively simple case, a direct approach is hopeless.

Kepler and Abbott [5] have calculated the overlap after one time step as a function of the initial overlap for any matrix $\mathbf{J}$. This function is given by

$$m_1(m_0) = 1 - 2\int \mathrm{d}\gamma\, \rho(\gamma) H\left(\frac{m_0\gamma}{\sqrt{1-m_0^2}}\right). \tag{3}$$

where H is defined by

$$H(a) = \frac{1}{\sqrt{2\pi}}\int_a^\infty \exp(-\tfrac{1}{2}x^2)\,\mathrm{d}x. \tag{4}$$

By analogy with spin systems I have defined the *field strengths* γ_i^μ, by

$$\gamma_i^\mu = \frac{\xi_i^\mu \sum_{j=1}^N J_{ij}\xi_j^\mu}{\sqrt{\sum_{j=1}^N J_{ij}^2}}. \tag{5}$$

and $\rho(\gamma)$ is the distribution function for these fields over the sites i (for the fixed point under discussion). Note that a state is stable if and only if its field at each site is positive.

This calculation is not readily extended to subsequent time steps. The reason for this and the reason for the rapid multiplication of order parameters in the Gardner calculation are the same. The average over initial states $\boldsymbol{S}(0)$ satisfying $m(\boldsymbol{S}(0), \boldsymbol{\xi}^1) = m_0$ can be performed without difficulty because these states are statistically independent. The states $\boldsymbol{S}(1)$, however, are not independent, but have acquired correlations in the dynamics through correlations within $\mathbf{J}$.

It is an interesting, directly observable feature of neural networks that the system may initially flow toward the memory state (m_n increasing), but then after some time, flow away (m_n decreasing). So in general, one may not write

$$m_{n+1}(m_0) = m_1(m_n(m_0)). \tag{6}$$

The only case where no flow reversal of this type occurs, and where equation (6) is valid, is when the matrix symmetry σ, defined by

$$\sigma^2 = \frac{\sum_{i,j=1}^N J_{ij}J_{ji}}{\sum_{i,j=1}^N J_{ij}^2} \tag{7}$$

vanishes [6]. It is also known that in this case the number of *static metastable states* no longer scales as $e^{\alpha N}$, but is at most polynomial in N (see below). Static metastable states, or spurious fixed points, are those states which are fixed points under the dynamics of equation (1), but which are not desired memory states. There is significance in the correlation between the existence of static metastable states and the non-iterative nature of the overlap dynamics. It usually is the case that those initial points which move first toward the memory state and then away from it get stuck in one of the metastable states; the metastable states impede flow to the attractor. The spin-glass phase described in the thermodynamic analysis of Amit *et al* [7] is a manifestation of this effect.

This close connection between flow restriction and static metastable states suggests that a study of the density of static metastable states over m with analysis of its dependence on the distribution function, ρ, and the symmetry, σ, may prove fruitful†.

Some previous results on the distribution of metastable states are available. It is well known from spin-glass theory [8] that the number, $\mathcal{N}$, of metastable states for a 'typical' symmetric matrix (one whose entries above or below the main diagonal are independent and distributed normally) is exponentially large, $\mathcal{N} = \exp(0.199N)$, and that this is self-averaging; for $N \to \infty$ almost every matrix gives this result and therefore, so does the average over such matrices. This calculation can be extended to arbitrary symmetry [9], and one sees the typical decrease of $\mathcal{N}$ as σ decreases.

Gardner [10] extended this approach to look at the dependence of $\mathcal{N}$ on m for the separable Hopfield–Hebb model in order to investigate the structure of the attractors, which in this case are not the memory patterns themselves, but instead, are clustered around them. Treves and Amit [11] calculated the density of static metastable states in asymmetrically diluted Hopfield–Hebb networks to determine the effectiveness of such a dilution for suppressing the spin-glass phase.

Krauth, Nadal and Mézard [6] introduced the single-pattern iterated neural network (SPINN), in which a single memory pattern is learned and the matrix symmetry is constrained to some specified value (the matrix is otherwise random) to investigate the effect of asymmetry on domain size. To this system, they applied Gardner's technique for calculating overlap dynamics and were able to calculate the overlap functions out to four time steps before the explosion of order parameters proved overwhelming. They also performed numerical simulations and found that symmetry does indeed play a dominant role in determining domain size.

I have adopted their system and will calculate the density $d\mathcal{N}/dm$ of static metastable states for a SPINN with arbitrary field distribution, ρ, and symmetry, σ.

2. The density of static metastable states

For a given matrix, $\mathbf{J}$, the number of states stable to single-spin flips is

$$\mathcal{N}(\mathbf{J}) = \mathrm{Tr}_{S} \prod_{i=1}^{N} \theta\left(S_i \sum_{j=1}^{N} J_{ij} S_j \right). \tag{8}$$

† There are, of course, non-static metastable states, i.e. limit cycles, in asymmetric networks, and their presence may affect the retrieval of memory states. Nevertheless, the heuristic motivations above refer specifically to the static states, and since the consideration of stable limit cycles of arbitrary length would render the problem intractable, I will restrict my attention to the study of static metastable states.

I will average this quantity with the appropriate measure on matrix space to get the typical number of static metastable states

$$\bar{\mathcal{N}} = \mathrm{Tr}_{\boldsymbol{S}} \int \mathrm{d}\mathbf{J}\, P(\mathbf{J}) \prod_{i=1}^{N} \theta\Big(S_i \sum_{j=1}^{N} J_{ij} S_j\Big). \tag{9}$$

The measure is chosen by introducing the matrix $\mathbf{T}$ related to $\mathbf{J}$ by

$$J_{ij} = \frac{1}{\sqrt{2(1+a^2)}}[(T_{ij} + T_{ji}) + a(T_{ij} - T_{ji})] \tag{10}$$

where a is related to σ through

$$a^2 = \frac{1+\sigma^2}{1-\sigma^2}. \tag{11}$$

I impose on $\mathbf{T}$ the normalization condition

$$\sum_{j=1}^{N} T_{ij}^2 = N \tag{12}$$

and choose the matrices with the *prior* probability $f(\gamma)$ on the fields associatied with each matrix. The function f is introduced here simply to make the calculation easier to formulate and will be eliminated in favour of the physically relevant distribution ρ shortly (see [12]). I will make the convenient gauge choice of letting the single pattern be ferromagnetic, i.e. $S_i = +1$ for each i, so that

$$\gamma_i = \frac{1}{\sqrt{N}} \sum_{j=1}^{N} J_{ij}. \tag{13}$$

Now the matrix distribution function P is given by

$$P(\mathbf{J})\,\mathrm{d}\mathbf{J} = C_\sigma[f] \prod_{i=1}^{N} \Big[\delta\Big(\sum_{j=1}^{N} T_{ij}^2 - N\Big) f\Big(\frac{1}{\sqrt{N}} \sum_{j=1}^{N} J_{ij}\Big)\Big]\,\mathrm{d}\mathbf{T} \tag{14}$$

where $C_\sigma[f]$ is a normalization constant, calculated to be

$$C_\sigma[f] = \exp N\Big(-\frac{1}{2} + \frac{1}{2}w^2 - \log \int \frac{\exp(-\frac{1}{2}x^2)\,\mathrm{d}x}{\sqrt{2\pi}} f(x - \sigma w)\Big). \tag{15}$$

The order parameter w is given by its saddle-point equation

$$w - \frac{\partial}{\partial w} \log\Big(\int \frac{\exp(-\frac{1}{2}x^2)\mathrm{d}x}{\sqrt{2\pi}} f(x - \sigma w)\Big) = 0. \tag{16}$$

Now ρ is related to f by

$$\rho(\gamma) = \int \mathrm{d}\mathbf{J} P(\mathbf{J}) \delta\Big(\gamma - \frac{1}{\sqrt{N}} \sum_{j=1}^{N} J_{1j}\Big) \tag{17}$$

so

$$\rho(\gamma) = \frac{\exp[-\frac{1}{2}(\gamma - \sigma w)^2] f(\gamma)}{\sqrt{2\pi}\int[\exp(-\frac{1}{2}x^2)/\sqrt{2\pi}]\,\mathrm{d}x f(x - \sigma w)}. \tag{18}$$

Note that in equation (9), $\tilde{\mathcal{N}}$ is invariant under a rescaling of f. Therefore, in that equation, one may use equation (17) to replace $f(\gamma)$ by $\rho(\gamma)\exp[\frac{1}{2}(\gamma - \sigma w)^2]$. Equation (16) becomes

$$w = -\frac{\sigma}{1+\sigma^2}\langle\gamma\rangle_\gamma \tag{19}$$

where the angle brackets denote averaging over γ using $\rho(\gamma)$:

$$\langle g(\gamma)\rangle_\gamma \equiv \int \mathrm{d}\gamma\, \rho(\gamma) g(\gamma). \tag{20}$$

The calculation of equation (9) is performed by making repeated use of the Fourier representation of the delta function

$$1 = \frac{1}{2\pi}\int \mathrm{d}x\, \mathrm{d}\hat{x}\, \exp(\mathrm{i}x\hat{x}) \tag{21}$$

to introduce several order parameters including the overlap with the single (ferromagnetic) memory state

$$m = \frac{1}{N}\sum_{i=1}^{N} S_i. \tag{22}$$

The resulting integral factorizes over sites, leaving an expression whose evaluation by saddle-point integration becomes exact as $N \longrightarrow \infty$.

Let $F(m)$ be defined by

$$\frac{\mathrm{d}\tilde{\mathcal{N}}}{\mathrm{d}m} = \mathrm{e}^{NF(m)}. \tag{23}$$

Then

$$\begin{aligned} F(m) = &-\frac{1}{2}\sigma^2 t^2 + t(s + m\sigma^2 u) - \frac{1}{2}(1+\sigma^2)u^2 - \frac{1}{2}v^2 + \frac{1}{2}(1+\sigma^2)w^2 \\ &+ \left\langle \log\left\langle \mathrm{e}^{\sigma\gamma(w+\tau t-u)} H\left(\frac{-\tau m(\gamma + \sigma(u - \tau t)) + \sigma(v + \tau s)}{\sqrt{1-m^2}}\right)\right\rangle_\gamma \right. \\ &\left. - \log\left(\frac{1+\tau m}{2}\right)\right\rangle_\tau \end{aligned} \tag{24}$$

where the subscripted angle brackets denote averages over the subscript variable as above for γ and

$$\langle g(\tau)\rangle_\tau \equiv \sum_{\tau=\pm 1}\frac{1+\tau m}{2} g(\tau). \tag{25}$$

The order parameters s, t, u and v are given by saddle-point equations, e.g.

$$-v + \frac{\partial}{\partial v}\left\langle \log\left\langle e^{\sigma\gamma(w+\tau t-u)} H\left(\frac{-\tau m(\gamma + \sigma(u - \tau t)) + \sigma(v + \tau s)}{\sqrt{1-m^2}}\right)\right\rangle_\gamma\right\rangle_\tau = 0 \tag{26}$$

etc. w is still given by equation (19).

It should be noted that $F(m)$ is not self-averaging over all parameter space, and so the calculation given here must, strictly speaking, be interpreted as providing an upper bound on the value of $F(m)$ rather than the typical value itself. This said, I believe it unlikely that the features of importance in the following are misrepresentented in this approximation.

3. Discussion

Equation (24) is the main result of the computational part of this paper. But in this general form, it leaves little for one's intuition to grasp. One can get a feeling for its content by evaluating it for specific functions in the place of ρ. The two that I will consider are

$$\rho(\gamma) = \delta(\gamma - \gamma_0) \tag{27}$$

and

$$\rho(\gamma) = (1/\sqrt{2\pi}) \exp[-\tfrac{1}{2}(\gamma - \gamma_0)^2]. \tag{28}$$

These choices are made because they are the field distributions for two well known matrix learning rules (the former is that of the pseudo-inverse matrix [13] and the latter, that of the Hebb matrix) and though their properties when $|\gamma_0|$ is large may be expected to be similar, they have strongly contrasting behaviour when $|\gamma_0|$ is small. As γ_0 passes through zero from below, the pattern with the delta-function distribution suddenly becomes stable, whereas the pattern with the normal Gaussian distribution becomes stable only asymptotically. The value of these considerations will be evident when I analyse the features of equation (24) thus constrained.

The former choice yields

$$F(m) = -\frac{1}{2}v^2 - \frac{1}{2}\mu t^2 - tmv - \left\langle \log\left(\frac{1+\tau m}{2}\right)\right\rangle_\tau + \left\langle \log H\left(\frac{-\tau m\gamma_0 + \sigma v(1-\tau m) + \sigma t(m - \tau\mu)}{\sqrt{1-m^2}}\right)\right\rangle_\tau \tag{29}$$

where

$$\mu = m^2 - \sigma^2(1 - m^2). \tag{30}$$

The latter gives, quite simply

$$F(m) = -\frac{1}{2}v^2 + \left\langle \log H(\sigma v - \tau m\gamma_0) - \log\left(\frac{1+\tau m}{2}\right)\right\rangle_\tau. \tag{31}$$

$F(m)$ for the delta-function distribution is plotted in figure 1 for two values of γ_0 and of σ. F begins, at $m = 0$, at some positive value (this value depends on σ, but not on ρ), and then declines, cut off by the phase space term. For γ_0 positive, the curve reverses direction, increases and becomes positive, reaches a maximum, and crosses back through zero and then remains negative. Referring to equation (23), note that for N large, one expects to find static metastable states where F is positive. Where it is negative we expect none.

The appearance of a band of metastable states at values of m near 1 may be thought of as resulting from the 'energetically' favourable condition of being near the stable state. However, it is apparently not good to be *too* close, perhaps because one is then *in the domain of attraction* of the stable state, and metastable states cannot exist here. This suggests that the boundary of this band may coincide reasonably well with the edge of the attractor's domain. But one must be aware that even when the density is exponential in N, the *fraction* of metastable states among all states may be vanishingly small as $N \longrightarrow \infty$. A direct and unambiguous interpretation in these

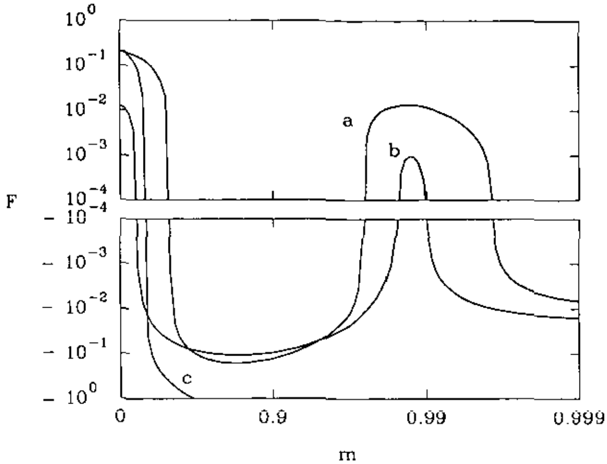

Figure 1. A plot of $F\ (= (1/N)\log(\partial\mathcal{N}/\partial m))$ against m for the case $\rho(\gamma) = \delta(\gamma-\gamma_0)$. Curve a, $\sigma = 1$, $\gamma_0 = 0.4$; curve b, $\sigma = 0.2$, $\gamma_0 = 0.4$; curve c, $\sigma = 1$, $\gamma_0 = -0.4$.

terms is thus precluded. I will examine this somewhat further below. Nevertheless it seems reasonable to analyse the boundary where static metastable states disappear. Thus, one is led to study $m_c(\gamma_0, \sigma)$, the zeros of F.

Figures 2 and 3 plot the locus of the zeros of $F(m)$ as a function of γ_0 for $\sigma = 0.2$ and $\sigma = 1$. These curves bound the region where F is positive. For the delta-function distribution, figure 2, one sees m_c increase with γ_0 for γ_0 negative, and then suddenly at $\gamma_0 = 0$ a second band of metastable states appears. This band widens as γ_0 increases but both edges recede from $m = 1$. I suggest that the recession of the upper edge is due to the growth of the memory state's domain of attraction at the expense of nearby metastable states. At higher values of γ_0, the two bands fuse into one.

A comparison of figures 2 and 3 lends some support to this supposition. Figure 3 gives $m_c(\gamma_0)$ for the Gaussian distribution given above. There is no sudden appearance of a band of metastable states, but instead the region where metastable states occur simply encroaches upon larger m smoothly as γ_0 increases. In particular, the upper edge of this band never recedes. The memory state is *not* stable in this case, and therefore has no domain of attraction to devour the metastable states. Above some σ-dependent critical value of γ_0, however, the band of metastable states splits. The upper band narrows and both of its boundaries approach $m = 1$ asymptotically while the boundary of the lower band recedes. The upper band of metastable states has become the attractor in lieu of a stable memory state. And now the recession of the lower band is a response to the increasing domain size of this attractor. Note in passing that the joining of the two bands for γ_0 below the critical value is closely related to the phase transition where recall gives way to the spin-glass phase in the thermodynamics of the Hopfield–Hebb network [10]. For the comparable system here (i.e., taking $\sigma = 1$) this critical value, from figure 3, is $\gamma_0 = 2.19\cdots$ compared with $\gamma_0 = 2.68\cdots$ in the Hopfield–Hebb case [7]. I expect these numbers to be similar, but not identical. Gardner's computation, done specifically for the Hopfield–Hebb construction, found that the bands split at a capacity corresponding to $\gamma_0 = 3.02$, rather than the value of 2.68 at which retrieval fails. The further difference found here is likely due to the Hopfield–Hebb matrix being a rather unusual (and sub-optimal)

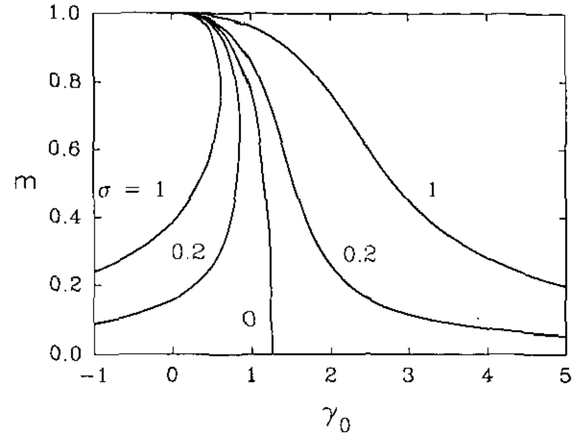

Figure 2. The locus of the zeros of $F(m)$ for $\rho(\gamma) = \delta(\gamma - \gamma_0)$. The curves are labelled by the value of σ used. The region between the two curves with the same label is the region where static metastable states exist. This region narrows with decreasing σ and vanishes as the curves collapse into one at $\sigma = 0$. By the interpretation advanced here, the rightmost curve gives the critical overlap for the domain of attraction.

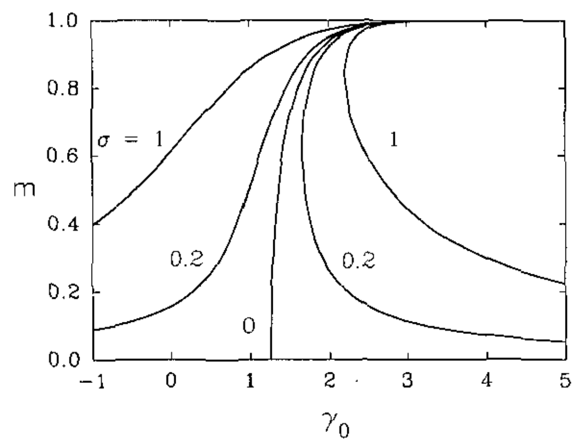

Figure 3. As figure 2, but for $\rho(\gamma) = (1/\sqrt{2\pi}) \exp[-\frac{1}{2}(\gamma - \gamma_0)^2]$.

member of the class of symmetric standard Gaussian matrices [12].

In further support of my interpretation of these features, I would like to consider the case $\sigma = 0$. Here the width of the band shrinks to zero, but its limiting location is still well defined. This case is of interest because it is the one instance where one does know the domain size [6]. It is given simply by the unstable fixed point of equation (6), where the function $m_1(m_0)$ is given by equation (3). Going back to general ρ and equation (24), note that equation (26) shows that vanishing σ implies that s, t, u, v and w likewise vanish, and I have that

$$F(m) = \frac{1+m}{2}\left[\log\left\langle H\left(-\frac{m\gamma}{\sqrt{1-m^2}}\right)\right\rangle_\gamma - \log\left(\frac{1+m}{2}\right)\right]$$

$$+\frac{1-m}{2}\left[\log\left\langle H\left(\frac{m\gamma}{\sqrt{1-m^2}}\right)\right\rangle_\gamma - \log\left(\frac{1-m}{2}\right)\right]. \tag{32}$$

It is easy to see that the fixed point of equation (3), when inserted in equation (32), causes the vanishing of each bracketed term separately, and obviously gives $F = 0$. That is, the two methods agree when $\sigma = 0$.

When σ is not zero, one must turn to numerical simulation. Figure 4 shows a comparison of the critical overlap as calculated in the foregoing, and computer simulations of a 100-node network. The data was obtained using matrices trained on a variant of the perceptron learning algorithm [14] that produces matrices with sharply peaked field distributions, and large but not perfect symmetry. I have therefore used equation (29) with $\sigma = 0.9$ for this comparison. Also included in this figure is a curve giving the prediction of the phenomenological rule of [5]. The agreement here is not unreasonable, though in and of itself, it provides less than complete confirmation. There is, however, some difficulty in interpreting numerical results of this type; for finite N the domain boundaries are not sharp, and there is some arbitrariness in choosing how to locate m_c. For the data presented here, I took m_c to be the value at which 90% of initial points make it to the memory state. Taking this fraction to be 100% adds significantly to the value of m_c at intermediate values of γ_0. Bearing in mind the caveat expressed in the discussion following equation (31), for finite N one might expect, and does indeed see, a small region of overlap between the measured domain and the area supporting static metastable states. This may be taken into account by imagining the flow toward the memory pattern of an initial state located in the region containing the static metastable states as a percolation process, with initial states located just within the edge of the metastable region having a non-zero probability of escape to the memory pattern. This would require one to account for the domains of attraction of the metastable states themselves, but would likely prove both informative and interesting.

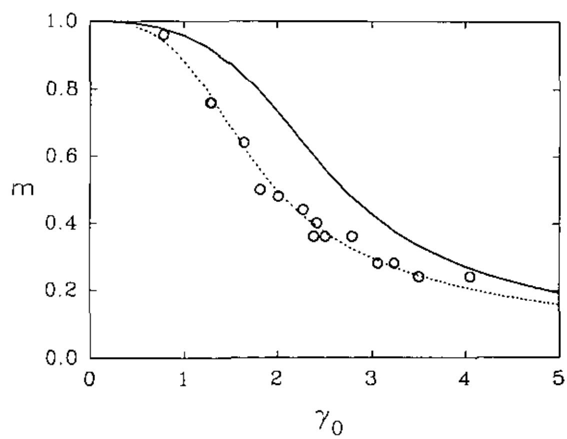

Figure 4. A comparison of m_c computed as the locus of zeros for $F(m)$ (full curve), as computed by the phenomenological rule of [5] (dotted curve), and as numerically determined (open circles). For the first calculation I have used the delta function field distribution, and $\sigma = 0.9$. For the numerical data (from [5]) one has $N = 100$.

In conclusion, I have related the distribution of static metastable states to the domain of attraction of the attractor in a single-pattern neural network. The analysis suggests that the boundary of the region which supports metastable states also demarcates the boundary of the domain of attraction of the memorized pattern, in the sense that *all* points in this domain flow to the memory state, although a non-negligible fraction of those outside it may likewise complete the pattern successfully.

Acknowledgments

I would like to thank L F Abbott for useful discussions and valuable suggestions. This work was supported in part by National Institutes of Health grant T32NS07292.

References

[1] Hopfield J J 1982 *Proc. Natl Acad. Sci. USA* 79 2554
[2] Gardner E 1988 *J. Phys. A: Math. Gen.* **21** 257
[3] Forrest B M 1988 *J. Phys. A: Math. Gen.* **21** 245
[4] Gardner E, Derrida B and Mottishaw P 1987 *J. Physique* **48** 741
[5] Kepler T B and Abbott L F 1988 *J. Physique* **49** 1657
[6] Krauth W, Nadal J-P and Mézard M 1988 *J. Phys. A: Math. Gen.* **21** 2995
[7] Amit D, Gutfreund H and Sompolinsky H 1987 *Ann. Phys., NY* **173** 30
[8] Bray A J and Moore M A 1980 *J. Phys. C: Solid State Phys.* **13** L469
Tanaka F and Edwards S F 1980 *J. Phys. F: Met. Phys.* **10** 2471
De Dominicis C, Gabay M, Garel T and Orland H 1980 *J. Physique* **41** 923
[9] Crisanti A and Sompolinsky H 1988 *Phys. Rev.* A **37** 4865
Kepler T B 1989 *PhD Thesis* Brandeis University
[10] Gardner E 1986 *J. Phys. A: Math. Gen.* **19** L1047
[11] Treves A and Amit D 1988 *J. Phys. A: Math. Gen.* **21** 3155
[12] Abbott L F and Kepler T B 1989 *J. Phys A: Math. Gen.* **22** 2031
[13] Kohonen T 1984 *Self Organization and Associative Memory* (Berlin: Springer)
Personnaz L, Guyon I and Dreyfus G 1985 *J. Physique* **46** L359
Kanter I and Sompolinsky H 1987 *Phys. Rev.* A **35** 380
[14] Abbott L F and Kepler T B 1989 *J. Phys. A: Math. Gen.* **22** L711