Toy model for protein folding

Frank H. Stillinger, Teresa Head-Gordon,* and Catherine L. Hirshfeld†
AT&T Bell Laboratories, Murray Hill, New Jersey 07974
(Received 1 March 1993)

A conceptually simple model for protein-folding phenomena has been created: it is two-dimensional and has only two different "amino acids." Ground-state conformations have been determined for all of its flexible polypeptides containing seven or fewer monomers. This complete database displays a wide geometric variety of folded shapes and shows that single point mutations in some cases induce substantial folding modifications. Neural-network concepts have been employed to analyze results. The simplest static neural networks required to act as error-free read-only memories provide a way to visualize the logical structure of underlying folding principles. The topologies of optimal networks found thus far suggest that protein folding has a more complex cooperative character than has been embodied previously in theoretical approaches.

PACS number(s): 87.10.+e, 87.15.By, 42.79.Ta

I. INTRODUCTION

Protein-folding phenomena present a daunting group of scientific challenges. Perhaps this is inevitable, since only a complex and diverse family of molecules could fulfill proteins' assigned roles in basic biological processes. The large and still rapidly growing literature on the subject of protein folding [1–3] chronicles many remarkable advances in both experiment and theory, yet this remains an open problem. Given an arbitrary but fully specified sequence of amino acids, we cannot yet predict the folding pathway of the corresponding polypeptide, the conformation of the final state, nor even verify in all cases whether that final state is one of lowest free energy or simply a metastable "trap" in the kinetic folding pathway.

The strategy selected for the present paper seeks a few modest insights by introducing and exploiting a highly simplified "toy model." Its motivation is roughly analogous to that behind the Ising [4] or Heisenberg [5] models for magnetism: namely to strip away distracting detail in the hope of attaining more penetrating insights. One major advantage of the toy model is that it becomes feasible to determine a complete database of ground-state structures for all "polypeptides" up to some modest (but nontrivial) degree of polymerization. In this respect our approach avoids uncertainties stemming from incompleteness of the real protein-structure database [1]. It also permits an application of neural-network concepts [6] to the interpretation of our model's behavior.

Section II defines our "toy model;" Sec. III presents some of its general properties. Ground-state energies and structures have been determined for all species containing seven or fewer residues, and Sec. IV covers the results of this comprehensive search. Section V introduces the concept of "optimal neural network" for exact representation of our gap-free database, as a means of uncovering the deep logic of the protein-folding patterns. Section VI contains a reprise and assessment of the toy model, its results, and prospects for evolution toward greater realism.

II. THE MODEL

Our model incorporates only two "amino acids," to be denoted by A and B, in place of the 20 that occur naturally. They will be linked together by rigid unit-length bonds to form linear unoriented polymers that reside in two dimensions. As Fig. 1 illustrates, the configuration of any n-mer is specified by the $n-2$ angles of bend $\theta_2,\ldots,\theta_{n-1}$ at each of the nonterminal residues. We adhere to the conventions that

$$-\pi \le \theta_i < \pi \ , \tag{2.1}$$

that $\theta_i = 0$ corresponds to linearity of successive bonds, and that positive angles indicate counterclockwise rotation.

We postulate that two kinds of interactions compose the intramolecular potential energy for each molecule: backbone bend potentials (V_1) and nonbonded interactions (V_2). The former will be independent of the A,B sequence, while the latter will vary with that sequence and will receive a contribution from each pair of residues not directly attached by a backbone bond.

Residue species along the backbone can be conveniently encoded by a set of binary variables $\xi_{1,\ldots,}\xi_n$. If $\xi_i = 1$, the ith residue is A; if $\xi_i = -1$, it is B. The intramolecular potential-energy function Φ thus can be expressed as follows for any n-mer:

$$\Phi = \sum_{i=2}^{n-1} V_1(\theta_i) + \sum_{i=1}^{n-2} \sum_{j=i+2}^{n} V_2(r_{ij},\xi_i,\xi_j) \ . \tag{2.2}$$

The distances r_{ij} can be written as functions of the intervening angles (recalling that backbone bonds have unit length):

$$r_{ij} = \left\{ \left[\sum_{k=i+1}^{j-1} \cos\left[\sum_{l=i+1}^{k} \theta_1 \right] \right]^2 + \left[\sum_{k=i+1}^{j-1} \sin\left[\sum_{l=i+1}^{k} \theta_1 \right] \right]^2 \right\}^{1/2} . \tag{2.3}$$

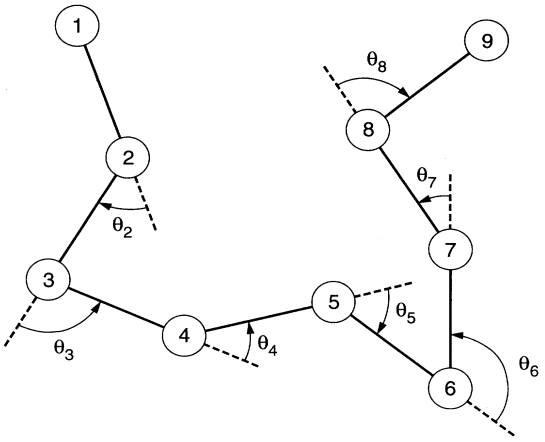

FIG. 1. A schematic diagram of a generic 9-mer, with serially numbered residues, and backbone bend angles.

In the following we shall not consider intermolecular interactions.

Our model assigns a simple trigonometric form to V_1:

$$V_1(\theta_i)=\tfrac{1}{4}(1-\cos\theta_i)\ . \tag{2.4}$$

The nonbonded interactions V_2 have a species-dependent Lennard-Jones 12,6 form:

$$V_2(r_{ij},\xi_i,\xi_j)=4[r_{ij}^{-12}-C(\xi_i,\xi_j)r_{ij}^{-6}]\ , \tag{2.5}$$

$$C(\xi_i,\xi_j)=\tfrac{1}{8}(1+\xi_i+\xi_j+5\xi_i\xi_j)\ . \tag{2.6}$$

On account of Eq. (2.4), successive bonds would tend toward linearity ($\theta_i=0$), if nothing else mattered.

The coefficient $C(\xi_i,\xi_j)$ is $+1$ for an AA pair, $+\frac{1}{2}$ for a BB pair, and $-\frac{1}{2}$ for an AB pair. Consequently the first of these pairs may be regarded as strongly attracting, the second as weakly attracting, and the third as weakly repelling. This diversity mimics in a simple way that of real amino-acid residues, which vary in size, polarity, and degree of hydrophobicity [1]. In fact, results presented below imply that A and B behave respectively as hydrophibic and hydrophilic residues. As will become clear in Sec. III and IV, the interplay between the backbond bend interaction that tends to produce linear structures, and the various combinations of attractive and repulsive nonbonded pair interactions, generates a wide range of ground-state geometries. It is in this last respect that our toy model remains faithful to the character of real proteins.

The description given so far for our model makes no mention of solvent. However, one has the option to include the effects of a solvent medium, at least implicitly, by interpreting Φ in Eq. (2.2) as an intramolecular potential of mean force [7,8].

III. GENERAL FEATURES

A. Enumeration

Two choices exist for each residue of an n-mer, namely A or B. If this linear polymer had an intrinsic directionality to its backbone, the number of distinguishable molecules would be simply 2^n. But that directionality is missing. Any molecule $(\xi_1,\ldots,\xi_n)$ is fundamentally invariant under sequence reversal $(\xi_n,\ldots,\xi_1)$, implying a reduction in the number of distinguishable n-mers.

Consider first the case of $n=2m$, an even integer. If the n-mer is cut at its middle bond, each half that results *has* directionality. Consequently there are 2^m distinguishable half-molecules. Rebonding each of these with its own kind yields 2^m different centrosymmetric molecules. Noncentrosymmetric molecules are reconstituted by bonding inequivalent halves; this can occur in $\frac{1}{2}(2^m)(2^m-1)$ ways. Adding these to the centrosymmetric count yields

$$N(2m)=2^{m-1}(2^m+1) \tag{3.1}$$

for the number of distinguishable $(2m)$-mers.

The number of distinguishable $(2m+1)$-mers follows trivially, since each of them corresponds uniquely to the insertion of a single residue in the center of a $(2m)$-mer. The insertion can be either of two possibilities, A or B. Consequently,

$$N(2m+1)=2^m(2^m+1)\ . \tag{3.2}$$

B. Interaction of parallel chains

In order to illuminate the tendency of the model molecules to fold into compact globular form, we now present an elementary calculation for the interaction of two parallel strands, each linear, with separation D large compared to the bond length. For present purposes these are viewed as portions of the same molecule with a very high degree of polymerization, connected through a remote turn; but they could have equally well been parallel linear strands from distinct molecules.

Let x and y, respectively, be the fractions of A residues in the two chains. Under the given assumption of wide separation, it is only these gross compositions that should matter, not the specific residue sequence details. Furthermore $D\gg 1$ requires only accounting for the r_{ij}^{-6} terms in the nonbonded interactions V_2 acting between the chains.

The net interaction of an A residue on the "x" strand with all residues on the "y" strand can be estimated by the following integral (s measures distance along the latter strand):

$$[-y+\tfrac{1}{2}(1-y)]\int_{-\infty}^{+\infty}\frac{ds}{(D^2+s^2)^3} =\frac{(1-3y)}{D^5}\int_0^{\infty}\frac{dt}{(1+t^2)^3}\ . \tag{3.3}$$

The corresponding result for a B on the "x" strand is

$$[\tfrac{1}{2}y-\tfrac{1}{2}(1-y)]\int_{-\infty}^{+\infty}\frac{ds}{(D^2+s^2)^3} =\frac{(2y-1)}{D^5}\int_0^{\infty}\frac{dt}{(1+t^2)^3}\ . \tag{3.4}$$

The interaction per unit length between the strands re-

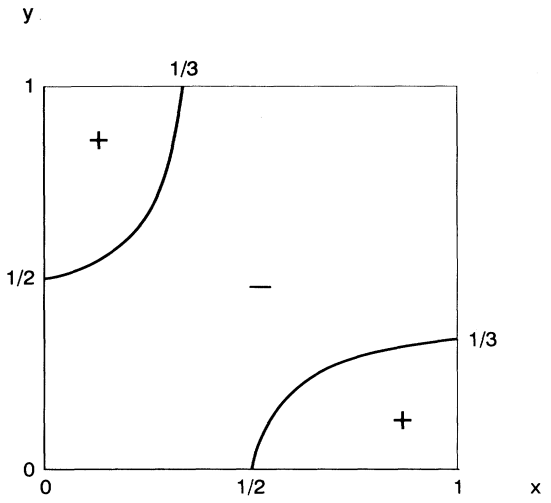

FIG. 2. Regions of net attraction (−) and of net repulsion (+) for two widely separated, linear, parallel strands. Variables x and y stand for the fractions of A residues in the two strands.

quires weighting (3.3) by x, (3.4) by $(1-x)$, and adding; the result is

$$\frac{[-1+2(x+y)-5xy]}{D^5}\int_0^\infty \frac{dt}{(1+t^2)^3} = \frac{3\pi}{16D^5}[-1+2(x+y)-5xy]\ . \quad (3.5)$$

The quantity shown in Eq. (3.5) is negative along the equal-composition line $x=y$, indicating net interstrand attraction. More generally, it changes sign across the hyperbolic locus:

$$y=\frac{1-2x}{2-5x}\ ,\quad x=\frac{1-2y}{2-5y}\ . \quad (3.6)$$

Figure 2 illustrates the resulting sign regions in the relevant unit square of the x,y plane.

If x,y lies anywhere in the net attractive region of Fig. 2, and if the two strands are long enough, the molecule will always find it favorable to fold. The energy required to produce a U turn is positive but fixed, whereas the stabilization between the strands grows linearly with their length. If x,y lies in either net repulsive region of Fig. 2, this scenario is not directly applicable. But if, as assumed, the strands are indeed very long each will fold back on itself to lower the energy because points x,x and y,y both lie in the net attractive region.

Consequently we must conclude that if the degree of polymerization is sufficiently high, the ground-state structure of all toy model proteins must be folded.

IV. GROUND STATES

For any number n of residues, and for any given sequence of those n residues specified by $\xi_1,..,\xi_n$, the potential-energy function Φ is precisely defined, and in principle can be minimized with respect to the conformational angles $\theta_2,\ldots,\theta_{n-1}$. In practice this is easy for small n (i.e., 3,4, and 5), but becomes increasingly tedious and demanding as n increases. Our goal has been to generate and interpret a database of ground states that is *complete* through the maximum feasible value of n, that is to say all molecular ground states were to be obtained for this or a smaller number of residues.

The search for ground states has been carried out through heptamers ($n=7$). A mixed strategy was used, relying both on steepest-descent minimizations on the Φ hypersurfaces from a wide variety of random and "hand-picked" starting points, as well as a Monte Carlo simulated annealing procedure. Repeated application of these approaches has provided checks on results, and we feel confident that indeed all $n\leq 7$ ground states have been correctly identified.

The simple case of the trimers ($n=3$) provides an introductory illustration. The six distinct molecules are *AAA*, *AAB*, *ABA*, *ABB*, *BAB*, and *BBB*. Each has only a single bend degree of freedom θ_2. Furthermore, the potential energy as specified in Sec. II above depends only on the species of the terminal residues, and not on that of the central residue. Consequently there are just three distinct cases to consider: *AXA*, *AXB*, and *BXB*. Figure 3 shows the potential-energy curves vs θ_2 for each of these.

It is obvious from Fig. 3 that the *AXB* and *BXB* trimers are linear in their ground states. That is certainly expected for *AXB*, where the terminal residues repel each other at all separations. Even though modest terminal residue attraction exists for *BXB*, the bend potential energy is sufficiently costly that the possibility of an absolutely stable bent shape is eliminated. Only when both terminals are *A* is the nonbonded interaction sufficiently at-

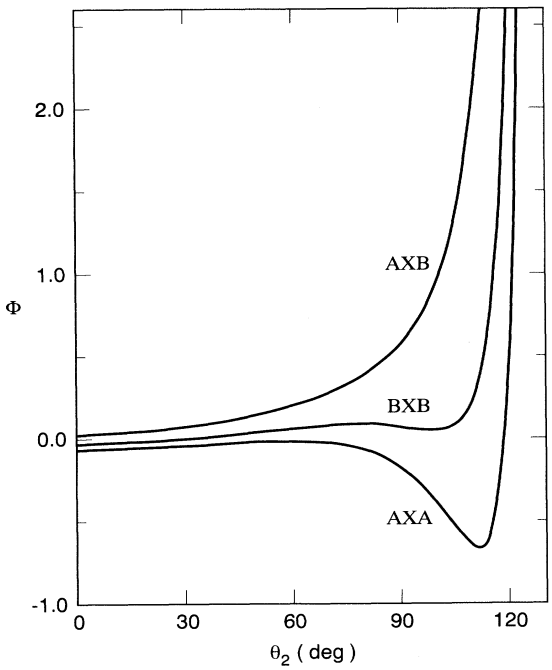

FIG. 3. Bend potential-energy curves for trimers. Results are independent of the central residue's species, denoted here by *X*.

tractive to generate a bent ground state; θ_2 is approximately $\pm 111.4°$ in this nonlinear structure. But notice also that the AXA molecules retain the linear form as a metastable (relative) Φ minimum, and *BXB* molecules have a bent metastable minimum; this is the first appearance of a multiple minimum problem that magnifies dramatically in severity as the molecules increase in residue number.

Table I collects and displays ground-state energies and backbond bend angles for all trimers (6), tetramers (10), and pentamers (20). Several protocols control the presentation in Table I in order to simplify and standardize the information:

(a) Molecules are listed in alphabetical order for each number of residues, and, in the case of sequences differing only by reversal, only the first in alphabetical order appears. Thus *AAAB* appears in the tetramer grouping while *BAAA* does not; and *AAAB* precedes *AABA*.

(b) The residues are numbered sequentially from left to right (e.g., $A_1B_2A_3B_4B_5$).

(c) Angles in radians are normalized by π for convenience of applications to the neural networks discussed in Sec. V.

(d) In some cases, the ground states appear as mirror-symmetric pairs, or even quadruplets. The bent *AXA* trimers supply the simplest examples of pairs, *AAAA* being an example of the quadruplet. We cope with this degeneracy by consistently choosing the potential-energy minimum with the most positive θ_2 value.

Extension of Table I to include hexamers (36) and heptamers (72) would be unwieldly. However, the information is available from the authors on request.

Notice that the fraction of molecules with linear ground states ($\theta_2=\theta_3=\cdots=\theta_{n-1}=0$) declines with an increasing number of residues. These represent $\frac{2}{3}$ of the trimer entries in Table I, $\frac{2}{5}$ of the tetramer entries, and $\frac{1}{10}$ of the pentamer entries. Furthermore, only a single linear ground state appears among the hexamers ($\frac{1}{36}$), and none among the heptamers. In connection with the argument presented in Sec. III B above, it appears likely that no linear ground states for any sequence will be found for $n>6$. In spite of this, we find that every sequence has the linear structure as a local Φ minimum.

Although only two distinct shapes, bent and linear, appear among the trimer ground states, the tetramers present far greater diversity. Figure 4 illustrates this point, distinguishing four families: "linear," "symmetric globular," "asymmetric globular," and "switchback." Even greater diversity (Fig. 5) appears in the pentamer group (a trend that continues for hexamers and heptamers). In many cases we have found that the molecules experience interfamily transitions upon excitation from the ground state to metastable states. Note also that

TABLE I. Ground-state properties of toy-model polypeptides. Angles θ_i are measured in radians. Hexamer and heptamer results are not listed, but are available upon request from the authors.

Molecule	Φ	θ_2/π	θ_3/π	θ_4/π
AAA	−0.658 21	0.618 66		
AAB	0.032 23	0.000 00		
ABA	−0.658 21	0.618 66		
ABB	0.032 23	0.000 00		
BAB	−0.030 27	0.000 00		
BBB	−0.030 27	0.000 00		
AAAA	−1.676 33	0.618 39	0.339 20	
AAAB	−0.585 27	0.617 59	−0.051 30	
AABA	−1.450 98	0.332 70	0.621 80	
AABB	0.067 20	0.000 00	0.000 00	
ABAB	−0.649 38	0.617 67	−0.066 70	
ABBA	−0.036 17	0.476 90	0.476 90	
ABBB	0.004 70	0.000 00	0.000 00	
BAAB	0.061 72	0.000 00	0.000 00	
BABB	−0.000 78	0.000 00	0.000 00	
BBBB	−0.139 74	0.558 28	0.351 80	
AAAAA	−2.848 28	0.335 97	0.620 22	0.045 43
AAAAB	−1.589 44	0.618 98	0.337 48	−0.068 94
AAABA	−2.444 93	0.297 23	0.333 06	0.621 76
AAABB	−0.546 88	0.617 56	−0.053 73	−0.001 68
AABAA	−2.531 70	0.329 43	0.623 54	0.045 51
AABAB	−1.347 74	0.332 69	0.621 33	−0.545 74
AABBA	−0.926 62	0.167 22	0.482 28	0.473 27
AABBB	0.040 17	0.000 00	0.000 00	0.000 00
ABAAB	−1.376 47	0.622 22	0.331 10	−0.063 03
ABABA	−2.220 20	0.619 00	0.047 39	0.619 00
ABABB	−0.616 80	0.617 65	−0.071 04	−0.002 24
ABBAB	−0.005 65	0.478 80	0.473 41	−0.141 84
ABBBA	−0.398 04	0.245 76	0.555 51	0.245 76
ABBBB	−0.065 96	0.054 89	−0.342 37	−0.561 78
BAAAB	−0.521 08	0.039 24	−0.616 71	0.039 24
BAABB	0.096 21	0.000 000	0.000 00	0.000 00
BABAB	−0.648 03	0.053 28	−0.616 82	0.053 28
BABBB	−0.182 66	0.569 20	0.335 74	0.266 59
BBABB	−0.240 20	0.317 73	0.576 42	0.097 38
BBBBB	−0.452 66	0.343 45	0.565 01	0.093 18

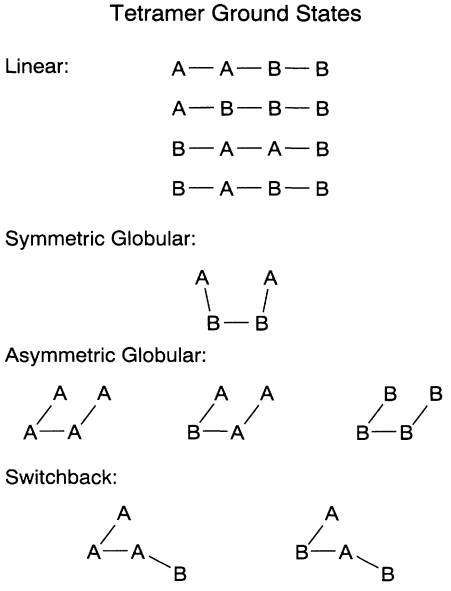

FIG. 4. A schematic illustration of the shape families of the tetramer ground states.

point mutations (single monomer change) can either preserve or drastically alter ground-state structure.

V. NEURAL NETWORKS

We have drawn upon the neural-network concept to aid in the interpretation of the ground-state database reported in Sec. IV. Specifically we have considered static feed forward networks [9], and have required that they perform perfectly as read-only memory devices for the backbone bend angles of the toy polypeptides. For each complete group of n-mers, subject to the protocols invoked in Sec. IV, inputting the binary sequence code $\xi_1, \ldots, \xi_n$ must produce the correct set of normalized angles $\theta_2/\pi, \ldots, \theta_{n-1}/\pi$ as output.

The "neurons" which compose our network all possess the same "clipped linear" response function:

$$R(x)=\begin{cases}-1 & (x \le -1) \\ x & (-1<x<1) \\ 1 & (1 \le x)\ .\end{cases} \tag{5.1}$$

Neuron 1 receives inputs I_{kl} from sources k (either externally supplied ξ_i's, or outputs of neurons preceding it in the network); each of these is delivered with weight w_{kl}. The output of neuron 1 is

$$O_1 = R\left[\sum_k w_{kl} I_{kl} + b_l\right], \tag{5.2}$$

where b_1 is the bias for that neuron. The last layer of the network will contain $n-2$ neurons, each with the task of delivering the correct value of a θ_i/π as its output.

While surveying alternative architectures (i.e., connection schemes) for perfectly performing neural networks, it is useful to assign a figure of merit f to each candidate. This is simply the number of neurons plus the number of interconnections. It is thus also the number of weights w_{kl} and biases b_1 that must be determined to make the network perform in an error-free manner. Note that a connection with vanishing weight is no connection at all and does not contribute to f. We have been interested in determining optimal neural networks for each n-mer group, namely those with the smallest possible value of f.

The optimal network for trimers is easily found, and is almost trivial. It has $f=2$, and is illustrated in Fig. 6. There the single neuron required has been symbolized by a circle, and the quantity within that circle is its bias $b=\theta^*/2\pi$, with θ^* denoting the stable bend angle for the nonlinear AXA trimers. The single input channel to the neuron carriers weight $w=\theta^*/2\pi$. As a result of the data protocol, θ_2 has a value 0 or θ^* depending *only* on the residue species variable ξ_3, and the network graphically illustrates this point by failing to connect to the ξ_1 and ξ_2 input nodes.

It must be stressed that we only require the networks to perform properly for the sets of input sequences $\xi_1, \ldots, \xi_n$ admitted by the imposed protocol. The numbers $N(n)$ of these allowed sequences, Eqs. (3.1) and (3.2), are substantially smaller than 2^n, the numbers of unrestricted sequences. Indeed it is clear that even the simple trimer network in Fig. 6 will "misperform" if fed a disallowed sequence (such as -1, 1, 1). However, this represents no loss of basic folding information.

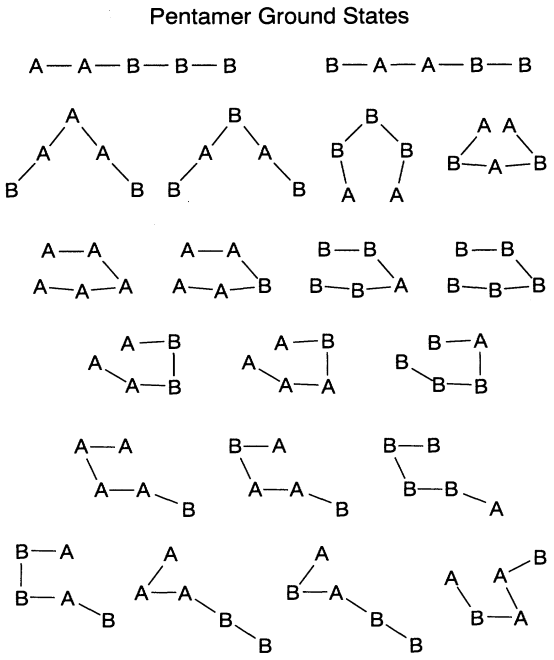

FIG. 5. Shape families of the pentamer ground states.

Although one expects perfectly performing networks to be more complicated for larger n, at least assurance exists that *some* finite architecture is available for every n. The Appendix presents the argument, and leads to an upper bound for the optimal figure of merit:

$$f_{\rm opt}(n) \le (n+4)N(n)+n-2\ . \tag{5.3}$$

Our limited experience indicates that $f_{\rm opt}$ tends to be considerably smaller than this bound; for trimers, we have seen $f_{\rm opt}(3)=2$, whereas (5.3) merely specifies $f_{\rm opt}(3) \le 43$.

The searches for optimal tetramer and pentamer networks have employed a combination of Monte Carlo pro-

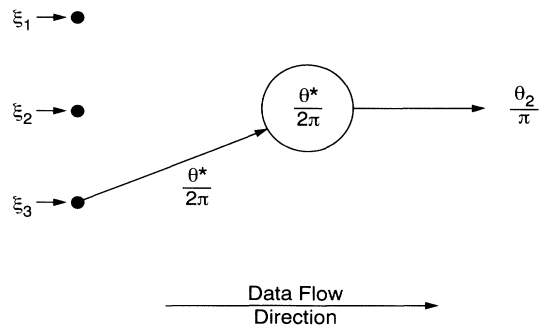

FIG. 6. The optimal network for trimers (figure of merit $f=2$). The stable bend angle for AXA trimers (111.4°, 1.9436 rad) has been denoted by θ^*.

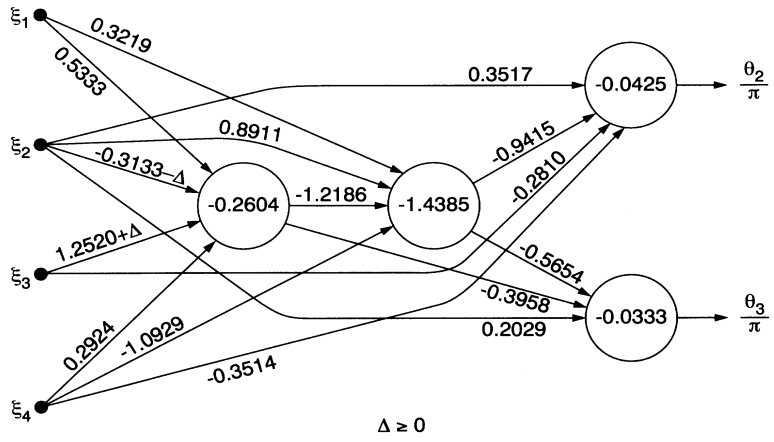

FIG. 7. The optimal tetramer network, with $f=19$. Notice the invariance with respect to Δ, subject to the inequality shown.

cedures. These have started with a variety of candidate network topologies, and have respectively generated approximately 10^2-10^3 random sets of weights and biases. For each, small random changes in weights and biases were produced, with a test to see if the rms difference between the database and the network predictions of all elements of the database was thereby reduced. This process was continued until either a perfectly performing network emerged (rms error below 10^{-9}), or trapping in a local rms minimum with unacceptable error occurred. The majority of cases suffered the latter fate. In those few cases that were successful, attempts to reduce f were initiated, pruning out weakly weighted connections one by one with adjustments of remaining weights and biases to restore perfect performance. This evolutionary streamlining would eventually terminate, as further pruning became impossible.

Figure 7 displays the simplest tetramer network discovered by this search. Although we have no proof, we believe it is optimal. It possesses two "hidden layers," each of a single neuron, interposed between the four input nodes and the two output neurons. The figure of merit $f=19$ may be compared with the upper bound of 82 given by Eq. (5.3). This dramatic reduction in f below that value for the banal "overkill" network (see Appendix) lends credence to the proposition that optimal networks image the fundamental logic of folding.

Weights and biases in the tetramer network have been determined to at least ten significant figures of accuracy by the rms reduction procedure. However, they have been rounded off to four decimal places for convenience in Fig. 7. Notice the degeneracy regarding connection weights from the ξ_2 and ξ_3 input nodes to the first hidden-layer neuron: an arbitrary positive number Δ can, respectively, be subtracted from and added to these connection weights without changing the error-free network performance. The existence of this invariance stems from the nonlinearity of the neural response function R, Eq. (5.1).

Bend angles θ_i never attain values $\pm\pi$ on account of the r^{-12} term in V_2, Eq. (2.5). Consequently, output neurons always operate in their linear ramp regimes. By contrast, hidden-layer neurons are frequently saturated (i.e., in the flat response regimes). As the tetramer network in Fig. 7 runs through the ten independent allowed inputs, the first hidden-layer neuron is saturated in six cases, the second hidden-layer neuron is saturated in five.

The simplest pentamer networks uncovered by our Monte Carlo search and refinement exhibit the generic architecture indicated in Fig. 8. Three hidden layers, each composed of two neurons, reside between the five input nodes and three output neurons. If all possible connections were present in this design, the figure of merit would be $f=84$. Once again this lies well below the upper bound in Eq. (5.3), namely 183.

A tedious search for the lowest figure of merit in this family of pentamer networks has reduced f to 70 by pruning out 14 connections. This apparent optimum is attained by a continuous subfamily of networks with parametric variation of weights and biases, analogous to that found for tetramers. But in contrast to that former case, the variations affect all but six weights and one bias, and do not appear to have a simple linear representation as before. Table II provides the weights and biases for one subfamily member.

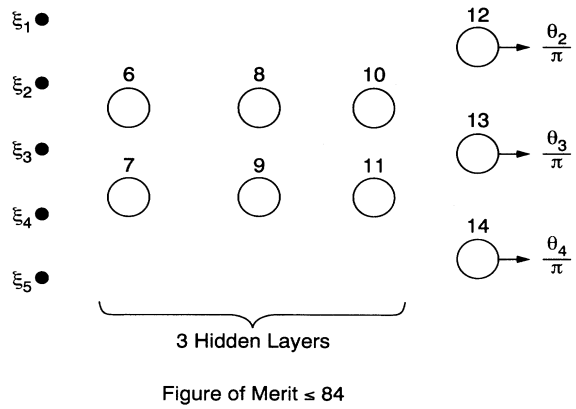

FIG. 8. The generic architecture of the simplest pentamer networks discovered.

VI. CONCLUSIONS

By limiting attention just to two dimensions, and to two kinds of monomers ("amino acids"), it has been possible to formulate a very simple but useful toy model for protein folding. A wide variety of ground-state structures arise, dependent on the number of monomers present and on their sequence. This diversity stems from the competition between the backbone bend portion of the potential energy and the nonbonded interactions between pairs of amino-acid residues. A complete database of ground-state energies and structures has been created for all model polypeptides with seven or fewer residues.

Neural networks offer a way to analyze and to visualize the folding database for the toy model. Specifically we have created static feed-forward networks that act as error-free read-only memories; upon being fed the residue sequence as binary input, they yield the correct backbone bend angles for that case. These networks are not unique, but we have stressed the importance of those with optimal architecture, namely a minimum number of neurons plus interconnections.

Neural networks thought to be optimal for trimers, tetramers, and pentamers were presented in Sec. V. We have proposed that their topologies constitute an image, or logical flow diagram, of the basic physical principles that underlie the folding patterns in the model. In particular, they indicate that generally the folding is a nonlocal,

TABLE II. Weights and biases for an optimized pentamer network ($f=70$). The numbering scheme conforms to Fig. 8.

Element	Value	Element	Value
$w(1,6)$	−0.6154	$w(6,8)$	−0.1799
$w(1,7)$	−0.7188	$w(6,9)$	−0.1610
$w(1,8)$	0.7134	$w(6,11)$	−0.9275
$w(1,9)$	0.8246	$w(6,12)$	0.0048
$w(1,10)$	−1.2748	$w(6,13)$	0.7030
$w(1,11)$	−0.2623	$w(6,14)$	0.0965
$w(1,12)$	−0.8640		
$w(1,13)$	−0.2825	$w(7,8)$	0.3389
		$w(7,9)$	0.2747
$w(2,6)$	1.0515	$w(7,10)$	−1.0650
$w(2,7)$	0.3240	$w(7,11)$	−0.3724
$w(2,8)$	−0.1779	$w(7,12)$	0.1606
$w(2,9)$	−0.8761	$w(7,13)$	0.2045
$w(2,10)$	0.8235	$w(7,14)$	−0.7284
$w(2,11)$	0.7100		
$w(2,13)$	−0.3508	$w(8,11)$	−0.6252
$w(2,14)$	−0.0495[a]	$w(8,12)$	0.4413
		$w(8,13)$	0.5943
$w(3,6)$	1.1435		
$w(3,7)$	1.5005	$w(9,10)$	−0.3902
$w(3,8)$	−0.8562	$w(9,11)$	0.4274
$w(3,10)$	−0.6807	$w(9,12)$	0.1303
$w(3,11)$	−0.2358	$w(9,13)$	0.8388
$w(3,14)$	0.6088		
		$w(10,12)$	0.0515
$w(4,7)$	−0.5580	$w(10,14)$	0.6196
$w(4,8)$	−0.2574		
$w(4,9)$	−0.5930	$w(11,12)$	−0.2059
$w(4,10)$	1.1252	$w(11,13)$	0.7112
$w(4,12)$	0.1604	$w(11,14)$	−0.6132
$w(4,13)$	0.4078[a]		
$w(4,14)$	−0.4097	$b(6)$	−0.1143
		$b(7)$	0.2566
$w(5,6)$	−0.6054	$b(8)$	0.3667
$w(5,7)$	0.1978	$b(9)$	1.4137
$w(5,8)$	−0.5262	$b(10)$	0.8513
$w(5,9)$	−0.3006	$b(11)$	−0.2154
$w(5,12)$	0.0954[a]	$b(12)$	0.0735
$w(5,13)$	0.4489[a]	$b(13)$	−0.1862
$w(5,14)$	0.4038[a]	$b(14)$	0.0714[a]

[a]Invariant values over the optimal network subfamily.

highly cooperative process.

Suppose, hypothetically, that the bend angle θ_i were to depend only on the numerical position i along the chain, as well as the species of the residue at that position. Then moving backward through the network from the θ_i/π output location, one would only have a path to the ξ_i input node, and would encounter no branching that would allow access to other input nodes. More generally, one might suppose that θ_i depended on numerical position i, as well as the local species sequence at $i-1$, i, and $i+1$; this would admit path branching from the θ_i/π output to the three successive input nodes for ξ_{i-1}, ξ_i, and ξ_{i+1}, but not others. Disregarding the trivial trimer case, the apparently optimal tetramer and pentamer networks fail to display this simplifying feature, but instead show that each output is connected, if only indirectly, to each input node.

It strikes us as unlikely that optimal networks for hexamers, heptamers, etc. will be any simpler in this input-output connection sense. Whatever the specific details of the hidden-layer architecture, we expect every pair of input and output channels to have at least one direct or indirect connection pathway. The information processing carried out within the optimal error-free network defines and utilizes collective species variables that incorporate information about the entire sequence. This in turn suggests that any procedure which utilizes only local sequence information to predict folded structure is intrinsically imprecise, whether applied to a simple model such as that considered in this paper, or applied to real proteins [10–14].

Recent theoretical literature for protein folding contains other examples of simplified models relying on just two "amino acids." Some of these utilize a discrete lattice space, freely jointed chains, and simplified nearest-neighbor interactions [15,16]; this tends to produce considerable degeneracy in ground states for at least some sequences. Other versions incorporate more elaborate backbone (local) interactions, but retain a lattice space, and nearest-neighbor interactions for nonbonded residue pairs [17].

The elementary toy model studied here is one of the simplest nontrivial possibilities that might have been considered. Obviously it permits extensions or generalizations in any of several directions that would move it closer to the real world of proteins. These include use of more amino acids, more degrees of freedom per residue (both backbone and side chain), and transfer from two to three dimensions. For clarity this should be implemented stepwise, with repeated attention to the architecture of optimal neural networks to reveal the logical structure of folding principles. The end results of a systematic study along these lines should substantially strengthen our insights into the behavior of real proteins.

APPENDIX

Choose a specific n-mer sequence σ permitted by the protocol and denote its binary code by $\xi_{\sigma 1},\ldots,\xi_{\sigma n}$. Its bend angles will be $\theta_{\sigma 2},\ldots,\theta_{\sigma,n-1}$. Then consider the modular network unit illustrated in Fig. 9(a). The connections bear weights shown next to their directed lines, and the biases of the neurons appear inside their symbolic circles. One easily verifies that the output of this modular unit is always 0, except for its special sequence $\xi_{\sigma 1},\ldots,\xi_{\sigma n}$, which stimulates output 1. Hence each modular unit is a detector for its own special sequence. If one such detector is present for each of the $N(n)$ allowed σ's, their outputs can then be supplied simultaneously to each of the $n-2$ output neurons. As Fig. 9(b) indicates, the latter have vanishing biases, and will invariably report the correct θ_i/π if the $N(n)$ input channels carry weights $\theta_{\sigma i}/\pi$ from each of the σ-detector modular units.

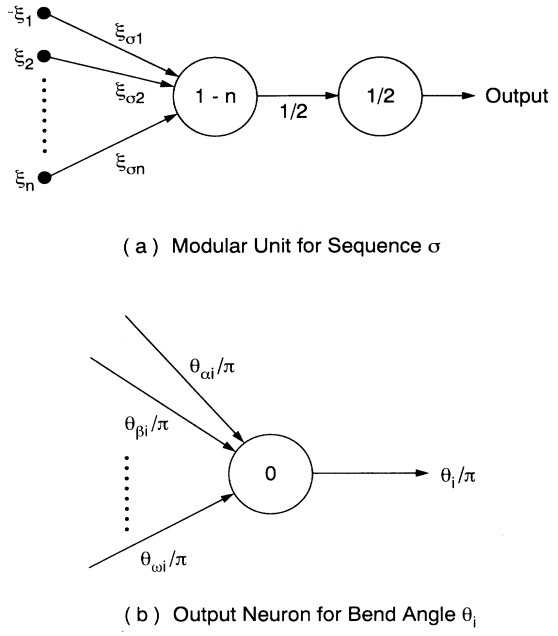

FIG. 9. Structures used for the "overkill" networks.

Each sequence detector requires two neurons, so the total, including the output neurons, is $2N(n)+n-2$ neurons. The number of weighted connections is $n+2$ per detector, hence $(n+2)N(n)$ in all. The corresponding figure of merit is defined to be the sum of these:

$$f=(n+4)N(n)+n-2\ . \tag{A1}$$

The optimal network for n-mers cannot have its f larger than this, so we have written inequality (5.3) in the text to state this fact.

The family of networks just described represents substantial "overkill." Indeed, their operation, though error free by construction, might be viewed as a blind recitation of the given database. The massive architecture embodies no simplifying features based upon the molecular folding principles. Only after stripping an n-mer network down to its optimal figure of merit, can one expect to see the architecture reveal the logic of folding.

*Present address: Donner Laboratory, Lawrence Berkeley Laboratory, Berkeley, CA 94720.
†Present address: Dept. of Physics, Harvard University, Cambridge, MA 02138.

[1] T. E. Creighton, *Proteins, Structures and Molecular Principles* (Freeman, New York, 1984), Chaps. 5 and 6.
[2] *Protein Folding,* edited by L. M. Gierasch and J. King, (AAAS, Washington, D.C., 1990).
[3] B. T. Nall and K. A. Dill, *Conformations and Forces in Protein Folding* (AAAS, Washington, D.C., 1991).
[4] E. Ising, Z. Phys. **31**, 253 (1925).
[5] W. Heisenberg, Z. Phys. **49**, 619 (1928).
[6] J. Hertz, A. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computation* (Addison-Wesley, Redwood City, CA, 1991).
[7] W. G. McMillan and J. E. Mayer, J. Chem. Phys. **13**, 276 (1945).
[8] T. L. Hill, *Statistical Mechanics* (McGraw-Hill, New York, 1956), pp. 262–277.
[9] Ref. 6, p. 90.
[10] H. Bohr, J. Bohr, S. Brunak, and R. M. J. Cotterill, FEBS Lett. **261**, 43 (1990).
[11] J. D. Hirst and M J. E. Sternberg, Protein Eng. **4**, 615 (1991).
[12] N. Qian and T. J. Sejnowski, J. Mol. Biol. **202**, 865 (1988).
[13] L. H. Holley and M. Karplus, Proc. Natl. Acad. Sci. USA **86**, 152 (1989).
[14] D. G. Kneller, F. E. Cohen, and R. Langridge, J. Mol. Biol. **214**, 171 (1990).
[15] K. F. Lau and K. A. Dill, Macromolecules **22**, 3986 (1989).
[16] K. Yue and K. A. Dill, Proc. Natl. Acad. Sci. USA **89**, 4163 (1992).
[17] A. Sikorski and J. Skolnick, Proc. Natl. Acad. Sci. USA **86**, 2668 (1989).