Reports

Observation of Phase Transitions in Spreading Activation Networks

Jeff Shrager, Tad Hogg, Bernardo A. Huberman

Phase transitions, similar to those seen in physical systems, are observed in spreading activation networks. Such networks are used both in theories of cognition and in artificial intelligence applications. This result confirms a predicted abrupt behavioral change as either the topology of the network or the activation parameters are varied across phase boundaries.

CCO PREADING ACTIVATION" REFERS to a class of algorithms that propagate numerical values (activation levels) in a network for the purpose of selecting the nodes that are most closely related to the source of the activation. This process of tracing chains of connections can be found in early psychological theories dating back to the associationist models of thought attributable to Aristotle, in experimental psychology research, and in the theories of Freud and Pavlov (1). Quillian (2), who introduced spreading activation as a computational process, has shown how search in a semantic network can be accomplished by the use of this technique. More recently, it has been applied in various artificial intelligence systems (2-4) and is a component of a number of computational models of memory in cognitive psychology (1, 3,5).

The networks in which spreading activation is thought to occur (human memory, for example) contain very many nodes. There exist proposals to build very large databases that incorporate this technique (6), and various proposed applications of the Connection Machine (7) involve spreading activation through extremely large networks. However, real computational experiments are usually conducted with small networks from which the properties of the larger cases are then extrapolated. The few theoretical analyses of very large nets that have been attempted (1) are normally restricted to highly regular topologies with dubious simplifying assumptions (one-dimensional networks, for example). These simplified experiments and standard analyses assumed that the findings continue to hold when the systems are scaled up without bound. For this assumption to be true, the range of causal interactions between nodes in space-time, which defines an "event horizon" in these nets, must remain relatively small. Recently Huberman and Hogg (8) predicted that as parameters such as the number of links in a large network are modified, highly nonlinear behavioral changes can lead to sudden explosive growth in the size of the event horizon. As in physical phase transitions (9), these global changes are characterized by macroscopic singularities. This in turn implies that any system whose functioning can be construed as spreading activation should show explosive changes in its event horizon as the predicted phase boundaries are approached. This is a reflection of the fact that for large systems, topology is a dominant determinant of behavior, irrespective of local computational details.

In addition to the aforementioned applications, spreading activation is used in psychological models as a mechanism for deciding which of several possible related memories or actions should be selected in ambiguous situations. For instance, the Adaptive



Fig. 1. Predicted phase diagram for a simple spreading activation net (assuming no limits or bounds). Dashed lines represent the experimental trajectories investigated: the horizontal line for experiments 1 and 2, and the vertical lines for 3 and 4.

Control of Thought (ACT*) model of Anderson (1) uses spreading activation in two related ways: to select items from a memory arranged in a network of nodes (concepts) and links between them; and to select among competing behaviors represented by production rules in performance situations. In addition, Anderson's associative model of memory includes learning. Specifically, declarative learning can take the form of the addition of new links between existing memory items represented by nodes. Procedural learning can take the form of new connections in a network whose nodes represent parts of the rules of performance for skills that the person has acquired. Thus, both of these sorts of learning can be viewed in terms of graph dynamics.

If the predicted topological phase transitions are confirmed by experimental test, they have particular implications for cognitive theories of learning and memory, and for computational applications (especially in artificial intelligence), which depend on the dynamic behavior of a network whose connectivity is changing. This report describes four computational experiments that confirm the existence of such phase transitions in spreading activation networks, and it analyzes the corrections due to their finite size.

Spreading activation networks typically consist of a set of potentially active nodes representing various states or items, variously interconnected by weighted links, and a local relaxation rate at which the activity of an isolated node decays to zero. The dynamic behavior of these networks is controlled by three parameters. The first, specifying their topology, is the average number of links per node, μ . The second, α , is a positive number that describes the relative amount of activity that flows from a node to its neighbors per unit time. The third parameter is the relaxation rate γ , which can have a value between 0 and 1. In a typical application, some nodes (the sources) are activated by external inputs and these in turn cause others to become active with varying intensities.

Computationally, the spread of activation takes place in discrete time steps. Let A(N)be a vector whose *i*th element is the activation of the *i*th node at time step N, let C(N)be a vector whose elements specify the external inputs at the same time step. In the standard model of activation plus relaxation (1), the time evolution of the net is given by

$$\mathbf{A}(N) = \mathbf{C}(N) + (1 - \gamma)\mathbf{A}(N - 1) + \alpha \mathbf{R} \mathbf{A}(N - 1)$$
(1)

Intelligent Systems Laboratory, Xerox Palo Alto Research Center, Palo Alto, CA 94304.



Fig. 2. Experimental results from networks with 100 nodes, as the number of links is changed from 20 to 100 (path 1&2 on phase diagram $\mu = 0.4$ to 2.0; $\alpha = 0.4$, $\gamma = 0.6$). Curves show the topological transition involving the growth in the number of affected nodes. Each curve represents 25 trials with e = 0.001 and B = 10 (solid curve) or B = 1000 (dashed curve). (**A**) Fraction of active nodes; (**B**) relative fluctuations. The arrow shows the predicted transition point on the phase diagram in Fig. 1.

where \boldsymbol{R} is a matrix with diagonal elements of value zero and off-diagonal elements R_{ii} that represent the weight of the link from node *j* to node *i*, which is defined to be zero for nonexisting links. We consider the case where the activation from each node is divided among the attached nodes according to the weight of their connections. This means that those columns of **R** that have any nonzero elements sum to one. Columns that are all zero correspond to nodes with no links. This procedure for spreading activation through a net is analogous to current flow in an electrical network. Thus the transitions discussed here are similar to percolation transitions seen in these electrical networks (9) but occur in graphs of arbitrary dimensionality.

Sudden phase changes manifest themselves in explosive variations in the range of interactions between nodes of the net as well as in the relaxation rate of the net toward equilibrium. For systems with infinitely many nodes and unbounded growth of activation, the various possible phases predicted by the Huberman and Hogg theory (8) are shown in Fig. 1. First, consider the case where α/γ is small (region I). In this regime, the activation of the net relaxes quickly to its asymptotic value and it is localized in space. This implies that events at the source node do not significantly affect the activation of the far regions of the net. Moreover, they will have little effect on the source itself at later times. Because of its limited extent in time, this behavior can be thought of as taking place in a finite temporal cluster.

For computational purposes, the existence of such a phase implies that one can effectively ignore a priori both far regions of the net and its ancient history. Thus, in this regime the event horizon affecting a given node at a given time is localized in both space and time. The existence of these localized event horizons means that simple analyses that rely on the ability to ignore far regions of the net (spatial localization) as well as eventual equilibration (time localization) will be applicable in this phase.

As α/γ increases, the overall relaxation of the net toward its stable fixed point becomes increasingly sluggish, with its characteristic relaxation time diverging to infinity as the transition is approached. This leads to sharp phase transitions into phases II and III (Fig. 1). In phase II, one encounters a regime where the event horizon grows indefinitely in time but remains localized in space. This means that unlike phase I, ancient history matters in determining the activation of any node, and that since the net reverberations never settle down (that is, the activity keeps increasing) the assumption of an equilibrium between the net and the time variations at the source no longer holds.

Finally in region III of Fig. 1 the event horizon is extended not only in time but also in space. This means that the amount of spreading grows indefinitely, and therefore far regions of the net can significantly affect each other. This region is separated from phase II by a sharp transition in which the number of nodes with activation values above a given positive threshold grows explosively.

The finite size of any real network, as well as the limited time over which experimental observations are made, introduce corrections to this theory that result in a smoothing of the transitions. For example, slightly below the phase boundaries observations might not last long enough to distinguish between continued activity growth and the eventual reaching of the fixed point. Moreover, near the transition points, relative fluctuations in behavior become very large, which can obscure the identification of the average behavior predicted by the theory. On the other hand, we show below how this can be exploited in identifying the phase transition.

We have examined the applicability of the Huberman and Hogg theory to actual computational situations in four experiments, each consisting of 25 spreading activation trials. In each trial we randomly generated links in a graph containing 100 nodes. Spreading activation is computed according to Eq. 1. One node in the network was selected as a constant source. Thus, C(N) was a constant vector with a single nonzero element. The total activation of the network, which is the sum of the activities in all of the nodes, as well as the number of activated nodes were recorded at each iteration.

Two additional parameters were used to



Fig. 3. Experimental results for paths 3 and 4 on the phase diagram (that is, increasing α/γ from 2/3 to 1.5). Curves show the dynamic transition involving growth of the fixed point and sluggishness of relaxation. Each curve represents 25 trials with e = 0.001, B = 100, and the number of links = 20 (solid, $\mu = 0.4$) or 70 (dashed, $\mu = 1.4$). (A) Total activity; (B) completion time. The arrow shows the predicted phase transition point.

decide on the eventual behavior of the network. The network was considered to have settled into a fixed point when the total activity on two consecutive iterations differed by less than the settling threshold e, which was set to the value 0.001 in all these experiments. An upper bound, B, determined when the network failed to settle. Computation was stopped when the total activity in the network either settled or reached this bound. In practice, we used relatively small networks and values of B, thus limiting the number of iterations. Since limiting the total number of iterations smoothes the expected transition we varied B in order to examine its effect.

The results of the first two experiments are depicted in Fig. 2. They explore left-toright movement on the phase diagram as the value of μ increases. This path corresponds to the trajectory labeled 1&2 in Fig. 1. For both of these experiments $\alpha = 0.6$ and $\gamma = 0.4$, so $\alpha/\gamma = 1.5$. In experiment 1 (solid curve), B = 10 and in experiment 2 (dashed curve), B = 1000. Figure 2A shows the extent of the spatial event horizon by displaying the average fraction (over the 25 trials) of active nodes in the network at the time the network either settles or reaches the value B. As expected from the limited size of our network, the transition is smoothed out. To highlight the transition, we examined the relative fluctuations in the network, computed by dividing the standard deviation of the fraction of active nodes by their average. This quantity, shown in Fig. 2B, peaks at approximately the predicted point of 50 links ($\mu = 1$).

The results depicted in Fig. 3 are from the remaining two experiments. We investigated movement in the vertical direction on the phase diagram, as shown by transitions 3 and 4 of Fig. 1. In both of these cases, α and γ begin at 0.4 and 0.6, respectively; α is increased by 0.02 while γ is decreased by the same amount, until their ratio reaches the value 1.5. The vertical transition was tried at two different values of μ , 0.4 and 1.4, corresponding to 20 links (solid curve) and 70 links (dashed curve). Figure 3A shows the total activity of the network when the network either settles or reaches the bound B = 100.0. One can clearly see the predicted sudden growth in activity when $\alpha/\gamma = 1.0$. Below this point the network always settles. Above it, the activity continues to grow and eventually reaches B. The difference in the upper part of these curves (about 30 for the solid curve and about 75 for the dashed curve) is a result of the overall greater number of connected nodes in the latter case.

Figure 3B shows the time taken for the network to either settle or reach the value B.

In this figure we have separated each curve into the settling region (left) and bounded region (right). The peak is a result of different phenomena on either side. On the left, the curve results from the time required for the network to settle to its fixed point, whereas on the right instead of settling the network continues to grow faster as α/γ increases, reaching the value *B* more rapidly.

The theory developed by Huberman and Hogg (8) provides a way of understanding the robust global features of a system. Such features, which are hard to notice in small systems, can play a large role in the behavior of larger ones. Our results clearly show that the predicted phase transitions can be observed even in relatively small spreading activation networks. Moreover, the existence of such transitions has immediate implications for the predictions of memory models and the behavior of artificial intelligence systems which incorporate learning. Specifically, instead of the fairly smooth transitions in performance that have been generally assumed in these situations, we have shown that abrupt transitions can be expected. More generally, these experiments show that statistical models provide a useful way to understand the behavior of large

systems. They also emphasize the dominating influence of topological properties. These are particularly important implications for the behavior of any network with a dynamic topology.

REFERENCES AND NOTES

- 1. See J. R. Anderson [*The Architecture of Cognition* (Harvard Univ. Press, Cambridge, MA, 1983)] for a discussion and analysis of spreading activation. Also see J. R. Anderson and G. H. Bower [Human Associative Memory (Lawrence Earlbaum, Hillsdale,
- 1968). Spreading activation is used to compare and contrast word-senses in a network-structured dictioiary database
- nary database.
 A. Howe, in Proceedings of the Canadian Society for Computational Studies of Intelligence (London, On-tario, 1984), pp. 25-27.
 P. R. Cohen and P. M. Stanhope, in Proceedings of the 6th International Workshop on Expert Systems and their Applications, Avignon, France, 28-30 April 1986. Spreading activation is used to find suggested granting agencies to which proposals can be sent.
 A. M. Collins and E. F. Loftus, Psychol. Rev. 82, 407 (1975).
- (1975)
- 6. D. Lenat, M. Prakash, M. Shepard, Artif. Intell.
- D. Hendy, M. Hakash, M. Ghard, Arug. Inite. Mag. 6, 65 (Winter 1986).
 D. Hillis, The Connection Machine: Computing for the New Wave (MIT Press, Cambridge, MA, 1986).
 B. A. Huberman and T. Hogg, Artif. Intell., in
- 9. G. Deutscher, R. Zallen J. Adler, Eds., Percolation Structures and Processes, Annals of the Israeli Physical Society (Hilger, Bristol, 1983), vol. 5.
- 10. We have benefited from discussions with A. Howe. 27 October 1986; accepted 10 April 1987

Single-Channel and Genetic Analyses Reveal Two Distinct A-Type Potassium Channels in Drosophila

CHARLES K. SOLC, WILLIAM N. ZAGOTTA, RICHARD W. ALDRICH*

Whole-cell and single-channel voltage-clamp techniques were used to identify and characterize the channels underlying the fast transient potassium current (A current) in cultured myotubes and neurons of *Drosophila*. The myotube (A_1) and neuronal (A_2) channels are distinct, differing in conductance, voltage dependence, and gating kinetics. The myotube currents have a faster and more voltage-dependent macroscopic inactivation rate, a larger steady-state component, and a less negative steady-state inactivation curve than the neuronal currents. The myotube channels have a conductance of 12 to 16 picosiemens, whereas the neuronal channels have a conductance of 5 to 8 picosiemens. In addition, the myotube channel is affected by Shaker mutations, whereas the neuronal channel is not. Together, these data suggest that the two channels are separate molecular structures, the expression of which is controlled, at least in part, by different genes.

HE OPPORTUNITY TO COMBINE GEnetic, molecular, and single-channel analyses makes Drosophila an ideal preparation for the study of ion-channel function. By analyzing the effects of genetic mutations on the gating behavior of single ion channels, the influence of small changes in structure on channel function can be examined. A number of behavioral mutations proposed to affect ion channels have been isolated (1-3). Mutations at the Shaker locus alter or eliminate A-type potassium currents in larval and adult muscle cells, suggesting that the mutation lies in a structural gene for the A-current channel (4-7). Shaker mutations also cause hyperexcitability of larval presynaptic terminals (8, 9) and abnormally long action potentials in the adult giant axon (10), effects that are mim-

Department of Neurobiology, Stanford University Med-ical School, Stanford, CA 94305.

^{*}To whom correspondence should be addressed.