Complex Systems: Network Thinking

Melanie Mitchell

Portland State University and Santa Fe Institute

(To appear in *Artificial Intelligence*.)

> I am convinced that the nations and people who master the new sciences of complexity will become the economic, cultural, and political superpowers of the next century. —Heinz Pagels [50]
>
> When I hear the word "complexity," I don't exactly reach for my hammer, but I suspect my eyes narrow. It has the dangerous allure of an incantation, threatening to acquire the same blithe explanatory role that "adaptation" once did in biology. —Philip Ball [1]

1 Introduction

Like artificial intelligence in its earlier years, the field of complex systems has received a lot of press in the last two decades, both positive and negative, breezily hyped and unjustly derogatory. As in the early years of AI, complex systems scientists have at times made promises and set up expectations that are hard for any young science to fulfill. Is it true that the "laws of complexity spontaneously generate much of the order of the natural world" [28]? Is the fact that simple programs can produce highly complex behavior "one of the more important single discoveries in the whole history of theoretical science" [66]? On the other hand, is it fair for a major science magazine to ask, on its front cover, if complexity science is "a sham" [53]?

In spite of the hype and the slurs, it turns out that some very interesting and worthwhile work has been done in complex systems over the last several years, evidenced not only by the many books, journals, conferences, and even whole institutes devoted to the field, but also by real changes in the way people approach science, medicine, economics, business, and even the study of intelligence in natural and artificial systems. Computer modeling of complex systems has become widely accepted as a valid scientific activity. "Systems biology," in which complex-systems ideas are central, has become a keyword of the post-human-genome-project

Email address: mm@cs.pdx.edu (Melanie Mitchell).

era. The pursuit of general laws governing complex networks is now center stage in fields ranging from search engines to sociology.

In this article, I discuss some recent ideas in complex systems on the topic of networks, contained in or inspired by three recent complex systems books that seem particularly relevant for people in AI. The general science of networks is the subject of Albert-László Barabási's *Linked* [2] and Duncan Watts' *Six Degrees* [63]. Commonalities among complex biological networks—e.g., immune systems, social insects, and cellular metabolism—and their relation to intelligence in computational systems are explored in the proceedings of a interdisciplinary conference on "Distributed Autonomous Systems" [57].

The ideas discussed in the third book have led to me to propose four general principles of adaptive information processing in decentralized systems. These principles, and the relevance of "network thinking" for AI (and vice versa), are the subject of the last two sections of the article.

2 Complex Systems

Often-cited examples of complex systems in nature and society include the brain, the immune system, biological cells, metabolic networks, ant colonies, the Internet and World Wide Web, economic markets, and human social networks.

There is no generally accepted formal definition of "complex system." Informally, a complex system is a large network of relatively simple components with no central control, in which emergent complex behavior is exhibited. Of course, the terms in this definition are not rigorously defined. "Relatively simple components" means that the individual components, or at least their functional roles in the system's collective behavior, are simple with respect to that collective behavior. For example, a single neuron or a single ant are complicated entities in their own right. However, the functional role of these single entities in the context of an entire brain or an entire colony is relatively simple as compared with the behavior of the entire system.

"Emergent complex behavior" is tougher to define. Roughly, the notion of emergence refers to the fact that the system's global behavior is not only complex but arises from the collective actions of the simple components, and that the mapping from individual actions to collective behavior is non-trivial. The notion of nonlinearity is important here: the whole is more than the sum of the parts. The complexity of the system's global behavior is typically characterized in terms of the patterns it forms, the information processing that it accomplishes, and the degree to which this pattern formation and information processing are adaptive for the system—that is, increase its success in some evolutionary or competitive context. In characterizing behavior, complex-systems scientists use tools from a variety of disciplines, including nonlinear dynamics, information theory, computation theory, behavioral psychology, and evolutionary biology, among others.

The field of complex systems seeks to explain and uncover common laws for the emergent, self-organizing behavior seen in complex systems across disciplines. Many scientists also believe that the discovery of such general principles will be essential for creating artificial life and artificial intelligence.

Complex systems, as their name implies, are typically hard to understand. Traditionally the more mathematically oriented sciences such as physics, chemistry, and mathematical biology have concentrated on simpler model systems that are more tractable via mathematics. The rise of interest in understanding general properties of complex systems has paralleled the rise of the computer, because the computer has made it possible for the first time in history to make more accurate models of complex systems in nature.

3 Networks

In recent years there has been a strong upsurge in the study of networks in many disciplines, ranging from computer science and communications to sociology and epidemiology. The mathematical study of networks arose from graph theory, which began as early as the eighteenth century with Euler's solution to the famous "Bridges of Königsberg" problem. In the 1950s Erodös and Rényi did influential work on the theory of random graphs. However, until recently, mathematical graph theory did not have a large impact on the study of real-world networks, since the latter's properties were quite different from those of random graphs. Quantitative sociologists have been studying social networks for a long time. In the last 10 years or so a number of applied mathematicians and physicists have joined them in developing a general theory of networks.

Two recent books that describe this work are *Six Degrees: The Science of a Connected Age* by Duncan Watts [63], and *Linked: The New Science of Networks* by Albert-Lázló Barabási [2]. Both are books for the "educated lay reader." Along with Mark Newman, Watts and Barabási together edited a more technical collection of recent papers in network theory: *The Structure and Dynamics of Networks* [46]. In this and the following two sections I sketch a few of the recent findings on networks that may be of particular interest to AI practitioners.

3.1 Examples of networks and questions about networks

A network (or graph) is simply a collection of nodes (vertices) and links (edges) between nodes. The links can be directed or undirected, and weighted or unweighted. Many—perhaps most—natural phenomena can be usefully described in network terms. The brain is a huge network of neurons linked by synapses. The control of genetic activity in a cell is due to a complex network of genes linked by regulatory proteins. Social communities are networks in which the nodes are people (or organizations of people) between whom there are many different types of possible relationships. The Internet and the World-Wide-Web are of course

two very prominent networks in today's society. In the realm of national security, much effort has been put into identifying and analyzing possible "terrorist networks."

All of these topics have been studied for some time, but only fairly recently has the study of networks in general become a major topic of research in complex systems. Among the reasons for this are fast computers, which make it possible to study real networks empirically, and the increased attention this field is getting from physicists looking to other fields to apply their powerful analysis techniques. As Watts points out, "No one descends with such fury and in so great a number as a pack of hungry physicists, adrenalized by the scent of a new problem" [63].

Here are some of the questions that network scientists are trying to address:

- What topological measures can be used to characterize properties of networks?
- What properties do different sets of real-world networks share, and why? How did these properties come about?
- How do we design efficient algorithms to determine these properties?
- How do these properties affect the dynamics of information (or disease, or other communications) spreading on such networks, and the resilience of such networks to noise, component failures, or targeted attacks?
- Given a network with certain properties, what are the best ways to search for particular nodes in this network?

Answering these questions could have a large impact not only on our understanding of many natural and social systems, but also on our ability to engineer and effectively use complex networks, ranging from better Web search and Internet routing to controlling the spread of diseases, the effectiveness of organized crime, and the ecological damage resulting from human actions.

3.2 Network models

The books by Watts and Barabási describe three general models of networks: *random*, *small-world*, and *scale-free*. These models are each characterized by the way in which networks are created, and by several resulting statistics, such as degree distribution, average path length between pairs of nodes, and degree of clustering.

In the models described below, assume there are n nodes that form a given network. Also assume, for simplicity, that links are undirected and unweighted. Two nodes are said to be *neighbors* if there is a link between them.

The *degree* of a node is defined as its number of neighbors. The *degree distribution* of a network is the distribution over the frequencies of different degrees over all nodes in the network. For example, in Figure 1, the node labeled n_1 has degree 3, and the degree distribution of the network is plotted next to the network. The average degree over all eight nodes is 2.

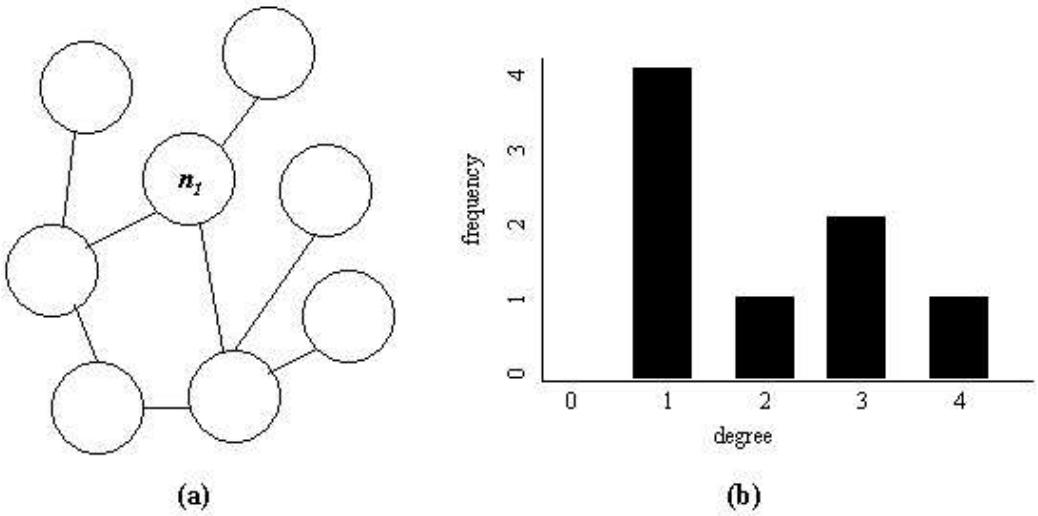

Fig. 1. (a) Example of a network. The node labeled n_1 has degree 3 (i.e., three links to other nodes). (b) Degree distribution of the network.

Random Networks

A *random network* is created by specifying that each pair of nodes is connected by a link with uniform probability p. Such networks were studied from a pure mathematics point of view by Erodös and Rényi [43]. In the limit of large n, many of the ensemble average properties of random networks can be expressed analytically, making them elegant mathematical objects. However, some key properties of Erodös-Rényi random networks turn out to be quite different from those seen in most real-world networks.

First, random networks do not exhibit the strong clustering of nodes (i.e., sub-networks of densely interconnected nodes) seen in the real world. Clustering occurs when, say, you are more likely to be a friend of your friends' friends than of a random person. The *clustering coefficient* of a network is the average probability that two neighbors of a given node are also neighbors of each other. The clustering coefficient of a random network is exactly p. However, empirical studies of real-world networks often show clustering coefficients that are orders of magnitude higher than the value on a random graph with the same number of nodes and links.

Second, it can be shown that the degree distributions of random networks are approximately Gaussian (actually, Poisson, in the limit of large n), which is very different from the power-law degree distributions often observed in real-world networks. A power-law degree distribution looks like $P(k) \sim k^{-\gamma}$, where $P(k)$ is the probability that a node will have degree k, and γ is a real number greater than zero. Figure 2 (left) shows a hypothetical power-law degree distribution with $\gamma = 2.1$ (the approximate empirical degree distribution of the World Wide Web [2]). The power-law distribution is right skewed: it has a long right tail with no cut-off point (i.e., where $P(k)$ reaches zero). For contrast, Figure 2 (right) shows a Gaussian distribution with short tails and clear cutoffs.

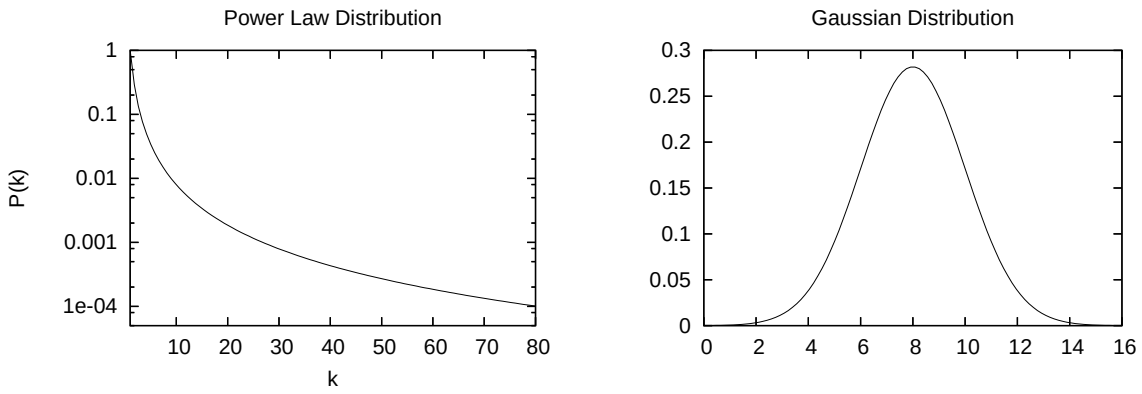

Fig. 2. Left: A hypothetical power law distribution $P(k) = k^{-2.1}$, plotted on a semi-log scale, where k is the degree of a node and $P(k)$ is the probability of a given node having that degree. Right: A hypothetical Gaussian distribution.

Small World Networks

In 1998, Watts and Strogatz introduced their *small-world* network model. To build a small-world network, start with a cellular-automaton-like "ring" lattice in which each node links to k of its nearest neighbors . For each link, with small probability p rewire one end of the link to a randomly chosen node in the network. For $p = 0$ the result is a completely regular network; for $p = 1$ the result is a completely random network. For low but nonzero p, the result is a network with many local connections and a few "long-distance" connections. Such networks have the so-called "small world property": even though relatively few long-distance connections are present, the shortest path length (number of link hops) between two nodes scales logarithmically or slower with network size n for fixed average degree. This means that even in small-world networks with many nodes, the shortest path length between two individual nodes is likely to be relatively small, hence the "small-world" designation. In contrast, in a regularly connected network, path lengths scale linearly with the size of the network. Watts and Strogatz discovered that the small-world property comes about even when p, the rewiring probability, is very small.

Small-world networks with low but non-zero p also exhibit a high clustering coefficient. In contrast, random networks have low average path lengths and low clustering coefficients.

However, the Watts and Strogtaz model typically results in degree distributions that do not match most studied real-world networks.

Scale-Free Networks

About the same time that Watts and Strogatz developed their small-world network model, Barabási and Albert developed an alternative network model, in which growth via "preferential attachment" produced so-called *scale-free* networks. A scale-free network is simply one with a power-law degree distribution, such as the distribution shown in Figure 2 (left). Note that the Gaussian distribution shown in Figure 2 (right) has "cutoff" values where the

distribution goes to zero. This gives the distribution an intrinsic scale—e.g., if adult human heights were plotted versus their frequencies in the population, no human would be, say, under 1 foot or over 15 feet tall. In contrast, a scale-free distribution has no cutoff; instances of all scales are present, hence the term "scale-free."

In a network with a Gaussian degree distribution, the vast majority of nodes have approximately the same number of links. It is very rare to find large deviations from the average. In contrast, in a scale-free network, there are large variations in the number of links per node. There are a small number of nodes ("hubs") with many links coexisting with a much larger number of nodes with a smaller number of links.

Many—perhaps most—real-world networks that have been studied seem to be scale-free; that is, they have power-law rather than Gaussian degree distributions. This is familiar in our own lives: you probably know a few people who each have a huge group of friends (or collaborators, or fans) and many more people with much smaller such groups.

In 1999 Barabási and Albert proposed the mechanism of "growth by preferential attachment" as the explanation for the empirical fact of scale-free networks [3]. Unbeknownst to them at the time, they had rediscovered suggestions made earlier by Simon and Price [44]. The simple idea is that "the rich get richer." Networks grow in such a way that nodes with higher degree receive more new links than nodes with lower degree. Intuitively this makes sense. People with many friends tend to meet more new people and thus make more new friends than people with few friends. Web pages with many incoming links are easier to find than those with few incoming links, so more new web pages link to the high-degree ones.

The network growth model proposed by Barabási and Albert is as follows: The network starts out with a small number of initial nodes. At each time step, one new node is added, and is connected to m existing nodes, for some m. The probability of making one of those connections to node i is proportional to the degree of node i.

In various studies of real-world networks, including the World Wide Web, the Internet, electrical power networks, networks of airline city connections, scientific collaboration networks, metabolic networks in biological organisms, protein networks in yeast, and several others, Barabási, Albert, and others found that the empirical degree distribution was explained well by the power law $P(k) \sim k^{-\gamma}$, with γ somewhere between 2 and 3 (depending on the particular network).

Barabási and Albert (echoing the earlier work by Price and others) conclude that their simple "growth with preferential attachment" mechanism is what drives the evolution of real-world networks. Barabási points out that in the real world, "small differences in ability or even purely random fluctuations can get locked in and lead to very large inequalities over time" [2].

It can also be shown that scale-free networks have the small-world property. Thus, some (but not all) small-world networks are also scale-free.

Table 1 gives a qualitative summary of the various characterizations of networks described

Network Model	**Degree Distribution**	**Clustering Coefficient**	**Average Path Length**
Regular	constant	high	high
Random	Poisson	low	low
Watts-Strogatz small world (low, nonzero p)	depends on p	high	low
Barabási-Albert scale free	power law	high	low
Empirical results on real-world networks	power law	high	low

Table 1
Qualitative comparison among characterizations of different network models and empirical results on real-world networks.

above with respect to the different network models.

Implicit in the discussion above are several simplifying assumptions that have been made in most studies of networks. The models discussed so far do not include new links appearing between existing nodes, or existing links disappearing. Likewise, no variation in "weights" on links or type of link is used—there no representation of different strengths or dimensions of relatedness. There is no cost associated with creating links. Moreover, it is unrealistic to assume that the probability of creating a link to a node is exactly proportional to its degree; in the real world, all nodes are not identical, but vary in intrinsic properties such as "quality" (e.g., perhaps some Web pages are linked to because they are high quality, not just because they already have many links). However, if the proportional probability assumption is modified, the degree distribution might no longer be a power law. There is much on-going research to investigate relaxing these and other assumptions. But the surprising fact remains that, in spite of all the assumptions, these simple models seem to capture some essential aspects of real-world networks.

4 Applications of Network Characterizations

Both the Watts and the Barabási books emphasize the importance of "network thinking" in dealing with complex systems in the real world. The purpose of characterizing networks according to degree distribution, clustering coefficient, average path length, and the like is both to better understand networks from a scientific point of view and to develop better technologies for designing and managing networks in desired ways. Some examples are: developing better methods to design and deliver vaccines and other drugs; being able to better manage

developing epidemics; designing improved algorithms for Web search, search in peer-to-peer networks, Internet communication, and computer security; mitigating the effects of power failures; and facilitating discovery of highly connected sub-networks of individuals, whether they be customers who share tastes in books or possible criminal organizations which should be more closely monitored.

Cohen, Havlin, and ben-Avraham [8] propose an intriguing public health strategy that employs "network thinking" for vaccination strategies in the face of limited resources: Instead of vaccinating people randomly, ask a large number of people to each name a friend, and vaccinate that friend. The idea is that in a social network with power-law degree distribution and high clustering coefficient, collecting names of friends of individuals will be more likely to produce hubs of the network (people with many friends) than choosing random individuals to vaccinate. And vaccinating a hub is more likely to slow the spread of a disease than is vaccinating a randomly chosen person.

Similarly, drug designers might focus on the hubs of genetic regulation networks in a virus as potential targets of new drugs. Web search engines might rank pages by focusing on both "authorities" (pages pointed to by many other pages) and "hubs" (pages connecting to many pages) [30].

More generally, characterizations of networks are needed to predict the conditions under which they will exhibit resilience to the loss or failure of nodes or links and to determine their vulnerabilities. It is know, for example, that in the face of random failures of nodes or links, scale-free networks are far more robust than random networks in the sense that they retain the basic structural and communications characteristics summarized in Table 1. This is why the Internet exhibits robustness even though individual routers fail all the time. However, scale-free networks are not resilient to failures of their hubs, and are thus quite vulnerable to accidental failures or targeted attacks. Bad weather or other problems at a hub airport, for example, can cause long delays in flights all over the country. While hubs are a reason for short path lengths in scale-free networks, this comes at the price of vulnerability. There is clearly a trade-off between efficiency and robustness.

4.1 Searching in networks

A major potential application of network thinking is improved algorithms for searching networks for nodes with certain characteristics. This is in general a difficult problem. In a large network, just because there exists a short path from one node to another does not mean that it is easy to find.

A classic series of sociology experiments performed in 1957 by Stanley Milgram [37] illustrates both the small-world property of social networks as well as the difficulty of finding short paths between people. Milgram's initial experiment involved volunteers living in Kansas and Nebraska. Each was given a letter, and the goal was for the letters to reach a particular target person in Boston who was not known by any of the Midwestern volunteers. Each volunteer

was told only the name and profession of the target, and was instructed to send the letter to someone whom they knew on a first name basis and whom they thought would be more likely than themselves to know the target person. Each recipient in the chain was given the same instructions. Each letter had a record of the path of people through which it passed. In several different experiments with varying protocols, Milgram found that between about 5 and 40 percent of the letters did get through to the targets. He found that the average path length (number of people passing on a letter) of these successful letters was six. This is the origin of the familiar notion of "six degrees of separation." Milgram's analysis also demonstrated the existence of hubs in the network: the successful letters all passed through a small number of particularly well-connected people.

Watts, in his book, describes this experiment in depth and points out several aspects of the original experiment that call Milgram's reported results into question. He also points out that while the average successful path-lengths were short, only a relatively small number of the letters were successful, which demonstrates the difficulty people have in searching for short paths with only local knowledge of the network.

More recently, the computer scientist Jon Kleinberg proved that in networks following the Watts-Strogatz model, in which the random rewirings are chosen with uniform probability over all nodes, there is no efficient decentralized search algorithm that will find short paths [31]. This shows, among other things, that Milgram's social networks are different in some way from those of the Watts-Strogatz model, since at least some short paths were discovered. Kleinberg proposed a variant of the Watts-Strogatz model that allows for efficient decentralized search. The model starts with set of nodes where the spatial distance between two nodes can be defined in d dimensions. The network is constructed by repeatedly adding links between two nodes with probability that decays with the dth power of their distance. This model yields a network in which each node is likely to have the same number of links at all distance scales. In Watts' words: "[Y]ou should expect to have as many friends who live in your neighborhood as live in the rest of the city, the same number of friends who live in the rest of the state, and the same number again in the remainder of the country, and so on, right up to the scale of the whole world."

Kleinberg showed that this condition of equal connectivity at all distance scales both ensures short average path lengths and allows an individual node i to efficiently find a short path to another node j. Node i simply forwards a message to the node linked to it whom it judges closest to the target. As Watts points out, "In order for social connections to be useful—in the sense of finding anything deliberately—they have to encode information about the underlying social structure....But what Kleinberg's model does not explain is how the world could actually be this way."

Of course, in the real world, the "distance" between two individuals can have many dimensions: physical space, shared interests, age, social class, gender, and so on. Watts gives evidence that making these different dimensions explicit makes peer-to-peer network search much easier, even when Kleinberg's condition of equal connectivity over distance scales is not met.

Another important application of network analysis is the problem of finding clusters, or community structures, in a given network. This is the problem of finding sub-networks ("communities') that contain dense interconnections, meaning that nodes in a given community are much more likely to be related to nodes in that same community than to nodes in other parts of the network. Finding such communities is related to the problem of graph partitioning in computer science, and to the problem of hierarchical clustering in sociology.

Influential approaches to this problem have been developed by Newman and Girvan [47], Newman [45], and Guimerà and Amaral [18]. In particular, Guimerà and Amaral's work is of interest because their algorithms discovered previously unknown community structure in metabolic networks. In such a network, the nodes represent metabolites (small molecules that are the substrates and products of metabolic processes in organisms). Two nodes i and j are connected by a link if there is a chemical reaction in which i is a substrate and j a product, or vice versa.

Guimerà and Amaral used an existing database of metabolic networks to search for community structure in the networks of 12 different organisms. They used simulated annealing to search for a community structure that maximizes network modularity, as defined by Newman and Girvan. In the structures that were found, Guimerà and Amaral identified several types of "roles" played by nodes, including

- hubs of modules (nodes with many connections within their community or "module");
- ultra-peripheral nodes (all links are within the node's module);
- peripheral nodes (most links are within the node's module);
- non-hub connectors (nodes with many links to other modules);
- non-hub kinless nodes (links homogeneously distributed among all modules).

They also found three sub-roles for hubs: provincial (the vast majority of the node's links are within the node's module), connector (the node is both a hub in its module and has many links to most other modules), and kinless (the node's links are homogeneously distributed among all modules).

Guimerà and Amaral's hypothesis was that nodes playing these roles in different organism had similar metabolic functions in their respective organism, and that the *structurally* relevant nodes (i.e., metabolites) should be conserved across species.

To investigate this, the authors assessed the degree to which different roles were conserved across different classes of organisms in the data. They found that ultra-peripheral nodes are the least conserved and that connector hubs are the most conserved, which is not surprising. However, they also found a counter-intuitive result: non-hub connectors are more conserved than provincial hubs, even though the degree of a hub was much higher than that of a non-hub connector.

The authors' interpretation is that connector nodes are responsible for communication be-

tween otherwise poorly connected modules. The modules themselves play functional roles, with functions that can be implemented by a variety of metabolites, so a particular set of metabolites implementing a function need not be conserved across species. However, connectors play a structural role, so must be preserved. In their words, "the global role of nodes in the network might be a better indicator of their importance than degree."

5 Computation in Cellular Automata

Up to this point, I have discussed only static structural properties of networks—e.g., degree distribution, hubs, clustering—and their implications for our ability to search and analyze those networks. No mention has yet been made of how these networks actually do something in and of themselves.

The books by Watts and Barabási focus largely on structural properties of networks, rather than on the dynamics of processes on networks, which is much less well understood. Watts emphasizes this point: "Next to the mysteries of dynamics on a network—whether it be epidemics of disease, cascading failures in power systems, or the outbreak of revolutions—the problems of networks that we have encountered up to now are just pebbles on the seashore."

I would extend this notion of dynamics to include information processing performed by networks. The main challenge is understanding the dynamics of the propagation of information (or diseases, computer viruses, fads, innovations, etc.) in networks, and how these networks process such information.

To understand and model information propagation, one has to take into account not only network structure, but also the details of how *individual* nodes and links propagate, receive, and process information, and of how the nodes, links, and overall network structure change over time in response to those activities. This is a much more complex problem than characterizing static network structure.

Complex networks with mainly local connections and no central control, such as the brain, the immune system, genetic regulatory networks, and so on have often been described in information processing terms, but little is known about how such networks can, and do, effect highly nontrivial and adaptive computations. Insight into how such computations work would not only be a tremendous increase in our understanding of these systems, but would also help us to design artificial complex adaptive systems with such collective computational properties. For example, a major challenge for computer science is to design adaptive information processing in distributed, decentralized sensor networks. Likewise, regular networks such as "cellular arrays" have been proposed as architectures for nanoscale or molecular computation because of their low-cost local connectivity and the potential of manufacturing such arrays using DNA self-assembly techniques [67].

Cellular automata (CAs) are the simplest form of regular networks. CAs have played a

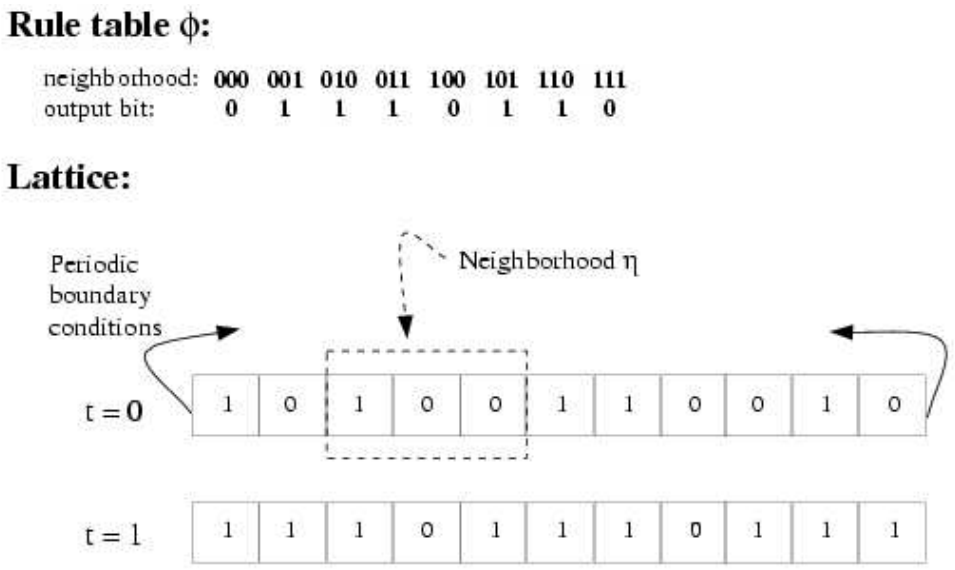

Fig. 3. Illustration of a one-dimensional, binary-state CA, with periodic boundary conditions shown iterating for one time step. At each time step, each cell updates its state as a function of its own state and those of its two nearest neighbors, according to the CA rule. Wolfram [65] proposed a numbering scheme for such *elementary* CAs in which the output bits are ordered lexicographically, as in the figure, and are read right-to-left (neighborhood 111 first) to form a binary integer between 0 and 255. In that scheme, the elementary CA pictured here is number 110.

considerable role as simple models for the study of complex systems, and have recently been brought into the limelight by Steven Wolfram in his well-publicized book *A New Kind of Science* [66].

A CA, in its simplest incarnation, is a one-dimensional lattice of sites or *cells*, each of which is colored either black or white. The color is called the *state* of the cell, and can change over time in discrete steps. At each time step every cell updates its state—it either maintains its color from the previous step or takes on the other color—as a function of its previous state and that of its two nearest neighbors. One can either consider the line of cells to extend without limit in both directions so there is no ambiguity at the boundaries, or one can define boundary conditions—for example, having the line wrap around in a circle.

The list of mappings of each three-cell neighborhood to the update state for the center cell is called the CA *rule* or *rule table*. A CA iterates in discrete time steps. At each time step, it synchronously updates the states of all cells according to its rule. Figure 3 illustrates such a rule and a circular 11-cell lattice in which that rule is applied over one time step. The rule ϕ is given as a table that maps each possible three-cell neighborhood (top row) to an update state for the center cell (bottom row), with white coded as "0" and black coded as "1."

There are 256 possible rules for a CA with two states and three-cell neighborhoods. Wolfram numbered these according to the binary integer given by stringing together the update states,

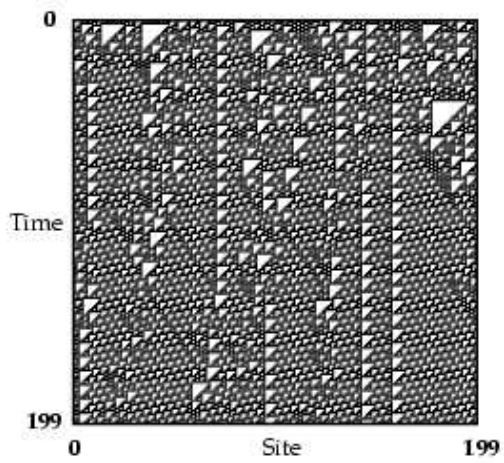

Fig. 4. A space-time diagram illustrating the typical behavior of Elementary CA 110 from a random initial configuration. The one-dimensional lattice of cells is displayed horizontally and time increases down the page. (This and other space-time diagrams in this paper were plotted using the lattice automaton simulation and graphing package la1d, written by James P. Crutchfield.)

with the state in the 111-neighborhood being the most significant bit [65]. The rule in Figure 3 is number 110 (01101110 in binary). The 11-cell lattice starts with a random initial configuration of states at $t = 0$, each of which is updated synchronously at $t = 1$ according to the given rule. Figure 4 gives the spatio-temporal behavior of this CA when run on a lattice of 200 cells for 200 time steps, also starting with a random initial configuration.

Such one-dimensional, two-state, two-neighbor CAs are called *elementary*; more complicated versions can have additional states per cell, larger neighborhoods to determine update states, additional dimensions, and asynchronous state updates. For example, the well-known "game of Life," invented by John Conway [5], is a two-dimensional CA in which each cell has two possible states (black and white), and each cell updates its state as a function of its current state and those of the eight surrounding cells.

Cellular automata were originally invented and studied by the mathematicians Stanislaw Ulam and John von Neumann. Von Neumann adopted CAs as a substrate for constructing a self-reproducing automaton [61].

Cellular automata have been labeled as "non-von-Neumann-style" architectures, to contrast their decentralized, spatially extended, homogeneous parallelism with the familiar "von-Neuman-style" architecture with its modular separation of an arithmetic logic unit, control unit, and random access memory, and its serial execution of instructions. This is an amusing irony, given that von Neumann played a major role in the invention of both types of computing architecture.

Because of their simplicity, cellular automata have been used by several research groups as vehicles to study collective computation in networks. When viewing a cellular automaton as a computer, one can view the CA rule as an encoding of a program and the initial configuration

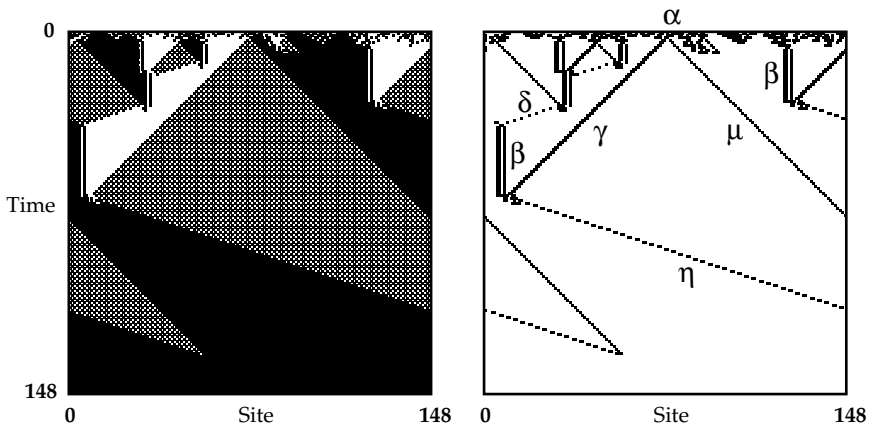

Fig. 5. **Left-hand side:** Space-time diagram displaying behavior of a CA evolved by the genetic algorithm for the density classification task, with state 1 colored black and state 0 colored white. The initial configuration has a majority of 1s. The grey-looking region consists of alternating 1s and 0s. **Right-hand side**: "Particles", which are the boundaries between regular regions.

as an encoding of an input. The computation is the step-by-step evolution of the states over time, with the output being some final pattern of states. An alternative view takes the initial configuration to encode both the program and the input, and the CA rule is the "hardware" that runs the program on the input to produce the output (again, some final pattern of states).

Because CAs are spatially extended, decentralized nonlinear systems, they can produce complicated, hard-to-predict behavior in spite of the simple transition rules by which they are defined. While there have been several demonstrations that some CAs can simulate universal Turing Machines [5,32,61,66], in practice it is challenging to design a CA rule and initial configuration that will result in a desired computation.

A number of groups have used evolutionary computing methods to automatically design CAs to compute. For example, Crutchfield, Mitchell, and Das [11] used a genetic algorithm to design rules for one-dimensional, binary state CAs to perform two tasks requiring collective computation: density classification and synchronization. In the density classification task, a CA, starting from any initial configuration of 1s (black) and 0s (white), must iterate (within a given number of time steps) to a fixed-point configuration of all 1s if there is a majority of 1s initially, and all 0s otherwise. In the synchronization task, a CA, again starting from any initial configuration, must iterate to a configuration that oscillates between all 0s and all 1s. Both tasks would be easy for systems with central controllers, but it is not trivial to design CA rules to perform these tasks.

The genetic algorithm was able, on some runs, to evolve high-performance rules to perform these two tasks. (Even better success has been obtained via co-evolutionary algorithms [27,41].) However, the mechanisms by which the evolved rules performed the computations

was not immediately obvious. It is, in general, not possible to predict from the rule table alone what a CA's temporal behavior will be or how, exactly, it is able to perform a given computation. One can gain some insight by examining plots of the space-time behavior of the CA for different initial configurations. An example from a rule evolved for the density-classification task is given in the left-hand diagram of Figure 5. Here, the one-dimensional lattice starts with a randomly generated initial configuration with a majority of cells in state 1. It can be seen that the CA correctly iterates to a fixed-point of all 1s (the behavior of this CA is likewise correct for most initial configurations, no matter which state is in the majority initially). The transient behavior before the fixed point is the period when the computation is being performed, but how exactly is it accomplished? It is not obvious from the space-time patterns.

Crutchfield and Hanson [9] proposed a framework for understanding computation in such spatially extended, decentralized systems, and Crutchfield, Mitchell, and Das [10,12] applied this framework to the evolved CAs.

In the left-hand diagram of Figure 5 it can be seen that over short times, local high-density regions are mapped to all 1s, local low-density regions are mapped to all 0s, with a vertical boundary in between them. This is what happens when a region of 1s on the left meets a region of 0s on the right. However, when a region of 0s on the left meets a region of 1s on the right, rather than a vertical boundary being formed, a checkerboard region (alternating 1s and 0s, appearing "gray" in the figure) is propagated. When the propagating checkerboard region collides with the black-white boundary, the inner region is cut off and the outer region is allowed to propagate. In this way, the CA uses local interactions and geometry to determine the relative sizes of adjacent low- and high-density regions that are larger than the neighborhood size. For example, in the left-hand diagram, the large inner white region is smaller than the large outer black region; thus the propagating checkerboard pattern reaches the black-white boundary on the white side before it reaches it on the black side; the former is cut off, and the latter is allowed to propagate.

The black-white boundary and the checkerboard region can be thought of as "signals" indicating "ambiguous" regions. The creation and interactions of these signals can be interpreted as the locus of the computation being performed by the CA—they form its emergent "algorithm." Crutchfield, Mitchell, and Das formalized these signals as propagating *particles* (Figure 5, right-hand diagram) which follow specific laws of kinetics and interaction in a given CA. They were able to explain the computation in terms of these particles. moreover, using this "particle" framework, Hordijk, Crutchfield, and Mitchell [25] were able to predict the computational behavior of a given evolved CA, and explain its errors on the density-classification task.

This work has been extended, in part, to two dimensional CAs [26]. Computation in evolved two-dimensional CAs is harder to analyze, since the spatio-temporal patterns formed are much more complex.

Interestingly, a group studying plant physiology and behavior have found that the dynamics

of stomata networks (networks of apertures that can open and close) on leaves has characteristics similar to the dynamics of particles in evolved CAs observed by Crutchfield, Mitchell, and Das [51]. The group proposes that stomata networks are doing a similar form of collective computation that controls the opening and closing of individual stomata so as to optimize the balance between carbon dioxide uptake and water loss for the plant as a whole.

It is still an open question whether evolutionary methods can be used to efficiently design CAs and other network architectures to perform more complicated computational tasks. Likewise, the effects of varying the network's topology on computational ability or evolvability are largely unknown.

Some initial work along these lines has been done, independently, by Watts [62]; Tomassini, Giacobini, and Darabos [33]; and Teuscher [58], demonstrating that small-world topologies are more efficient and robust than CAs for performing the density-classification and synchronization tasks.

Along the same lines, Oikonomou and Cluzel [48] demonstrated that scale-free networks are more "evolvable" than random networks for a task requiring the network to display a desired temporal behavior. Several groups have also explored the computational abilities of random Boolean networks (e.g., [29,59,36])

There is, not surprisingly, a trade-off between computational efficiency and cost of long-range connections. As far as I know, this trade-off has not been precisely quantified in general terms, and no analysis has yet been done of the spatio-temporal mechanisms underlying these various non-CA network computations. I believe this is a very important area of research, and is at this time wide open for advancement.

6 Adaptive Information Processing in Natural Complex Systems[1]

What principles govern information processing in complex networks in nature? In 1999 a group of theoretical immunologists met at the Santa Fe Institute to participate in a workshop, organized by Lee Segel and Irun Cohen, whose aim was to explore interdisciplinary approaches to this question. They also invited some people from computer science, biochemistry, and entomology to explore general principles underlying "distributed autonomous systems," both natural and artificial. The proceedings of this workshop was published in book form under the title "Design Principles for the Immune system and Other Distributed Autonomous Systems" [57]. In this section I describe three instructive examples inspired by this book; in the following section I present four proposed general principles of information processing in decentralized systems, which I believe apply in cognition as well.

[1] Parts of this section were adapted from [40].

6.1 Adaptive information processing in the immune system via affinity maturation and diffuse feedback

Immunology is one field that has taken quite seriously the challenge of understanding the dynamics of information on networks. The immune system is a huge network consisting of many kinds of cells and molecules and their interactions, with very complex dynamics giving rise to continual adaptation and learning, memory, pattern recognition, and collective global control. How does it all manage to work?

An instructive example is the process called "affinity maturation," in which the immune system adaptively increases its specificity for detecting foreign pathogens such as viruses or bacteria.

Pathogens are detected by the immune system via the collective actions of trillions of lymphocytes (white blood cells) that continually circulate in the blood and tissues. A lymphocyte's surface is covered with identical receptors that bind to a particular range of molecular shapes. Within this range, there are particular shapes to which these receptors bind most strongly and others to which they bind more weakly.

In order to "cover" the huge space of possible pathogen shapes in a reasonable way, the population of lymphocytes in the body at a given time is enormously diverse: each individual lymphocyte recognizes a range of shapes that differs from the range recognized by other lymphocytes in the population. When a lymphocyte is born, a novel set of identical receptors is created via a complicated randomization process in the lymphocyte's DNA. Because of continual turnover of the lymphocyte population (about 10^7 new lymphocytes per day) and thus continual introduction of new receptors, the space of possible pathogen shapes is well covered in a dynamic way.

This means that any foreign pathogen is almost certain to encounter a lymphocyte whose receptors will match its shape to at least some degree. When this happens, a bond is formed. The better the match ("affinity"), the stronger the bond. If the number of strongly bound receptors on a lymphocyte exceeds a threshold, and the lymphocyte gets "go-ahead" signals from other lymphocytes with similarly bound receptors, the lymphocyte is activated. Once activated, the lymphocyte secretes antibody molecules, which bind to and neutralize the pathogens, and mark them for destruction by other immune system cells.

The activated lymphocyte then migrates to a lymph node, where it divides rapidly, producing large numbers of daughter lymphocytes. These daughters are subject to a very high rate of mutation, which results in some modifications in their receptors. These mutants are tested on pathogen molecules that are captured in the lymph node. The mutants that do not bind die after a short time.

The ones that bind to the pathogen are unleashed into the bloodstream, where some of them encounter and bind to pathogen molecules, in some cases more strongly then did their mother lymphocyte. These activated daughters themselves migrate to lymph nodes, and

create mutated offspring. This cycle continues, with the best-matching new lymphocytes themselves producing the most offspring cells. A Darwinian process ensues, evolving large numbers of new lymphocytes with increasingly better affinities for the pathogen.

This process of detection and destruction typically takes a few days to weeks to eradicate the pathogen from the body. Other mechanisms prevent the immune system (in most cases) from mistakenly attacking the body's own molecules. For more details on affinity maturation and other immune system mechanisms, see [20].

Because the immune system itself, in the process of attacking pathogens, causes inflammation and other harm to tissue, Segel and Orosz have argued that the immune system must balance an impetus to kill foreign pathogens with the obligation to prevent harm to the body as much as possible [55,49]. Segel has speculated that this balance requires an additional level of "self-watching." He hypothesizes that specialized signals, in the form of molecules, are created by three types of situations: harm to the body due to pathogens, the immune system's successful killing of pathogens, and self-harm to the body due to immune system actions. These signals communicate such harm or success to other immune system cells. The "pathogen-harm" and "successful-kill" molecules should up-regulate immune system activity in the region in which they are produced, while the "self-harm" molecules should down-regulate such activity. Up-regulation might mean, for example, speeding up the cell-division process in activated lymphocytes; down-regulation might mean slowing it down.

How is such up- or down-regulation achieved? Segel's hypothesized signaling molecules would affect the concentration levels of *cytokines*—chemicals that mediate all immune system functions, and can have the desired up-or-down-regulation effects. Cytokines are known to form a signaling network: all immune cells have receptors for cytokines and are affected by them. Cytokines are currently the object of intense study in immunology: "More than 10,000 papers a year are currently being written [concerning cytokines], yet 'practically nothing is know about the behaviour of the network as a whole'" [54]. Segel proposes that it is the spatial-temporal concentration of cytokines that provides a "diffuse feedback" to the immune system and cause it to up- or down-regulate by affecting the actions of individual lymphocytes at different spatial locations. Diffuse feedback is a statistical effect of spatially heterogeneous concentrations of cytokines that are sampled locally by lymphocytes.

6.2 Foraging and Task Allocation in Ant Colonies

Analogies have often been made between ant colonies and the brain [22, "...Ant Fugue"]; both can be considered as networks of relatively simple elements (neurons, ants) from which emerge larger-scale intelligent and adaptive behaviors. Two examples of such behavior in ant colonies are the ability to optimally and adaptively allocate resources (ants) in foraging for food, and the ability to adaptively allocate ants to different tasks as needed by the colony. Both types of behavior are accomplished with no central control, via mechanisms similar to those described above for the immune system.

In many ant species, foraging works roughly as follows [6,7]. Foraging ants in a colony set out moving randomly in different directions. When an ant encounters a food source, it returns to the nest, leaving a pheromone trail. When other ants encounter a pheromone trail, they are likely to follow it. The greater the concentration of pheromone, the more likely an ant will be to follow the trail. If an ant encounters the food source, it returns to the nest, reinforcing the trail. In the absence of reinforcement, a pheromone trail will dissipate. In this way, ants collectively build up and communicate information about the locations and quality of different food sources, and this information adapts to changes in these environmental conditions. At any given time, the existing trails and their strengths form a good model of the food environment discovered collectively by the foragers.

Task allocation is another way in which an ant colony regulates its own behavior in a decentralized way. Gordon [16,17] has studied task allocation in colonies of Red Harvester ants. Workers in these colonies divide themselves among four types of tasks: foraging, nest maintenance, patrolling, and midden (refuse sorting) work. The numbers of workers pursuing each type of task adapts to changes in the environment. Gordon found, for example, that if the nest is disturbed in some small way, the number of nest maintenance workers will increase. Likewise, if the food supply in the neighborhood is large and high quality, the number of foragers will increase. How does an individual ant decide which task to adopt in response to nest-wide environmental conditions, even though no ant directs the decision of any other ant and each ant only interacts with a small number of other ants?

The answer seems to be that ants decide to switch tasks both as a function of what they encounter in the environment and as a function of their rate of interaction with ants performing different tasks. For example, an inactive ant—one not currently performing a task—who encounters a foreign object near the nest has increased probability of doing nest-maintenance work. In addition, an inactive ant that encounters many nest-maintenance workers entering and leaving the nest will also have an increased probability of adopting the nest-maintenance task; the increased activity in some way signals that there are important nest maintenance tasks to be done. Similarly, a nest-maintenance worker who encounters many foragers returning to the nest carrying seeds will have an increased probability of switching to foraging; the increased seed delivery signals in some way that a quality food source has been found and needs to be exploited. Ants are apparently able to sense, through direct contact of their antennae with other ants, what task the other ants have been engaged in, by perceiving specific chemical residues associated with each task.

Gordon points out that the proposed mechanism of task switching based on interaction rate can explain an initially puzzling finding of her experiments. She found that perturbing a nest by placing toothpicks around it increased the number of nest maintenance workers, who proceeded to assist in removing the toothpicks. However, this increase was more reliably seen in older colonies (ones that have produced more generations of ants) than in younger colonies. The individual ants in each type of colony were the same age, and ants presumably do not have the ability to pass on complex information about tasks and task switching to their offspring. But it does turn out that older colonies are larger than younger colonies,

and in larger colonies ants relying on individual interactions for information would be able to obtain better statistics about the proportion of ants currently doing each task. Thus one would expect a stronger effect in larger colonies, which is what Gordon observed.

6.3 Metabolism in Biological Cells

A third example of information processing in complex networks is cellular metabolism. In a biological cell, nutrient molecules are processed to yield energy, and cellular components are built up via parallel "metabolic pathways." These components are needed for internal maintenance and repair, and for external functions and inter-cellular communication. The reactions that occur continually change the dynamics, causing different building processes to proceed at different rates.

At the microscopic level, the metabolic pathways function roughly as follows. At any given time, millions of molecules in the cell diffuse randomly in the cytoplasm via thermal energy. The molecules continually encounter one another at random. Occasionally (on a scale of microseconds), enzymes encounter molecules of matching shape, allowing reactions to occur. These reactions cause macromolecules to be built up gradually.

Metabolic pathways are complex sequences of these chemical reactions, controlled by self-regulating feedback. At any given time, hundreds of such pathways are being followed, some independent, some interdependent. The pathways result in new macromolecules, initiation of other metabolic pathways, and the regulation of themselves or other metabolic pathways.

The regulation mechanisms typically involve concentration-based activation or inhibition of particular pathways. For example, the infusion of lactose molecules in a cell increases the activation of usually low-activity enzymes that degrade lactose. In many cases, a excess of a product of metabolic path inhibits first step in that path, yielding negative feedback control. In other cases, regulation comes about via changing protein concentrations, which have effects on the genetic transcription of needed enzymes. In addition, the mechanism of "anchoring" enzymes in membranes can serve to make needed interactions more efficient.

As in the immune system and ant colonies, the fact that metabolic pathways proceed in small steps allows the system to continually control the speed and depth of ongoing pathways in response to activities throughout the system.

7 Four Principles of Adaptive Information Processing in Decentralized Systems

The chapters in [57] argue in various ways that the immune system, ant colonies, and cellular metabolism are best understood as distributed autonomous networks that process information in adaptive ways. Information processing in all these systems emerges from multiple feedback mechanisms, and allows the system to generate and use the right resources at the right place and right time in order to benefit the entire system. All this is done without

central control, in a continually changing environment and in the face of multiple, often conflicting requirements.

What common principles can be gleaned from these systems?

The following is a list of four principles I have abstracted from the above sketches of immune systems, ant colonies, and cellular metabolism. While the principles are given here in a serial list, in reality they overlap and interact. I believe that these principles carry over to intelligence, self-awareness, and self-control in other decentralized systems. Broadening their reach—and their number—is the subject of future work.

(1) **Global information is encoded as statistics and dynamics of patterns over the system's components.**

In the immune system, for example, the spatial distribution and temporal dynamics of lymphocytes can be interpreted as a representation of the continually changing population of pathogens in the body. Similarly, according to Segel's speculations, the spatial distribution and dynamics of cytokine concentrations encode global information about the immune system's success in killing pathogens and avoiding harm to the self.

In ant colonies, the colony's food environment is represented, in a dynamic way, by the statistical distribution of ants on various trails, and the colony's overall state is represented by the dynamic distribution of ants performing different tasks.

In cellular metabolism the current state and needs of the cell are continually reflected in the spatial concentrations and dynamics of different kinds of molecules.

In all these cases, information is "read" via sampling statistics obtained by system's components.

But *who* actually interprets this representation, and performs the corresponding adaptations of the system? It is the collection of cells, ants, or active molecules themselves, taking actions depending on statistics they sample locally, via other individuals they encounter, or via the local chemical environment. In order for statistics to be reliably interpreted, actions are taken based on thresholds: a lymphocyte undergoes affinity maturation only if its receptors have a sufficient number of strong bonds, plus additional verification signals from other cells. A nest-maintenance ant switches to foraging only if it has encountered a sufficient number of other successful foragers. A metabolic pathway is followed only if the concentration of a particular metabolite is above a threshold.

(2) **Randomness and probabilities are essential.**

All three of the systems described above use randomness and probabilities in essential ways. The receptor shape of each individual lymphocyte has a randomly generated com-

ponent, so as to allow coverage by the population of many possible pathogenic shapes. The spatial distribution of lymphocytes has a random component, due to the distribution of lymphocytes by the blood stream, so as to allow coverage of many possible spatial distributions of pathogens. The detailed thresholds for activation of lymphocytes, their actual division rates, and the mutations produced in the offspring all involve random or noisy aspects.

Similarly, the movement of ant foragers has random components, and these foragers are attracted to pheromone trails in a probabilistic way. Ants also task-switch in a probabilistic manner.

Cellular metabolism relies on random diffusion of molecules, and on probabilistic encounters between molecules, with probabilities changing as relative concentrations change in response to activity in the system.

It appears that such intrinsic random and probabilistic elements are needed in order for a comparatively small population of simple components (ants, cells, molecules) to explore an enormously larger space of possibilities, particularly when the information to be gained is statistical in nature and there is little a priori knowledge about what will be encountered.

However, randomness must be balanced with determinism: Self-regulation in complex adaptive systems continually adjusts probabilities of where to move and how deeply to explore particular pathways in these large spaces.

(3) **The system carries out a fine-grained, parallel search of possibilities.**

Many, if not all, complex systems in biology have a fine-grained architecture, in that they consist of large numbers of relatively simple elements that work together in a highly parallel fashion.

Several possible advantages are conferred by this type of architecture, including robustness, efficiency, and evolvability. On additional major advantage is that a fine-grained parallel system is able to carry out what Hofstadter [21, Chapter 2] has called a "parallel terraced scan." This refers to a simultaneous exploration of many possibilities or pathways, in which the resources given to each exploration at a given time depends on the perceived success of that exploration at that time. The search is parallel in that many different possibilities are explored simultaneously, but is "terraced" in that not all possibilities are explored at the same speeds or to the same depth. Exploration is "on-line" in the sense that information is used as it is gained to continually reassess what is important to explore.

The fine-grained nature of the system not only allows many different paths to be explored, but it also allows the system to continually change its exploration paths, since only relatively simple micro-actions are taken at any time. Employing more coarse-grained actions would involve committing time to a particular exploration that might turn out not to be warranted. In this way, the fine-grained nature of exploration allows the system to fluidly and continuously adapt its exploration as a result the information it obtains.

For example, at any given time, the immune system must determine which regions of the huge space of possible pathogen shapes should be explored by lymphocytes. Each of the trillions of lymphocytes in the body at any given time can be seen as a particular mini-exploration of a range of shapes. The shape ranges that are most successful (i.e., bind strongly to pathogens) are given more exploration resources, in the form of mutated offspring lymphocytes, than those shape ranges that do not pan out (i.e., lymphocytes that do not bond strongly). However, while exploiting the information that has been obtained, via the process of affinity maturation, the immune system continues at all times to generate new lymphocytes that explore completely novel shape ranges. Thus the system is able to focus on the most promising possibilities seen so far, exploring these most quickly and to the most depth, while never neglecting to explore new possibilities. This is closely related to Holland's general principles for optimally balancing "exploitation" and "exploration" in adaptive systems [24]. Orosz's principle of "dynamic engagement" for the immune system also requires a fine-grained architecture [49].

Similarly, ant foraging uses a parallel-terraced-scan strategy: many ants initially explore random directions for food. If food is discovered in any of these directions, more of the system's resources (ants) are allocated, via the feedback mechanisms described above, to explore those directions further. At all times, paths are dynamically allocated exploration resources in proportion to their promise (the amount and quality of the food that has been discovered at those locations). However, due to the large number of ants and their intrinsic random elements, unpromising paths continue to be explored as well, though with many fewer resources.

Note that the redundancy inherent in fine-grained systems allows the system to work well even when the individual components are not perfectly reliable and the information available is only statistical in nature. Redundancy allows many independent samples of information to be made, and allows fine-grained actions to be consequential only when taken by large numbers of components.

(4) **The system exhibits a continual interplay of bottom-up and top-down processes.**

The notions of "bottom-up" and "top-down" processes are often used to describe aspects of cognition: for example, perception is an interplay between what is perceived "bottom-up" by the retina (or other early perceptual systems) and the "top-down" expectations derived from concepts activated by these bottom-up sensations.

A similar interplay can be found in the immune system. Bottom-up explorations are carried out by a continual patrol of lymphocytes with different receptors, collectively prepared to approximately match any pathogen (what Segel has called "a broad spectrum initial response" [56]). Top-down explorations consist of focused lymphocytes, which, when activated by a match, create offspring that are variations of themselves in order to zero in on a particular pathogen shape.

Likewise ant foraging consists of bottom-up explorations by ants moving at random, looking for food in any direction, and top-down explorations in which ants follow existing pheromone trails.

In cellular metabolism, bottom-up processes of random exploration by molecules are combined with top-down activation or inhibition driven by chemical concentrations and genetic regulation.

As in all adaptive systems, maintaining a correct balance between these two modes of processing is essential. Indeed, the optimal balance shifts over time. Early explorations, based on little or no information, are largely random, unfocused, and bottom-up. As information is obtained and acted on, exploration gradually becomes more deterministic, focused, and top-down, in response to what has been perceived by the system. Exhibiting such an appropriate shift in exploration mode has been hypothesized to be a general property of adaptive and intelligent systems [23,24].

8 Conclusion: Relevance for Artificial Intelligence

The general science of networks and its various applications has significant relevance for AI practitioners. For example, understanding the structure of the Internet and the World Wide Web will be important for intelligent routing, load balancing, search of all kinds, and the deployment of intelligent agents who assist human users in their network-related tasks. Network thinking of the kind described here will also be fundamental for developing effective decentralized algorithms for our increasingly distributed and linked computing, communications, and sensing networks, as well as effective security methods for these increasingly vulnerable systems. All of these are areas in which AI and machine learning research have played and will play major roles.

The understanding of information processing in biological and social networks has inspired novel AI techniques. In addition to the long-standing areas of neural- and evolutionary-

inspired computation, several new nature-inspired techniques have been developed over the past decade or so: ant-colony optimization [13], or more generally swarm intelligence [7]; computer and network immune systems [14,19]; economics-inspired AI [4,64]; and methods inspired by cellular and developmental biology [42,60]; among others. All of these approaches employ one or more of the principles I sketched in the previous section.

In my own work and that of my colleagues, these principles have guided the design of the Copycat, Metacat, Tabletop, and Letter Spirit projects on modeling high level perception, analogy making, and "self-watching", all performed under the guidance of Douglas Hofstadter [15,21,34,35,38,39,52]. I believe that similar principles, gleaned from mechanisms that produce adaptive behavior and self-awareness in natural systems, will be necessary components of future artificial intelligence systems with sophisticated cognitive abilities.

Much has been written about how information in the brain, including information about its own state, is represented in a decentralized, statistical, and dynamical fashion, allowing for pattern recognition, sophisticated memory, and learning. One of the most important research challenges for AI is to elucidate common principles of such representation and information processing more broadly across natural adaptive systems and to use these principles in the development of machines with similar abilities.

Acknowledgments

I am grateful to the following people for discussions on some or all of the work and ideas presented here: Doyne Farmer, Stephanie Forrest, Douglas Hofstadter, John Holland, Mark Newman, and Lee Segel.

This work has been supported by a 21st Century Science Initiative grant in Complex Systems from the J. S. McDonnell Foundation.

References

[1] P. Ball. Material witness: Designing with complexity. *Nature Materials*, 3:78, 2004.

[2] A.-L. Barabási. *Linked: The New Science of Networks.* Perseus, New York, 2002.

[3] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 2002.

[4] E. Baum. Toward a model of intelligence as an economy of agents. *Machine Learning*, 35:155–185, 1999.

[5] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *What is Life?*, chapter 25. Academic Press, London, 1982.

[6] E. Bonabeau. Control mechanisms for distributed autonomous systems: Insights from social insects. In L. A. Segel and I. R. Cohen, editors, *Design Principles for the Immune System and Other Distributed Autonomous Systems*, pages 281–292. Oxford University Press, New York, 2001.

[7] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence*. Oxford University Press, New York, 1999.

[8] R. Cohen, S. Havlin, and D. ben Avraham. Efficient immunization strategies for computer networks and populations. *Physical Review Letters*, 91:247901, 2003.

[9] J. P. Crutchfield and J. E. Hanson. Turbulent pattern bases for cellular automata. *Physica D*, 69:279–301, 1993.

[10] J. P. Crutchfield and M. Mitchell. The evolution of emergent computation. *Proceedings of the National Academy of Science U.S.A.*, 92:10742–10746, 1995.

[11] J. P. Crutchfield, M. Mitchell, and R Das. Evolutionary design of collective computation in cellular automata. In J. P. Crutchfield and P. K. Schuster, editors, *Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function*, pages 361–411, New York, 2003. Oxford University Press.

[12] R. Das, M. Mitchell, and J. P. Crutchfield. A genetic algorithm discovers particle-based computation in cellular automata. In Y. Davidor, H.-P. Schwefel, and R. Männer, editors, *Parallel Problem Solving from Nature—PPSN III*, volume 866, pages 344–353, Berlin, 1994. Springer-Verlag (Lecture Notes in Computer Science).

[13] M. Dorigo and T. Stützle. *Ant Colony Optimization*. MIT Press, Cambridge, MA, 2004.

[14] S. Forrest, A.S. Perelson, L. Allen, and R. Cherukuri. Self-nonself discrimination in a computer. In *Proceedings of the 1994 IEEE Symposium on Research in Security and Privacy*, pages 202–212, Los Alamitos, CA, 1994. IEEE Computer Society Press.

[15] R. M. French. *The Subtlety of Sameness*. MIT Press, Cambridge, MA, 1995.

[16] D. M. Gordon. Task allocation in ant colonies. In L. A. Segel and I. R. Cohen, editors, *Design Principles for the Immune System and Other Distributed Autonomous Systems*, pages 293–301. Oxford University Press, New York, 2001.

[17] D. M. Gordon. The regulation of foraging activity in Red Harvester ant colonies. *The American Naturalist*, 159(5):509–518, 2002.

[18] R. Guimerà and L. A. Nunes Amaral. Functional cartography of complex metabolic networks. *Nature*, 433:895–900, 2006.

[19] S. Hofmeyr and S. Forrest. Architecture for an artificial immune system. *Evolutionary Computation*, 8(4):443–473, 2000.

[20] S. A. Hofmeyr. An interpretive introduction to the immune system. In L. A. Segel and I. R. Cohen, editors, *Design Principles for the Immune System and Other Distributed Autonomous Systems*, pages 3–26, New York, 2001. Oxford University Press.

[21] D. Hofstadter. *Fluid Concepts and Creative Analogies*. Basic Books, New York, 1995.

[22] D. R. Hofstadter. *Gödel, Escher, Bach: an Eternal Golden Braid*. Basic Books, New York, 1979.

[23] D. R. Hofstadter and M. Mitchell. The Copycat project: A model of mental fluidity and analogy-making. In K. Holyoak and J. Barnden, editors, *Advances in Connectionist and Neural Computation Theory, Volume 2: Analogical Connections*, pages 31–112. Ablex, Norwood, NJ, 1994.

[24] J. H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, MA, 1992. First edition, 1975.

[25] W. Hordijk, J. P. Crutchfield, and M. Mitchell. Mechanisms of emergent computation in cellular automata. In A. E. Eiben, editor, *Proceedings of the Fifth International Conference on Parallel Problem Solving From Nature—PPSN V*, New York, 1998. Springer.

[26] F. Jimenez-Morales, J. P. Crutchfield, and M. Mitchell. Evolving two-dimensional cellular automata to perform density classification: A report on work in progress. *Parallel Computing*, 27(4):571–585, 2001.

[27] H. Juillé and J. B. Pollack. Coevolving the 'ideal' trainer: Application to the discovery of cellular automata rules. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, San Francisco, CA, 1998. Morgan Kaufmann.

[28] S. Kauffman. *At Home in the Universe*. Oxford University Press, New York, 1995.

[29] S. A. Kauffman and R. G. Smith. Adaptive automata based on Darwinian selection. *Physica D*, 22:68–82, 1986.

[30] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the Association for Computing Machinery*, 46:604–632, 1999.

[31] J. M. Kleinberg. Navigation in a small world. *Nature*, 406:845, 2000.

[32] K. Lindgren and M. G. Nordahl. Universal computation in a simple one-dimensional cellular automaton. *Complex Systems*, 4:299–318, 1990.

[33] Tomassini M., Giacobini M., and Darabos C. Evolution and dynamics of small-world cellular automata. *Complex Systems*, 15(4):261–284, 2005.

[34] J. Marshall. Metacat: A self-watching cognitive architecture for analogy-making. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, pages 631–636, Mahwah, NJ, 2002. Lawrence Erlbaum Associates.

[35] G. McGraw and D. R. Hofstadter. Perception and creation of alphabetic style. In *Papers of the AAAI 1993 Spring Symposium on Intelligence and Creativity*, Menlo Park, CA, 1993. AAAI Press.

[36] B. Mesot and C. Teuscher. Deducing local rules for solving global tasks with random Boolean networks. *Physica D*, 211:88–106, 2005.

[37] S. Milgram. The small world problem. *Psychology Today*, 2:60–67, 1967.

[38] M. Mitchell. *Analogy-Making as Perception: A Computer Model*. MIT Press, Cambridge, MA, 1993.

[39] M. Mitchell. Analogy-making as a complex adaptive system. In L. A. Segel and I. R. Cohen, editors, *Design Principles for the Immune System and Other Distributed Autonomous Systems*, pages 335–359. Oxford University Press, New York, 2001.

[40] M. Mitchell. Self-awareness and control in decentralized systems. In *Working Papers of the AAAI 2005 Spring Symposium on Metacognition in Computation*, Menlo Park, CA, 2005. AAAI Press.

[41] M. Mitchell, M. D. Thomure, and N. L Williams. The role of space in the success of coevolutionary learning. In L. M. Rocha, editor, *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 118–124, Cambridge, MA, 2006. MIT Press.

[42] R. Nagpal. Programmable self-assembly using biologically-inspired multiagent control. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 418–425, New York, 2002. ACM Press.

[43] M. E. J. Newman. Random graphs as models of networks. In S. Bornholdt and H. G. Schuster, editors, *Handbook of Graphs and Networks*, pages 147–169, Berlin, 2003. Wiley-VCH.

[44] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

[45] M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences, USA*, 103(23):8577–8582, 2006.

[46] M. E. J. Newman, A.-L. Barabási, and D. J. Watts. *The Structure and Dynamics of Networks*. Princeton University Press, Princeton, NJ, 2006.

[47] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Physical Review E*, 69, 2004. Article number 026113.

[48] P. Oikonomou and P. Cluzel. Effects of topology on network evolution. *Nature Physics*, 2(8):532–536, 2006.

[49] C. G. Orosz. An introduction to immuno-ecology and immuno-informatics. In L. A. Segel and I. R. Cohen, editors, *Design Principles for the Immune System and Other Distributed Autonomous Systems*, pages 125–149. Oxford University Press, New York, 2001.

[50] H. Pagels. *The Dreams of Reason*. Simon and Schuster, New York, 1988.

[51] D. Peak, J. D. West, S. M. Messinger, and K. A. Mott. Evidence for complex, collective dynamics and emergent, distributed computation in plants. *Proceedings of the National Academy of Sciences, USA*, 101:918–922, 2004.

[52] J. H. Rehling. *Letter Spirit (Part Two): Modeling Creativity in a Visual Domain*. PhD thesis, Indiana University, Bloomington, IN, 2001.

[53] *Scientific American*, June, 1995.

[54] L. A. Segel. Controlling the immune system: Diffuse feedback via a diffuse informational network. In G. Bock and J. Goode, editors, *Complexity in Biological Information Processing*, pages 31–44. Chichester; New York: John Wiley and Sons, 2001. Novartis Foundation Symposium.

[55] L. A. Segel. Diffuse feedback from a diffuse informational network in the immune system and other distributed autonomous systems. In L. A. Segel and I. R. Cohen, editors, *Design Principles for the Immune System and Other Distributed Autonomous Systems*, pages 203–226. Oxford University Press, New York, 2001.

[56] L. A. Segel. How can perception of context improve the immune response? In L. Steinman, editor, *Autoimmunity and Emerging Diseases*. Center for the Study of Emerging Diseases, Jerusalem, 2001.

[57] L. A. Segel and I. R. Cohen, editors. *Design Principles for the Immune System and Other Distributed Autonomous Systems*. Oxford University Press, New York, 2001.

[58] C. Teuscher. On irregular interconnect fabrics for self-assembled nanoscale electronics. In *Proceedings of the 2nd IEEE International Workshop on Defect and Fault Tolerant Nanoscale Architectures (NANOARCH 2006)*, Boston, MA, 2006.

[59] C. Teuscher and M. Capcarrere. On fireflies, cellular systems, and evolware. In A. M. Tyrrell et al., editor, *Lecture Notes in Computer Science: Proceedings of the Fifth International Conference on Evolvable Systems: From Biology to Hardware, ICES 2003*, volume 2606, pages 1–12. Springer-Verlag, 2003.

[60] C. Teuscher, D. Mange, A. Stauffer, and G. Tempesti. Bio-inspired computing tissues: Towards machines that evolve, grow, and learn. *Biosystems*, 68(2-3):235–244, 2003.

[61] J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, IL, 1966. Edited and completed by A. W. Burks.

[62] D. Watts. *Small Worlds: The dynamics of networks between order and randomness*. Princeton University Press, 1999.

[63] D. Watts. *Six Degrees: The Science of a Connected Age*. Gardner's Books, New York, 2003.

[64] M. P. Wellman. The economic approach to artificial intelligence. *ACM Computing Surveys*, 27(4):360–362, 1995.

[65] S. Wolfram. Universality and complexity in cellular automata. *Physica D*, 10:1–35, 1984.

[66] S. Wolfram. *A New Kind of Science*. Wolfram Media, Inc., Champaign, IL, 2002.

[67] S. Xiao, F. Liu, A. E. Rosen, J. F. Hainfeld, N. C. Seeman, K. Musier-Forsyth, and R. A. Kiehl. Assembly of nanoparticle arrays by DNA scaffolding. *Journal of Nanoparticle Research*, 4(4):313–317, 2002.