

# 1

## Mathematical Modeling of Stem Cells: A Complexity Primer for the Stem-Cell Biologist

**Mark d’Inverno**

*Centre for Agent Technology, Cavendish School of Computer Science, University of Westminster, London, U.K.*

**Neil Theise**

*Division of Digestive Diseases, Departments of Medicine and Pathology, The Milton and Carroll Petrie Division of Beth Israel Medical Center, New York, New York, U.S.A.*

**Jane Prophet**

*CARTE, University of Westminster, London, U.K.*

### INTRODUCTION

Many of us have been fascinated by the straight line of ants stretching from food sources to the anthills we found in our gardens. Taking a closer look, we found that this straight line was made from hundreds of individual, industrious ants, each behaving energetically. If we next focus on an individual ant’s behavior (from a modeling perspective, looking at the individual elements of a system is often called the *micro view*), it is very easy to interpret the behavior of the individual as unfocused and chaotic, and it is certainly very difficult to interpret its behavior as being *purposeful* when taken in isolation. It is only when we take a step back and look at the behavior of the entire group (called the *macro view*) that we can observe a purposeful global system of behavior. This purpose, bringing food back to the anthill, *emerges* from the total of the individual behaviors and interactions of the apparently undirected individual. Somehow, the sum of the local interactions of each individual—responding only to their local environment—produces a stable, surviving system, even though individual ants get lost or die.

This system is a commonly cited example of a large class of systems known as *complex adaptive systems*. In this paper, we provide details of several examples of such systems that illustrate some fundamental qualities and issues, before describing the properties of such systems in general. We will argue that this is precisely the paradigm to understand the global behavior of stem cells. We also discuss some of the details of our current project to model and simulate stem-cell behaviors using this paradigm.

## Simulating Ants

Using techniques from computer science, we can build an artificial model of ants to simulate the essential behavior of an ant colony by defining the behavior of each ant using the same simple computational algorithm. Some readers might be unclear on what an algorithm is, but conceptually it is no more than an automatic process, a clearly defined sequence of explicit instructions. This process defines how the ant should behave, precisely and completely. Let us consider one possible way to codify an individual ant's behavior as detailed in the rules that follow. (This example is adapted from many examples of the ant algorithm. The interested reader is encouraged to look at [education.mit.edu/starlogo/](http://education.mit.edu/starlogo/) and [cognitn.psych.indiana.edu/rgoldsto/complex/](http://cognitn.psych.indiana.edu/rgoldsto/complex/) for some excellent examples of complex systems such as ant colonies.) There are three ordered rules that determine behavior:

- Rule 1. Wander around randomly.
- Rule 2. If you do happen to find some food, take it back to the anthill and leave a trail of pheromones that will evaporate over time. Once you have done this go back to Rule 1.
- Rule 3. If you find a pheromone trail, then follow it in the direction that takes you away from the anthill until either:
  - (i) You find food and can perform Rule 2, or
  - (ii) The trail disappears, in which case, go back to Rule 1.

It is a simple matter to set up this simulation. First, we build an artificial model of an environment that is typically a two-dimensional space. This environment will contain the ant colony and several sources of food scattered around within it. We then program each of our ants with exactly the same algorithm described earlier. Lastly, we provide some initial conditions for this simulation as follows:

1. Place the anthill at some location.
2. Place the food sources at other locations.
3. Create a reasonably large (typically greater than 100) number of ants.
4. Position each individual ant at some point close to the anthill.

The program can then be *run*. This means that each ant starts processing the rules; it searches randomly and changes its behavior if it discovers entities in the environment, such as food or pheromones. What is extraordinary about the first time most observers see such an experiment is just how *intelligent* and *sophisticated* the simulated ant colony behaves as a whole. They soon find the food sources and bring them back, bit by bit, to the anthill. Some sources are found and depleted more quickly than others but, in most cases, all the food is discovered and brought back to the anthill.

This example is often used as an exemplar of a *complex system* (1). A complex system is often described as one in which the overall behavior of a system is somehow qualitatively different from that of the individual parts; the parts themselves behave as though they are not goal-directed, whereas the global behavior of the system is *meaningful*. In our current example, it is a simple matter to perceive the goal of the colony as searching for, and returning home with, food. This meaning only arises or *emerges* when looking at the whole system and comes from the collective behavior and interactions of a set of simple elements called *agents* (2). Because each individual agent's behavior or interaction may have significant ramifications for the entire future of the system, the global behavior of any such system cannot be predicted. Currently, there is no mathematical system that can model complex systems in a sufficiently detailed way for such predictions to be made. Only by *running* the simulation can we see what will happen in the future,

given a specific initial state. It is practically impossible to predict what global behavior will arise in general. Before we discuss agents, complex systems, and emergence in more detail, we will first give descriptions of two other examples of systems where sophisticated global behavior arises through the very simple behavior of a large number of interacting computational entities.

### Simulating Populations: The Game of Life

One of the first and, at the time, most extraordinary examples of an algorithm-based system that is often described as an example of a complex system is the so-called *Game of Life* by John Conway (3), which he first proposed in the late 1960s. The aim of constructing this artificial system was to show how the simplest set of rules for individual agents could be used to generate sophisticated global behavior that modeled some very basic principles of birth, death, and survival in a homogenous population. By homogenous, we mean that all agents (as with our simulated ants) had the same set of behavioral rules. In this simulation, agents would survive or die depending on the conditions of their *local environment*. If they were isolated or overcrowded, the agent would die. If neither was the case, then the agent would survive and, furthermore, in some appropriate conditions where there was neither overcrowding nor isolation, agents would be born into empty locations.

The rules of this system were chosen such that the behavior of each population was unpredictable and, after some revisions, the following set of criteria were developed:

1. There should be no initial configuration for which there is a simple proof (law) that the population can grow without limit.
2. Even without any such laws some initial configurations can grow without limit.
3. Depending on the initial conditions, the population would either:
  - (i) change and grow but ultimately die,
  - (ii) settle into a stable equilibrium that would never change, or
  - (iii) enter an oscillating phase in which they repeat a cycle of two or more periods endlessly.

Once again, there are simple agents (called counters) that can inhabit locations in a two-dimensional grid environment. On each successive *clock tick* of the system (sometimes called a *move* and sometimes a *generation* in the system's life history) there are very simple rules about whether counters survive or die or whether new ones are born. These rules are applied simultaneously as follows:

1. Survival. Every counter with either two or three neighboring counters survives for the next generation.
2. Death.
  - (i) Each counter with four or more neighbors dies (i.e., is removed) because of overcrowding.
  - (ii) Every counter with one or zero neighboring counter dies from isolation.
3. Birth. Every empty cell adjacent to *exactly* three neighbors is a birth cell and a counter is created in this space on the next move.

What was so extraordinary about this system was that it was one of the first to show how complex global behavior could emerge through the behavior and interaction of a number of very simple agents. Depending on the initial conditions (which grid squares had counters at the beginning of time, referred to as  $t = 0$ ), whole societies would seemingly rise then die or oscillate between different stable states, cycling endlessly. The

interested reader is encouraged to visit (4) to download and experiment with the game itself.

### Engineering Complex Systems: Robotic Rock Collection on Mars

Rather than wondering in amazement at the emergent global behavior of various systems, it is often the case that we wish to *design* a system that has key global properties that arise from a set of simple interacting agents. One of the best examples of this was in the work of Steels (5), that was subsequently described excellently by Wooldridge (6). The problem Steels set himself can be paraphrased as follows:

Suppose we want to collect samples of a particular type of rock on another planet such as Mars. We don't know where it is, but it's typically clustered together. A number of vehicle agents are available that can drive around the planet and later re-enter a mother spaceship and go back to earth. There is no detailed map of the planet but it is known that there are rocks, hills, and valleys that prevent the agents from communicating.

In a solution to this problem, Steels makes use of an agent architecture, first proposed by Brooks, called the subsumption (7,8) architecture. In the subsumption architecture, the different behaviors of an agent are layered with those lower-level behaviors that are most critical to the agent and that take precedence over any higher-level behaviors. At the time, Brooks' work was revolutionary because most researchers believed that the only way to build robots capable of sophisticated behavior was to use techniques from artificial intelligence such as symbolic representation and reasoning. He suggested that intelligent behavior (from the perspective of the individual or group) does not necessarily require agents to have a sophisticated model of the world. He argued that intelligence could emerge from simple systems that responded to the environment by stimulus-response rules, as in the two examples we have discussed previously. In such systems, the environment provides a stimulus that causes a rule to fire and the agent responds with some behavior that in turn affects the environment and so (possibly) the future behavior of itself and other agents sharing the same environment.

He pioneered the idea that intelligent systems could be engineered in this way and that intelligent behavior is an emergent phenomenon arising from the interaction of *societies of nonintelligent systems* (7) as stated earlier. Subsumption architecture comprises eight task-achieving behaviors, each of which is implemented separately. The hierarchy of layers reflects how specific the behavior is—the more specific the task, the higher the level. In the case of the mobile robot, there are eight levels from zero to seven that relate to contact avoidance, wandering, exploring, building maps, noticing change, distinguishing objects, changing the world according to goals, and reasoning about the behavior of others.

The first step in the agent's construction is to build the zeroth control level and, once this has been tested, to build the first control level on top of the zeroth level. The first level has access to the data in level zero and can also supply its own inputs to this layer to suppress the normal activity of the zeroth layer. The zeroth level continues to execute, unaware that there is a higher level intermittently influencing its behavior. This process is then repeated for each successive layer. Subsequently, each layer competes to control the behavior of the robot.

Before Steels designed his subsumption architecture for each of the agents, he introduced a gradient field, so that the agents could always locate the mother ship. For the agent to find the mother ship, all it needed to do was move up the gradient. Then he programmed the agent as follows. The first rule at level zero in the subsumption architecture and,

therefore, the rule with the highest priority was concerned with obstacle avoidance. The other rules can then be described in decreasing order of priority as follows. Note that there are five levels.

Rule 1. If you detect an obstacle, then avoid it.

Rule 2. If you are carrying samples and at the mother ship, drop the samples.

Rule 3. If you are carrying samples and not at the base, then travel up the gradient.

Rule 4. If you detect a sample, pick it up.

Rule 5. Move randomly.

Using this set of rules, the agents are *noncooperative*. There are no interactions between them and, there is certainly nothing that looks like global emergent system intelligence. However, inspired by the ant colony example described earlier, he introduced a new mechanism. The idea was that agents would carry radioactive crumbs that could be dropped, picked up, and detected by passing agents. Using this simple technique, sophisticated cooperation between the agents could now take place. The rules are almost identical except that Rule 3 is altered and there is a new level (Rule 5) introduced just before the highest level behavior to obtain a six-level architecture. These new rules introduce *cooperation* between the agents.

Rule 3. If you are carrying samples and you are not at the base, then *drop 2 crumbs* and travel up the gradient.

Rule 5. If you sense crumbs, then pick up one crumb and travel down the gradient.

What was extraordinary about this work was that it showed that near optimal performance could occur with a collection of very simple agents. Along with Brooks, Steels was one of the first to show how intelligent systems could be designed from the emergent behavior of simple interacting agents to achieve real tasks. From an engineering perspective, the solution was also significant because it was cheap on computational resources (these agents are really *very* simple!) and it was robust. And as with our ants, one or two agents breaking down would not impact significantly on the overall system performance.

In many senses, this was an attempt to harness the power of complex systems. However, there is a very important but subtle point to realize here: the complexity was predetermined; agents were specifically engineered to achieve an overall system behavior. However, we can recapture the emergence by realizing that if the simple specification (rule set) of each of the agents was shown to an observer, all but the most experienced programmer would not anticipate the optimal, collective, cooperative system behavior. To the observer then, the system would be displaying emergence. In short, there is some sense here of emergence being a personal phenomenon, that “emergence is in the eye of the beholder” if you like. This is quite important, and we shall discuss emergence and other issues relating to complex systems that we have touched on in our three examples in the next section.

## EMERGENCE

The systems we discussed earlier are some of the key original examples of computational *complex adaptive systems* (9,10). There are many other noncomputational systems that exhibit the same kind of emergent self-organization including (11) economies, social organisations [human (12) or animal], embryologic development, the weather, traffic, ecologies, growth of cities, the rise and extinction of species, and the diversity of

immune system responses. One key factor that is common to all these systems is that there is an *emergent* self-organization arising on the macro-scale from micro-scale interactions of the individuals constituting the system.

There has been a great deal of debate about what constitutes such systems, but it is essentially the notion that some kind of *order* or *structure* or *intelligence* occurs that is not predetermined. Perhaps, the best definition of emergence (and certainly one of the most cited) is given by Cariani (13). He first describes emergence as involving “the creation of qualitatively new structures and behaviors, which cannot be reduced to those already in existence” and then goes on to describe three kinds of emergence: computational, thermodynamic, and “relative to a model.”

### Computational Emergence

The three examples we first described in this chapter can be seen as examples of *computational emergence* in which complex global behaviors or structures arise from local computational interactions. Cariani makes the point that such emergence occurs only because of the observational frame through which the system is considered (i.e., the kind of person the observer is, and the degree of technical understanding of the system's underlying algorithms they have). He argues that because there is simply a set of initial conditions and behavioral rules, there is a sense that everything is predetermined and the consequence of this is that nothing is emergent. In Cariani's view, therefore, the game of life is not displaying emergence, because as soon as you encode it in a program, you have by default defined the set of possible states for that program. Introducing stochastic (random) elements into the program does not help either, he argues, because even random elements are at some lower level of the computational process deterministic. That is, the random parameters are themselves actually generated by deterministic algorithms.

A good way to understand how something at one level can be random but at another lower level completely determined is in the tossing of a coin before a football game. The home captain tosses the coin into the air and the opposition captain calls. From the perspective of both captains, as the coin is spinning in the air, whether the coin ends up heads or tails is totally up to chance. Half the time it will be heads and half the time tails, and that is all both captains know. If, however, the opposition captain was possessed with extraordinary powers of perception and mathematical ability, they could work out the rate of spinning, gravitational pull, air resistance, wind velocity, and so on and calculate with all certainty whether the coin will land on its head or not. At this lower level then, the tossing of a coin is a deterministic process completely decided by the laws of physics. As soon as we return to the higher level of everyday human modeling and perception, we lose this determinism, and nondeterminism is reintroduced.

In order to reintroduce emergence into work from multi-agent systems and related disciplines such as Artificial Life, Cariani then moves toward introducing a more pragmatic definition of emergence as being *relative to a model*.

### Emergence Relative to a Model

While there is no “system emergence” with any computational system, there is clearly emergence occurring from the point of view of the observer, and so Cariani provides this very pragmatic view. (Recall our previous comment “emergence is in the eye of the beholder.”) Emergence arises “relative to a model,” because an observer of a computational system does not typically have a detailed view of the processes that occur inside the system. That is, the observational frame is incomplete and the observed emergent

behaviors arise because they are based on issues that are *outside of this frame*. A more succinct definition of this type of emergence is the “deviation of the behavior of a physical system from an observer’s model of it.” Cariani summarizes this category as follows:

The emergence-relative-to-a-model view sees emergence as the deviation of the behavior of a physical system from an observer’s model of it. Emergence then involves a change in the relationship between the observer and the physical system under observation. If we are observing a device which changes its internal structure and consequently its behavior, we as observers will need to change our model to track the device’s behavior in order to successfully continue to predict its actions.

### Thermodynamic Emergence

This category is a much stronger, physical view of emergence and is essentially characterized as the emergence of order from noise in the physical environment. Again though, it is the phenomenon where nondeterministic processes at the micro-level lead to structures or behaviors at the macro-level. The typical example of this type of emergence is when considering a particular gas such as oxygen. The nondeterministic behavior of electrons, atoms, and molecules somehow leads to a stable gas with well-defined properties relating to pressure, temperature, and volume at a higher level.

However, there is also some notion here of emergence relative to a model, even though Cariani sees fit to distinguish it from computational emergence. Gas is only an emergent property of molecules, because we do not fully understand how that process works. If we could understand all the laws of the universe, then getting gas from molecules might seem pretty obvious to us. Taking this view, it would seem that the only way to understand the phenomenon of emergence is by having a model of the individuals (observers) who are perceiving it.

## COMPLEX ADAPTIVE SYSTEMS WITH MULTIPLE INTERACTING AGENTS

Although there is no agreed upon definition of exactly what constitutes either a complex system (14), emergence (15), or even what an agent is (16), there are a number of general properties that we list and outline here that have been instrumental in guiding our work in modeling the society of stem cells:

- Order is emergent rather than predetermined.
- The system’s future is, in general, unpredictable.
- The basic entities of a complex system are agents. There is a huge debate that has been raging for years about what constitutes an agent (16) but, in this context, they are autonomous or semi-autonomous entities that seek to maximize some measure of usefulness (this could relate to a goal, motivation, or utility) by responding to the local environment according to a set of rules that define their behavior. (This is sometimes referred to in more economically-biased accounts as the agent’s *strategy*.)
- The individuals are not aware either of the larger organization or its goals and needs. Clearly any single agent within the system cannot know the state and current behavior of every other agent and, as a result, cannot determine its behavior based on such complete global system information. Instead, the behavior of agents is governed by rules based on the local environment.

- Agents typically follow reactive rules that are typically of the form: if *condition* then fire *action*. For example, a possible rule for an agent might be *if there's a space next to me and I am currently too hot, then move into an empty space*.
- Agents can perceive aspects of the environment they are in and can act so as to change the state of the environment. Critically, in a complex system, agents must affect the environment in such a way that the environmental change can (i) be perceived by others and (ii) affect the behavior of others. (Recall the difference between the noncooperative and cooperative versions of the robot vehicle rock collectors.) That is, the agents must have a reasonable degree of *interaction* beyond that of, say, simple obstacle avoidance (17).
- The rules of an agent will often contradict and there must be some mechanism (possibly nondeterministic) for selecting from competing behaviors/rules. In general, the behavior of an agent will not be deterministic.
- Agents may be equipped with the ability to adapt and learn rules so as to have a more effective way of maximizing their usefulness in given situations. New rules would try, for example, to make the agent more able to act effectively in a wider variety of environmental situations.
- Rules may compete for survival. The more a rule is used in determining behavior, the greater the chance that it has of surviving in the future. Rules that are seldom or never used will have less chance of survival. Rules may change randomly or intentionally and may be integrated for more sophisticated action.
- Agents are resource-bounded and can only perceive (or experience) their *local* environment. Typically, they will also have the ability to determine what to do next when there is incomplete or contradictory sensory information about their local environment. It is also possible that agents might have access to some global information.
- A few individuals or agents will not make a sustainable complex system. It is only when the number of agents reaches a certain critical threshold that the system will exhibit global, meaningful behavior.

Using these basic principles, we have built and are currently implementing a model of stem cells and cell lineages. There are also several others who have done similar work that we discuss briefly here.

## MODELING STEM CELLS AND CELL LINEAGES AS COMPLEX ADAPTIVE SYSTEMS

Although mathematical modeling of stem-cell lineage systems is critical for the development of an integrated attempt to develop ideas in a systematic manner, it has not been a research area that has received a large amount of attention. Over the last year or so, there has been a noticeable climate change in this respect, and there is now a growing awareness of the need to use mathematical modeling and computer simulation to understand the processes and behaviors of stem cells in the body. Some reasons have been pointed out by Viswanathan and Zandstra (18) in an excellent survey of mathematical techniques for predicting behaviors of stem cells. We summarize the key points here:

- In the adult body, stem cells cannot be distinguished morphologically from other primitive nondifferentiated cell types.
- Extracting stem cells from an embryo means sacrificing it, posing serious ethical difficulties.



- There is no way to determine whether any individual isolated cell is a stem cell and to be able to model what its potential behavior might be. It is not possible to make any definite statements about this cell. At best, it can be tracked and its behavior observed, though clearly this behavior is simply one of many possible paths. The notion of a stem cell refers to the wide-ranging set of potential behaviors that it might have that are influenced by internal, environmental, and stochastic processes.
- The number of possible interactions and behaviors of a large number of stem cells makes the system extremely complex in all the senses described earlier. Theoretical simplifications are key to understanding fundamental properties.

There is, thus, a need for new theoretical frameworks and models that can be directly mapped to a computer simulation and that look at the dynamic self-organization of stem cells.

Before introducing a summary of our own work, we will consider some related theoretical investigations. We first introduce the recent work of Agur et al. (19) as it is very similar, algorithmically, to the game of life introduced earlier in this paper.

### A Simple Discrete Model of Stem Cells

In their recent work, Agur et al. used a model very similar to that of the game of life to understand what mechanism might be employed for maintaining the number of stem cells in the bone marrow and producing a continuous output of differentiated cells. This work is important, because it is one of the few examples where a mathematical model has been used to show what properties of stem cells might be required to enable the maintenance of the system's homeostasis.

Essentially, they model a niche as having the ability to maintain a reasonably fixed number of stem cells, to produce a supply of mature (differentiated) cells, and to be capable of returning to this state even after very large perturbations that might occur through injury or disease. The behavior of a cell is determined by both internal (intrinsic) factors (a local clock) and external (extrinsic) factors (the prevalence of stem cells nearby), as stated by the authors as follows:

1. Stem-cell behavior is determined by the number of its stem-cell neighbors. This assumption is aimed at simply describing the fact that cytokines secreted by cells into the micro-environment are capable of activating quiescent stem cells into proliferation and differentiation.
2. Each cell has internal counters that determine stem-cell proliferation, stem-cell transition into differentiation, and the transit time of a differentiated cell before it migrates to the peripheral blood.

The niche is modeled as a connected, locally finite, undirected graph, but for most intents and purposes, we can visualize this as a two-dimensional space made up of grid squares as in the game of life. Their model certainly applies for this topology.

Any grid square is either empty, or it is occupied by either a stem cell or a differentiated cell. A stem cell is able to interpret messages from neighboring locations (horizontal or vertical, not diagonal) such that it knows what is at those locations. Stem cells can divide into two stem cells (called proliferation) or become determined cells (no division takes place). Determined cells stay in the niche for a period and then eventually leave to enter the bloodstream.

There are three constant values (let us call them  $N_1$ ,  $N_2$ , and  $N_3$ ) that are used to reflect experimental observation. The first constant ( $N_1$ ) represents the time taken for a differentiated cell to leave the niche. The second ( $N_2$ ) represents the cycling phase of a stem cell; a certain number of ticks of the clock are needed before the cell is ready to consider dividing. Finally, the third ( $N_3$ ) represents the amount of time it takes for an empty space that is continuously neighbored by a stem cell to be populated by a descendent from the neighboring stem cell. The rules of the model, expressed in simple English, are as follows:

#### Rule for determined cells

1. If the internal clock has reached  $N_1$ , then leave the niche. Reset local clock to 0.
2. If the internal clock has not yet reached  $N_1$  then increment is 1.

#### Rule for stem cells

1. If the counter at a stem-cell location has reached  $N_2$  and all stems are neighbors, then become a differentiated cell. Reset the clock to 0.
2. If the counter of a stem cell is equal to  $N_2$ , but not all the neighbors are stem cells, then do nothing. Leave clock unchanged.
3. If the counter has not reached  $N_2$ , then do nothing except increment the clock.

#### Rule for empty spaces

1. If the counter at an empty grid has reached  $N_3$  and there is a stem-cell neighbor, then introduce (give birth to) a stem cell in that location. Reset clock.
2. If the counter at an empty grid has not reached  $N_3$  and there is a stem-cell neighbor, then increase the clock.
3. If there are no stem-cell neighbors at all, then reset the clock to 0.

A move is just as it is in the game of life; the *next state* of the system is a function of the clock, the state of the cell, and the state of the neighboring cells. All locations are then updated simultaneously as before. As with the game of life, there are no stochastic elements. The only real difference is that the agents have a local state (clock) that is not present in the game of life. What is remarkable here is that this simple model allows for sophisticated global behaviors to arise. All the basic *common sense ground rules* of stem cells to proliferate, to remain quiescent, and to produce continuous supplies of differentiated cells can be found in all the possible behaviors of this systems (That is, if you discount extreme situations, such as where all the grids are occupied by stem cells.). Moreover, there is always a sufficient density of stem cells in the niche and the system never dies out.

Although there are a number of difficulties with this work, in particular the fact that it requires spaces in the niche to have counters (rather than the cells as originally set out in the text), it is one of the few attempts to capture observable qualities of stem-cell systems in a simple mathematical model that can be simulated computationally. It is a very simple model and, as a result, the authors were able to mathematically *prove* many properties of their system, and their results are extremely important for paving the way for a more sophisticated analysis of various stem-cell-like properties in the future.

What is perhaps most extraordinary about this work is how similar the basic algorithms are to that of the game of life and the fact that it took 30 years to get from a cute mathematical game to cutting-edge work on the theoretical modeling of stem cells!

### Plasticity and Reversibility in Stem-Cell Properties

From a biological viewpoint, the model of Agur et al. does not allow any reversibility or plasticity in the basic properties of cells. For example, once a cell has differentiated, it cannot become a stem cell again (or, in a more continuous view, more plastic). Moreover, once a cell has left the niche, it cannot return. A recent example, an approach that uses a more sophisticated model and addresses these issues, is that of Loeffler and Roeder at the University of Leipzig, who model hematopoietic stem cells using various (but limited) parameters including representing both the growth environment within the marrow (one particular stem-cell niche) and the cycling status of the cell (20–22). The ability of cells to both escape and re-enter the niche and to move between high and low niche affinities (referred to as within-tissue plasticity) is stochastically determined. The validity of their model is demonstrated by the fact that it produces results in the global behavior of the system that exactly matches experimental laboratory observations. The point is that the larger patterns of system organization emerge from these few simple rules governing variations in niche-affinity and coordinated changes in cell cycle.

Another example, also from Loeffler, working with colleagues Potten and Meineke, models movement and differentiation of small intestinal stem cells from the stem-cell niche to the villous tip in a two-dimensional lattice-free cylindrical surface (23). In this model, cells interact by viscoelastic forces. Simulations were compared directly with experimental data obtained from observations of cells in tissue sections. These showed that the model is consistent with the experimental results for the spatial distribution of labeling indices, mitotic indices, and other observed phenomena using a fixed number of stem cells and a fixed number of transit cell divisions. Moreover, the model suggested a gradient, perhaps a diffusable protein, which could explain differentiation of cells as they moved up the villus. Thus, not only did the model fit experimental data already in hand, but it made predictions that could form the basis of new investigations.

### An Agent-Based Approach to Modeling Stem Cells

In our current work, we are building a more comprehensive formal model of cells as reactive agents responding to local environmental factors that can maintain some balance of cells under various conditions, using the criteria outlined in Section 3. The intention is to provide a toolkit for researchers and students to investigate behaviors of stem-cell systems, given a set of rules, environmental influences, and so on. As with Roeder and Loeffler (22), we will also allow reversibility, plasticity, and nondeterminism, but we model a greater number of internal and environmental parameters. As with the related work we have discussed in this chapter, there is something in common here with the “reverse engineering” approach of Steels and his design of robots to achieve an overall system behavior. What we wish to do is build our model of a stem cell in such a way that the overall system behavior has many of the observable qualities viewed in current medical experiments.

Currently, all cells are modeled as agents with *identical* abilities, perceptual capabilities, and rules. In line with Roeder and Loeffler (22), we see “stemness” not as a “yes” or “no” quality of any given cell, but as a continuum of potential behaviors. The more a cell has stemness, the more likely a cell is to behave in a stem-like way. The agent model details how the internal state, the local environment (proteins, populations, fluid pressure, and so on) affect the probabilities of behaving in certain ways, such as moving to or from a niche and cell division.

Their state will include information on how many divisions can occur, how likely the cell is to stay in the niche, whether it's more or less likely to divide, whether division is symmetric or asymmetric, how sensitive the cell is to protein signaling in the micro-environment, how likely it is to react to them once they are sensed, and so on. They can also perceive local environmental conditions, such as the relative concentrations of other stem cells and cells at various stages in the set of available lineages we model, as well as various signaling proteins such as SDF1. Their behavior will be nondeterministic and based on their current state and the state of the local environment. Emerging from this nondeterministic micro view, we expect a stable dynamic system that can be re-instantiated even after traumatic events.

Suppose, for example, that a certain stem cell is in a given environment. At any stage, it will have some probability of dividing and, if it does, there will be some probability of producing a daughter cell along one lineage and another probability of producing a cell along an entirely different lineage, and so on. Because we cannot say for certain what will happen—sometimes one action will happen and, in exactly the same situation, sometimes something else might happen—we introduce randomness or, more formally speaking, nondeterminism into the system. As we stated earlier, most commentators argue that some degree of nondeterminism is needed for a system to be a complex adaptive one. Self-organization fails to emerge in completely determined systems.

Once we simulate these agents and run the system so that it is in a kind of stable equilibrium (of the kind exemplified by the game of life example we described earlier), we can then consider the effects of disease and life-threatening environments. It will be possible to model and investigate what kinds of behavioral changes, even of a single stem cell because of a chance mutation in its rule base, might lead to system imbalance or collapse. In other words, do precisely defined rule changes mimic known clinical conditions?

This conceptual approach thus focuses on the fact that complex, adaptive systems typically have multiple equilibrium states where, for example, the number of agents of a particular kind may be kept constant. An equilibrium can be more or less stable: a very stable equilibrium needs massive events (either internal or external) to affect it while a nonstable one can be upset by relatively small events. These less-stable equilibria are more dangerous for the safety of a system, as a tiny event may lead to massive system change or total system collapse. Commonly cited examples are mass extinctions of species, collapse of stock markets, and the demise of cultures and civilizations. It is often changes in the interactions or behavior at the micro-level that affect phenomena such as mass extinctions.

Analogously, the failure of stem-cell systems is sometimes not merely due to the size of the internal or external change: it may be simply a necessary result of the generally high durability and sustainability, but complexity, of the cell system. Aplastic anemia, for example, a complete failure of the hematopoietic system, may not have a specific precipitating event. Likewise, acute hepatitis A is usually benign and self-limited, but a very few infected people suffer massive hepatic necrosis leading to death or the need for transplant. The unpredictability of these events may relate to our limited understanding of pathogenesis, but it might instead be inherent because the stem-cell system is a complex one.

In addition, we can use the formal model of our stem-cell complex system to build a computer simulation of a large system of stem cells and progeny. As we described earlier in this chapter, computer simulations have been very successful in showing how emergent, global, self-organizing properties can arise through very simple descriptions of individual behavior. However, the computational demands required to model large systems are enormous, and we will have to use a grid cluster to perform this simulation of hundreds of interaction agents.

The formal specification of our model is now being used as the blueprint to build the simulation. Using logic we can prove that the simulation implements the model exactly

and completely. As far as we are aware, this has not been achieved before and it means that observed events produced by running the simulation can be carefully interpreted within the semantic context of the formal structured model. Careful statistical analysis of our simulated clinical events may shed a new and very different light on these dire occurrences.

## RAMIFICATIONS FOR CURRENT THINKING

A new perspective from which to debate some of the key contested issues of adult stem-cell research arises when one considers cell lineages from a complex systems approach. For example, the longstanding debate as to whether stem-cell lineages are determined or stochastic processes becomes clear (22). As in the work of Loeffler and colleagues, the stochastic elements are required to obtain the observable results. Indeed, the increasing number of articles on the reversibility of gene restriction makes the stochasticity of lineage fate unavoidable in conceptualizing issues of cell plasticity (24,25). These theoretical notions are backed up by the results of both clinical studies (26) and single-cell culture and gene-expression experiments (27,28) where a greater variability of gene-expression pathways is revealed than would be expected from a complete determinism.

In this paper, we do not go into a detailed discussion of these findings, but we wish to make it clear that the evidence now strongly indicates a nondeterministic view. This is crucial for our complex-system interpretation to be appropriate: if we conceive of cell lineages as complex and adaptive, then stochasticity is implicit because fluctuations are necessary for self-organizing systems to explore new possibilities.

Another current controversy concerning adult stem-cell lineages relates the often low engraftment from bone marrow into other system organs: often less than five percent, sometimes less than one percent, in the absence of overt, severe injury (29,30). Some have argued that even if bone marrow plasticity can be demonstrated, such low levels of engraftment from the blood are physiologically trivial and insufficiently robust to be of relevance to tissue maintenance (31).

However, if we consider these alternate lineage phenomena as parts of a complex adaptive system, it reveals to us that the converse is more likely to be true. The documented low level of apparently random fluctuation, this “quenched disorder” that we mentioned earlier, is precisely what allows the system to be adaptive. Going back to our ants example, it is only the small percentage of ants straying from the main path that enables the formation of new paths to food in the event that the current line becomes interrupted or the food source runs out. From the complex system point of view, the low-level engraftment fluctuations are critical: without them, robust responses to injury might not be so efficient or even possible.

It is precisely this intermediate level of stochastic variation, somewhere between a fully determined system (where all events can be predicted; the behavior of each element is simply a function of the current state of the whole system) and a totally nondetermined system where any event can happen at any time (referred to as “chaos”) that makes cell lineage systems, and therefore our own bodies, complex, adaptive, and alive.

## CONCLUSION

We believe that recent experimental evidence makes it clear that it is increasingly necessary to use formal, computational models to investigate the nature of stem-cell systems rather than stem cells in isolation. There are several key reasons. First, adult stem cells

cannot be easily isolated; indeed, it may be that it is only by looking at their behavior in a system, not in isolation, that we can tell what kind of cell we were originally looking at. Second, even if we were able to track the behavior of a cell in the body, it would only tell us about one of the possible behaviors of the original cell; it tells us nothing about the potentially infinite array of behaviors that may have been possible if the environment and the chance elements had been different. Third, there is evidence to suggest that mechanical forces on cells are critical in determining stem-cell behavior. If this is the case, then any act of withdrawing cells from the original system would potentially affect that cell irrevocably (32). Fourth, by removing a cell from its original and natural habitat, the new environmental conditions will influence future behavior and lead to misleading results. Fifth, it is the totality of the stem cells as a *system* in the human body that is important. A key quality of the system is its ability to maintain exactly the right production of cells in all manner of different situations.

In response, therefore, we have developed a formal model that reflects many of the key experimental and recent theoretical developments in stem-cell research. Using techniques from multi-agent systems, we are currently building a complex, adaptive system to simulate stem-cell systems in order to provide a testbed from which to be able to investigate their key properties in general and to formulate new experiments to identify the underlying physiological mechanisms of tissue maintenance and repair.

## ACKNOWLEDGMENTS

The team of collaborators in this project (entitled CELL) also included the curator Peter Ride and the A-life programmer Rob Saunders (who was instrumental in helping us form some of the views here regarding emergence), both from the University of Westminster.

## REFERENCES

1. Luck M, d'Inverno M. A formal framework for agency and autonomy. In: Proceedings of the First International Conference on Multi-Agent Systems. San Francisco, CA: AAAI Press/MIT Press, 1995:254–260.
2. Odell J. Agents and complex systems. *J Object Technol* 2002; 1(2):35–45.
3. Gardner M. The fantastic combinations of John Conway's new solitaire game, 'Life.' *Sci Am* 1970; 223(4):120–123.
4. <http://www.bitstorm.org/gameoflife/>
5. Steels L. Cooperation between distributed agents through self-organization. In: Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World. Holland: Elsevier Science Publishers, 1990:175–196.
6. Wooldridge M. *An Introduction to MultiAgent Systems*. New York, NY: Wiley, 2002.
7. Brooks RA. Intelligence without reason, Proceedings of 12th International Joint Conference on Artificial Intelligence. Sydney, Australia, 1991:569–595.
8. Brooks RA. Robust layered control system for a mobile robot. *IEEE J Rob Autom* 1986; 2(1): 14–33.
9. Lewin R. *Complexity: Life at the Edge of Chaos*. 2nd ed. Chicago IL: University of Chicago Press, 2002.
10. Johnson S. *Emergence*. New York, NY: Scribner, 2001.
11. Walker A, Wooldridge M. Understanding the emergence of conventions in multi-agent systems. In: *ICMAS95* 1995:384–389.
12. Gilbert N, Conte R. *Artificial Societies: The Computer Simulation of Social Life*. London: UCL Press, 1995.

13. Cariani P. Emergence and artificial life. In: Langton CG, Taylor C, Farmer JD, Rasmussen S, eds. *Artificial Life II*. Boston, MA: Addison-Wesley Longman Publishing Co., 1991:775–797.
14. Holland J. *Emergence: From Chaos to Order*. Oxford: Oxford University Press, 2000.
15. Shoham Y, Tennenholtz, M. On the emergence of social conventions: modeling, analysis, and simulations. *Artif Intell* 1997; 94(1):139–166.
16. d’Inverno M, Luck M. *Understanding Agent Systems*. 2nd ed. Berlin Heidelberg, New York: Springer, 2003.
17. d’Inverno M, Luck M. Understanding autonomous interaction. In: Wahlster W, ed. *ECAI ’96: Proceedings of the 12th European Conference on Artificial Intelligence*. Budapest, Hungary: John Wiley and Sons, 1996: 529–533.
18. Viswanathan S, Zandstra PW. Toward predictive models of stem cell fate. *Cytotech Rev* 2004; 41(2/3):1–31.
19. Agur Z, Daniel Y, Ginosar Y. The universal properties of stem cells as pinpointed by a simple discrete model. *J Math Biol* 2002; 44(1):79–86.
20. Loeffler M, Roeder I. Tissue stem cells: definition, plasticity, heterogeneity, self-organization and models—a conceptual approach. *Cells Tissues Organs* 2002; 171:8–26.
21. Roeder I. *Dynamical Modelling of Hematopoietic Stem Cell Organisation*. Ph.D. dissertation, Leipzig University, 2003. [http://people.imise.uni-leipzig.de/ingo.roeder/diss\\_ingo.pdf](http://people.imise.uni-leipzig.de/ingo.roeder/diss_ingo.pdf) (accessed July 2004).
22. Roeder I, Loeffler M. A novel dynamic model of hematopoietic stem cell organization based on the concept of within-tissue plasticity. *Exp Hematol* 2002; 30:853–861.
23. Meineke FA, Potten CS, Loeffler M. Cell migration and organization in the intestinal crypt using a lattice-free model. *Cell Prolif* 2001; 34(4):253–266.
24. Theise ND, Krause DS. Toward a new paradigm of cell plasticity. *Leukemia* 2002; 16:542–548.
25. Theise ND. New principles of cell plasticity. *C R Biol* 2003; 325:1039–1043.
26. Thornley I, Sutherland R, Wynn R, Nayar R, Sung L, Corpus G, Kiss T, Lipton J, Doyle J, Saunders F, et al. Early hematopoietic reconstitution after clinical stem cell transplantation: evidence for stochastic stem cell behavior and limited acceleration in telomere loss. *Blood* 2002; 99:2387–2396.
27. Madras N, Gibbs AL, Zhou Y, Zandstra PW, Aubin JE. Modeling stem cell development by retrospective analysis of gene expression profiles in single progenitor-derived colonies. *Stem Cells* 2002; 20:230–240.
28. Krause DS, Theise ND, Collector MI, Henegariu O, Hwang S, Gardner R, Neutzel S, Sharkis SJ. Multi-organ, multi-lineage engraftment by a single bone marrow-derived stem cell. *Cell* 2001; 105:369–377.
29. Wagers AJ, Sherwood RI, Christensen JL, Weissman IL. Little evidence for developmental plasticity of adult hematopoietic stem cells. *Science* 2003; 297:2256–2259.
30. Weissman IL, Anderson DJ, Gage F. Stem and progenitor cells: origins, phenotypes, lineage commitments, and transdifferentiations. *Annu Rev Cell Dev Biol* 2001; 17:387–403.
31. Eastwood M, Mudera V, McGrouther D, Brown R. Effect of mechanical loading on fibroblast populated collagen lattices: Morphological changes. *Cell Motil Cytoskel* 1998; 40:13–21.
32. Ogawa M. Stochastic model revisited. *Int J Hematol* 1999; 69:2–5.