

James Odell
 jodell@compuserve.com
 www.jamesodell.com

Agents: Between Order and Chaos

NATURE HAS MOMENTS both of order and chaos. Interestingly enough, those forms that are considered most fit actually reside someplace in between. This phenomenon applies to business and software agents, as well.

Basic Agent Behavior One of the earliest form of agents is called a *cellular automata* (CA). The idea was originally conceived by the Polish mathematician Stanislaw Ulam in the early 1950s and further developed by John von Neumann and Arthur Brooks. Basically, a CA consists of a lattice of cells, or sites. Each cell has a state whose value is commonly expressed as 0 or 1, black or white, on or off, or a color selected from a set of colors. At discrete “ticks” of the CA clock, this value is updated according to a set of rules that specifies how the state of each cell is computed from its present value and the values of its neighbors.

The most familiar example is John Conway’s game, Life. As described in the October 1970 issue of *Scientific American*, only a checkerboard and an ample supply of markers are needed. The rules of Life are simple:

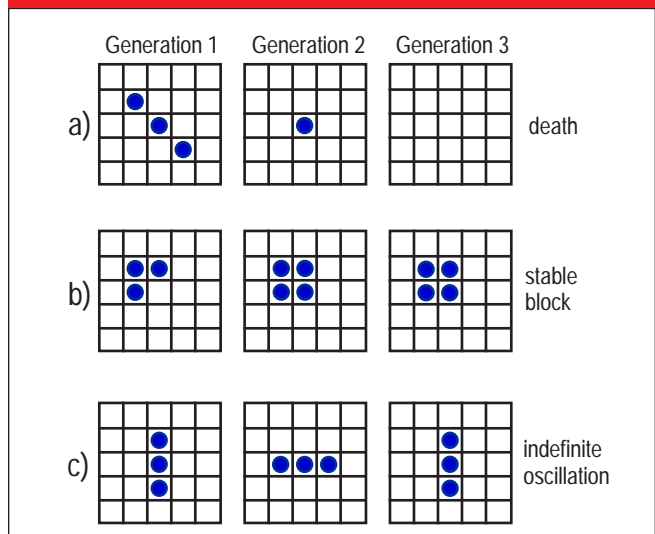
1. A dead cell (state 0), with exactly three of its eight immediate neighbors alive (state 1), is born. Under the right conditions, the cell comes alive.
2. A living cell with two or three living neighbors remains alive, that is, the cell stays alive when nurtured by its neighbors to the right extent.
3. All other cells die (or remain dead) due to overcrowding or loneliness.
4. Each cell is updated once per time period.

The checkerboard rules represent the laws of physics (or life) and, while the cells themselves are not mobile, an amazing amount of behavior emerges. Figure 1a depicts how a CA society can die out over three generations. Figure 1b, on the other hand, shows how a society can form a fixed configuration. Lastly, Fig. 1c illustrates how some patterns oscillate indefinitely.

Classifying Agent Behavior Over the long run, CA societies have similar kinds of emergent behavior. The patterns of Fig. 2 illustrate the four classes of behavior identified by Stephen Wolfram in 1983, when he was at Princeton’s Insti-

James Odell is a consultant, educator, and author. He works with IntelliCorp and James Martin & Co.

Figure 1. Some examples of Life patterns.

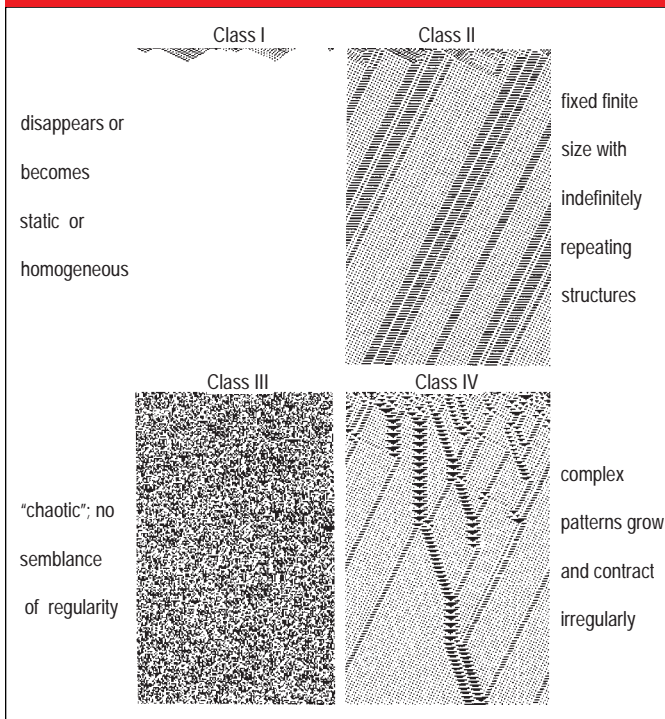


tute for Advanced Studies. Class I societies are those that exhibit a static, or *limit point*, behavior. Figures 1a and 1b are examples of this class, because the lattice will not change after generation 3. Class II societies exhibit periodic, or *limit cycle*, behavior which is the indefinite oscillation depicted in Fig. 1c.

Class I and II can be considered one extreme of CA behavior because everything is predictable and orderly. Class III on the other hand is aperiodic, or *chaotic*; that is, its structures display no obvious order or uniformity. In between these extremes, is a mysterious and complex class of behavior: Class IV. Such automata exhibit considerable local organization, yet also have areas of irregular behavior. In other words, Class IV automata are some place in between the two extremes: they exhibit orderly behavior as well as some chaotic behavior. (Images are courtesy of Andrew Wuensche, generated using his software “Discrete Dynamic Lab” from <http://www.santafe.edu/~wuensch/ddlab.html>.)

Order versus Chaos Cellular Automata offer a way to model natural and artificial processes, such as modeling crystallization, complex fluid flows, chemical reactions, and hardware architecture. Yet, CA involve an elementary form of agent. Imagine the kinds of systems that can be built with agents that are mobile and have sophisticated forms of communication and interaction. Such agent systems not only pro-

Figure 2. Wolfram's four classes of long-run behavior.



vide a richer way of modeling natural and artificial processes but provide a way of *implementing* such systems, as well.

Such mature agents systems are still subject to the same Wolfram behavior. You can build agents systems that are orderly (Class I and II), and such orderly behavior is appropriate for some kinds of systems. However, when agents are expected to learn and change their behavior, an orderly system discourages change. In a business example, all the jobs in an organization would be subdivided so that employees have no latitude and only do the job for which they are hired.

For an automated supply-chain system, the results would always be predictable. On the surface, such predictability would seem to be a good thing. However, when the business landscape changes (as it often does), the supply-chain operation would no longer suit the organization's needs. Instead, it would be predictably wrong. In both of these scenarios, everybody would benefit if the individual agents had the freedom to change. In short, orderly agent systems should become more fluid—and a bit closer to chaos.

Conversely, if agents are deep in a chaotic regime (Class III), they can never get the job done. For example, employees who do not know what they're supposed to do half the time end up working at cross purposes. A supply-chain system would not be able to deliver the right product, to the right person, at the right time. In both of these scenarios, if the individual agents could have tighter connections with fewer individuals, a greater degree of stability would be introduced. Chaotic agents, then, should become less fluid by adapting to what other agents are doing, resulting in aggregate behavior. This means pulling back from chaos.

The Edge of Chaos Neither order nor chaos seems to be the best place for complex systems—whether their agents are implemented using software, hardware, machines, or people. Instead, such agent systems need to be someplace in between. With too much order, the system stagnates and dies in the face of new competition that needs to be only a little bit better. With too much chaos, the system will not survive because it can not make useful products. The edge of chaos is on average where fitness is best (Fig. 3). Such systems can exploit what they have learned, as well as extend that learning through exploration.

Complex systems (including both living and business systems) are characterized by perpetual novelty. This approach can be scary: things can get out of control and errors will be made. Yet without this kind of approach, there will be no change—only *status quo*. To talk about complex adaptive systems being in equilibrium is meaningless because the system never gets there. It is always unfolding, always in transition. If a system ever reaches equilibrium, it is not just stable—it is dead.

Now, I am not suggesting that such complex systems be built immediately. In fact, this would probably result in chaos itself. Complex systems should be built simply at first, initially (and gradually) placing any edge-of-chaos processing with human agents rather than automated agents.

The reasons for this are technical and psychological. Technically, we do not yet fully understand how to build complex systems that function properly. We lack both a systematic methodology and industrial-strength, agent-system toolkits. Psychologically, living on the edge of chaos can make us uncomfortable. And, when we must delegate our tasks to automated agents, we will feel even more out of control. It's bad enough when people are intimidated by their home appliances. What will happen when automated agents choose the article we read, automatically answer our mail, and schedule appointments? On top of this, imagine how uneasy we will feel when automated agents begin making critical *business* decisions and acting on them. Confidence and understanding come slowly.

Figure 3. Systems poised between order and chaos are best able to carry out ordered yet flexible behavior.



Conclusion Stability is probably something valued in accounting and payroll systems. However, the next generation business systems should be operating on the edge of chaos. Order entry, inventory control, and supply chain systems are particularly appropriate. These are systems whose agents are people, machines, and software. To work effectively, these agents must work together as a living system: requiring flux, change, and the forming and dissolving of patterns.

Complex systems theory points us away from isolated units and toward interactions between individuals and their environment. Strategy focuses on the management of volatility, not the achievement of specific goals. Growth comes from agent relationships and rules rather than through a significant increase in size. Opposing thoughts or points of view are held simultaneously. Mild instability is encouraged. Build something workable, rather than “optimal.” Developing complex systems is not for the faint-hearted, which applies to both executives and IT system developers. We need to unleash our software and let it grow and learn like any living system. Only then can our systems mature beyond our limitations—and exceed our expectations.

A greater kind of courage and a different psychology is now required—to be willing to let go and experience the creativity, innovation and disturbance which comes about when we risk the outer boundaries of trying to maintain a balance and the excitement of living, developing and coaching at the “edge of

chaos.” Learning will perhaps ultimately prove less valuable in the third millennium than the skill and attitudes of unlearning—in the same way that knowing what to do may become far less important than knowing what to do when you no longer know what to do. (Petruska Clarkson, psychologist, a chartered consultant in UK) ♡

References

1. Kauffman, Stuart, *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*, Oxford University Press, New York, 1995.
- 2). Levy, Steven, *Artificial Life: A Report from the Frontier Where Computers Meet Biology*, Vintage Books, New York, 1992.

Acknowledgments The author would like to thank David Taylor for his helpful suggestions in producing this column.