How Complex Systems Learn and Adapt

Marilyn J. Darling, *Fourth Quadrant Partners*

This article was written to accompany a report on research conducted by Fourth Quadrant Partners (4QP): A Whole Greater than Its Parts: Exploring the Role of Emergence in Complex Social Change[1]. Its purpose is to give a fuller explanation of complex adaptive systems (CAS) theory that lies at the foundation of this research project. The article provides a layman's description of what CAS researchers have discovered about how complex systems learn and adapt, and provides an example to help readers see how the theory plays out in a social change context.

John Holland was a complexity scientist who devoted his career to answering the question: *What is the difference between those complex systems that adapt relatively quickly and those that do not?* He is credited with launching the field of complex adaptive systems (CAS) theory. He and his colleagues have studied everything from natural ecosystems to immune systems to human organizations to cities and economies.

Holland opened his 1998 book, *Emergence: From Chaos to Order*, by telling the story of how, in the 1950s, his colleague, Arthur Samuel, programmed one of the earliest digital computers to play checkers. Holland describes in great detail how Samuel was able to create a learning procedure that allowed the program to learn, through iteration, not only how to play checkers, but how to beat Samuel and, ultimately, to win against champion players.[2] (See sidebar: "CAS theory and deep learning.")

Holland started his discussion of emergence with the story about a computer learning algorithm because it illustrates what he observed through his research about how complex systems learn and adapt.

In complex adaptive systems, individual actors—referred to as *agents* in CAS theory—interact with each other and their environment as they seek to achieve a goal. As experience accumulates, they begin to notice patterns in their interactions. The more diverse the agents, the more diverse are their experiences. The more frequent the interactions among these diverse agents, the faster patterns become evident. Once a pattern has been discovered, it gives agents a way to begin to anticipate how their environment operates and develop hypotheses about cause and effect.[3] Agents can use this information to help them achieve their goal, even when everything else about the environment remains, in Holland's words, "perpetually novel."[4]

CAS theorists refer to these discoveries about cause and effect related to patterns as *building blocks*. Building blocks allow agents to learn to navigate hugely complex environments that are never the same twice. By collecting and experimenting with how to use combinations of building blocks, agents get better over time at consistently achieving their goal, even as conditions change. The more often a person travels internationally, the easier it becomes to navigate airports in different countries, with different procedures, even when she does not know the language.

Using building blocks like this gives us a platform to begin to understand other aspects of the system in a new way.

1 This research project was made possible by the generous support of the William and Flora Hewlett Foundation, the David and Lucile Packard Foundation, and the John S. and James L. Knight Foundation. Thanks also to Jillaine Smith, Heidi Sparkes Guber, and Ray Gordezky for their input. Readers may also want to read our article, "Emergent Learning: A Framework for Whole-System Strategy, Learning, and Adaptation" for an explanation of the difference between adaptive and emergent strategy.

2 John Holland, *Emergence: From Chaos to Order* (Perseus Books, 1998), 64.

3 Holland, *Hidden Order: How Adaptation Builds Complexity* (Perseus Books, 1995), 87-90.

4 Holland, *Hidden Order*, 34.

CAS theory and deep learning

Using today's massive computing power, Alphabet's subsidiary, DeepMind, created a "deep learning" program in 2014, AlphaGo, which started with a database of thousands of human Go games and a learning algorithm akin to Samuel's checkers-player program from the 1950s. Though Go is an infinitely more complex game, AlphaGo was able to beat champion players in 2015 and 2016.[5] Meanwhile, DeepMind, created AlphaGo Zero, which simply started to play the game randomly against itself, without the benefit of a database of previous games.[6] Counterintuitively, without the benefit of this large database of games, AlphaGo Zero was able to begin winning games within three days. Within 40 days, it had surpassed the skill of all of the previous versions of AlphaGo. Most recently, DeepMind launched Alpha Zero[7], which learned to master three *different* games (Shogi, Chess, and Go), outperforming existing game-playing programs within 24 hours.[8]

Holland asserts that the discovery of building blocks paves the way for innovation. He describes, for example, how the building blocks of a language—alphabets, punctuation, rules of grammar and syntax, "yield a never-ending flow of innovations."[9] Using what I know about the English language, I can say something that has never been uttered in the history of humankind and anticipate that I will be understood by someone who also speaks English. Another powerful example:

> Over the years, using real-time data, meteorologists began to see patterns in ocean temperatures, upper-air pressure centers, wind, and solar radiation. They began to study how these patterns interact, which led them to discover a powerful new building block: the jet stream. They began to observe how it behaved and hypothesize how it might affect weather patterns. This exploration became a platform for discovering new patterns and led to a level jump in their ability to predict local weather and longer-term forecasting of weather patterns, such as El Nino/La Nina cycles, drought and hurricane seasons. It has also become an important building block for aviation planning.

How does this relate to complex social change? Imagine an annual national conference on childhood literacy, in which change agents from across the United States and Canada who care about literacy—teachers, parent activists, nonprofit leaders, funders, clergy, school administrators, lawmakers—convene to hear the latest findings from researchers and presentations from their peers on how to increase the number of children who are reading at grade-level by the end of third grade.

If we asked individuals in this community of change agents what patterns they have observed in their work, they might name school issues such as classroom size, time constraints, school mandates, summer slide, or need for more innovative literacy programs and reading coaches to address learning styles; or more endemic social issues affecting early childhood development such as nutrition, effects of poverty and family stress, the number of parents in the home, early exposure to language and access to books, and other patterns that they hypothesize are predictors of literacy. Many could cite research data that validates these patterns and the cause-effect relationship they have on third grade literacy.

Yet childhood literacy remains an elusive goal. It's a nonlinear problem. The many factors that contribute to a child's being able to read at grade level interact in complex ways. Participants at that conference likely would not agree with each other on which patterns are most important, nor on their hypotheses regarding the most important actions to take to boost literacy. Speakers at the conference might advocate for their own discoveries and some participants might come away inspired to try a new approach to improve literacy in their own sphere of influence.

But what if the organizers of this annual conference explicitly used CAS theory to help improve the adaptive capacity

5 For more about AlphaGo, visit https://deepmind.com/research/alphago/, accessed April 4, 2018.

6 Described in detail in David Silver, et al., "Mastering the game of Go without human knowledge," in *Nature*, vol. 550 (19 Oct 2017); https://tinyurl.com/ybde9dcr

7 DeepMind Technologies Ltd., "AlphaGo Zero: Learning from scratch," (undated); https://deepmind.com/blog/alphago-zero-learning-scratch/, accessed April 4, 2018.

8 Samuel Gibbs, "AlphaZero AI beats champion chess program after teaching itself in four hours," in *The Guardian*, 7 Dec 2017, https://tinyurl.com/y88mhml5, accessed April 4, 2018.

9 Holland, *Emergence*, 214.

of this whole ecosystem of change agents? What difference might it make in the rate at which the entire ecosystem learns about how to increase literacy across North America? As we describe CAS theory below, we will use this conference to explore how these ideas might relate to complex social change.

The process we describe below is not unusual or new. As productive human beings, each of us is exquisitely skilled at learning about complex phenomena *as an individual*. We learn how to use language; how to drive a car; how to raise children; how to navigate the politics of the organizations in which we work. The challenge comes when we try to navigate these complex environments *together*.[10]

As Samuel's checkers-player program suggests, CAS theory focuses on what it takes for a whole system to learn and adapt primarily through the actions of individual agents within the ecosystem, rather than being guided by external design. This is not to say that external input impedes emergence, but that, as our research cases suggested, on an ongoing basis it is too cumbersome and expensive to serve as an ongoing source of direction in a multitude of places, all trying to respond to their unique challenges and opportunities.

As individual agents go about trying to achieve a goal (e.g., learning how to use a language to communicate), they formulate tacit if/then hypotheses: **If** I say "dada," **then** the big person in the room will pick me up. Over time, these hypotheses accumulate into what Holland refers to as an internal model.[11]

This tacit process works well for an individual. But for the larger system, the process needs to become more explicit or overt. Holland asserts that "when the model is tacit, the process of discovering and combining the building blocks usually proceeds on an evolutionary timescale. When the model is overt, the timescale may be **orders of magnitude** shorter."[12]

If Holland's proposal is true, this has huge implications for our conference organizers. How could they help the field of childhood literacy accelerate their progress by "orders of magnitude"?

Using building blocks as an organizing principle

As described above, as we navigate our perpetually novel environments, we begin to recognize and learn how to use patterns to navigate the complexity. Holland gives the example of getting a flat tire while driving on the highway.[13] We deconstruct this entirely new situation into familiar parts—how to slow down and move to the side of the expressway; where to find the spare tire; how to jack up the car (or, in the author's case, how to call AAA).

When the model is tacit, the process of discovering and combining the building blocks usually proceeds on an evolutionary timescale. When the model is overt, the timescale may be orders of magnitude shorter.

Our capacity to discover and then learn how to navigate these patterns is a critical component of learning in complex environments. Using massive computer power, a game-playing computer program can learn through rapid iteration. Living systems cannot iterate so quickly. CAS theory suggests that the more agents in living systems interact—the more they compare notes, so to speak—the faster these patterns become evident and useful; the more innovations they generate. Comparing the experiences of many diverse agents working towards the same goal is akin to running through many iterations of an experiment simultaneously. By comparison of many instances, we can also test to see which patterns are more common across contexts and which might be coincidental and unlikely to recur, so that we avoid learning the wrong lesson from a single success or failure.[14]

10 This is one way to think about the distinction between adaptive and emergent strategy. Adaptive strategy sets the stage for an individual organization to learn and adapt. Emergent strategy sets the stage for a whole ecosystem to learn and adapt.

11 Holland, *Hidden Order*, 57-60.

12 Holland, *Hidden Order*, 37. Emphasis added.

13 Holland, *Hidden Order*, 51.

14 Holland, *Emergence*, 242.

But how do we make this explicit, so that we can learn together? One obvious answer is that our conference organizers could design the conference around some well known and emerging building blocks: what are we learning from all of our experiences about the hypothesized relationship between classroom size and literacy? Between nutrition and literacy? Between the effects of poverty and literacy? What have each of us tried to do to mitigate challenges caused by large classrooms? Poor nutrition? Poverty? What has worked well to increase literacy and what has not? What does this tell us as an ecosystem of change agents about where we are today and what to try next?

Exploring multiple hypotheses

In human affairs, we have a tendency to favor consistency. CAS theory favors *diversity* of thinking. It suggests that consistency, in fact, impedes a system's ability to adapt by slowing down the process of discovering and learning about building blocks:

> The usual view is that the rules amount to a set of facts about the agent's environment. Accordingly, all rules must be kept consistent with one another. If a change is made or a new rule is introduced, it must be checked for consistency with all the other rules.
>
> There is another way to consider the rules. They can be viewed as hypotheses that are undergoing testing and confirmation. From this view, the object is to provide contradictions rather than to avoid them. That is, the rules amount to alternative, competing hypotheses. When one hypothesis fails, competing rules are waiting in the wings to be tried.[15]

CAS theory relies on the experience of a diverse set of individual agents working independently to achieve a goal; exploring their own hypotheses, but interacting with each other as much as possible. "Communication among agents can have a profound effect on the behavior of a complex system. The ability to communicate expands the behavioral repertoire of agents, introducing a variety of new opportunities…Communications can radically alter the performance of a social system; for example, ants leave pheromone trails that allow the colony to self-organize into a coherent mass for more efficient hunting."[16]

Coming back to our conference organizers, in a session on nutrition and literacy, this suggests reducing the space taken on the agenda by expert presentations and giving more opportunities for all of the change agents in the room to share their own experiences, discoveries, and hypotheses related to the connection between nutrition and literacy. "Here's what I see happening in our community; here's what I am trying; and here's the results I am getting. What about you?" In our report, *A Whole Greater than Its Parts,* we refer to this as *returning learning to the system*. The bee comes back to the hive and does a dance to communicate where it found nectar-rich flowers. Comparing the hypotheses and results from a whole range of experiences helps the community of change agents accelerate its ability to develop more powerful and nuanced hypotheses to inform their next set of actions.

Comparing the hypotheses and results from a whole range of experiences helps the community of change agents accelerate its ability to develop more powerful and nuanced hypotheses to inform their next set of actions.

Learning by credit assignment

CAS researchers share a challenge experienced by funders and evaluators: *attribution*. How do we know that the actions we took among all of the many interacting variables contributed in any way to moving the needle on a complex social change? In the example we started with, how does the game-playing program know that any given early move in a game of Go contributed to the program's win or loss?

Samuel's checkers-player program used credit assignment: after each game, each move (hypothesis) in a winning game received a small credit. If the program lost, each move lost a small amount of credit.[17] As the program continued to play games, some hypotheses gained a higher credit score than others. The hypotheses were listed according to their scores. (Holland described these lists as *bulletin boards*.[18])

15 Holland, *Hidden Order*, 53.

16 John Miller and Scott Page, *Complex Adaptive Systems: An Introduction to Computational Models of Social Life* (Princeton University Press, 2007) 242.

17 Holland, *Hidden Order*, 42.

18 M. Mitchell Waldorp, *Complexity: The Emerging Science at the Edge of Order and Chaos* (Simon & Schuster, 1992), 185-193.

A hypothesis with a higher score signified that it was more useful to the system and those hypotheses rose to the top—"a progressive confirmation of hypotheses . . ." as Holland describes it.

Credit assignment relies on being able to recognize success or failure. Importantly, a computer program doesn't get invested in victory; it's not afraid to admit a failure. It also doesn't get invested in proving its favorite hypothesis. In our human systems, all of us have observed how easy it is to glorify a success and ignore a failure, especially if it threatens a hypothesis that we have invested in using and promoting. CAS theory sheds a new light on the implications of this human tendency, which leads to over-crediting the hypotheses that may have contributed to a success and to not challenging the hypotheses that may have contributed to a failure. When this happens, the credit assignment 'bulletin board' can't do its job. *It slows our ability to learn and adapt.*

For our conference organizers, this reinforces the importance of talking about failures as well as successes, as honestly as possible, and about what participants think may have contributed to their results, for better or worse. Panel discussions that compare failures and successes among a series of literacy initiatives could give members of the audience an opportunity to develop their own more robust takeaways than presentations about a single success story.

Ultimately, the goal of the conference would be to accelerate the process of exploring the most important building blocks and discovering the most promising hypotheses, based on honest data from the field. But then what? What does CAS theory suggest about what needs to happen over the course of the year between conferences to accelerate learning across this ecosystem of diverse agents?

How does a complex system learn when the game never ends?

The essence of the learning process Holland describes is to continually improve the ability of a system of agents to recognize patterns and make predictions, based on hypotheses, about which move to make in a particular situation to move closer to a goal—absent an external designer or strategist. "Prediction makes improvement possible, even when there is no referee to distinguish 'right' from 'wrong.'"[19]

Computer games are an opportunity for CAS researchers to demonstrate what's possible. But let's summarize the differences between computer games and other kinds of complex systems and draw implications for our conference organizers and this ecosystem of change agents working on literacy.

Credit assignment is fairly straightforward in a game of checkers, chess, or Go. The game comes to an end when there is an obvious winner or loser. In many of the systems CAS researchers study, the goal is more complicated and long-term, primarily because the 'game' never ends. The goal of an immune system is to protect the identity and health of the living system. The goal of an economy is to allow agents in a community to use the diverse resources of the community to support a good life (however the community defines that).

Complex systems can't wait until the game is 'over' to progressively confirm their hypotheses.

Complex systems can't wait until the game is 'over' to progressively confirm their hypotheses. It would be akin to a sports team waiting until the end of a season to review its game films to figure out why it didn't win the championship. Complex systems need to be able to make predictions based on the available data that is "good enough." As Holland describes it, "when we face complex situations, our objective is almost always to 'do it better.'"[20]

In complex human systems, the connection between cause and effect is often distant in time and space, and data about results may be inconclusive or difficult to obtain. The goals themselves might not be clear or shared. On the flip side, human beings can do more with available data than simply assign a hypothesis a small credit. We can ask follow up questions and look for confirming or contradictory examples to test the meaning we make from our results. We can challenge our measures of success and reflect on and adjust the learning process itself. The more honest we can be about recognizing success or failure as we go about making meaning in the context of our ongoing work, the faster we will learn and adapt.

19 Holland, *Emergence*, 76.

20 Holland, *Emergence*, 216.

In this hugely complex and messy environment of social change, even without crystal clear data about results and what contributes to them, we still have to make decisions on a daily basis. Input from research and evaluation can get us started, but CAS theory suggests that anything we can do to make the real-time learning process explicit will help to accelerate our adaptation.

Coming back to our ecosystem of change agents working to improve childhood literacy, after the conference is done and everyone goes home is when the real work of this learning ecosystem gets started. Cause-effect relationships validated by research and evaluation can be a good starting point, but because of the complex and interacting variables, our community of change agents needs to experiment with how what they heard in the conference could be used to improve literacy in their own local environments.

CAS theory predicts that the more explicit these experiments are—naming and testing hypotheses about building blocks and being honest about assessing the results—the more they will accelerate learning and results in these local environments. Inasmuch as it is within their purview, our conference organizers can help the larger ecosystem learn and adapt by helping to return the learning to the system—perhaps by spreading news of the results of these experiments along with new research and evaluation data across the ecosystem; or holding webinars around the role of nutrition or reading coaches, or small regional gatherings to learn what new building blocks are being discovered related to low-income communities and literacy, what new hypotheses are being tested, and what results change agents are achieving; or organizing communities of practice around access to books or dealing with family stress.

Imagine how this might affect the conversations that happen at *next year's* conference. It might be akin to the first conference of meteorologists the year after the jet stream was discovered. The room would be a buzz with newly discovered patterns; new questions being asked; more nuanced hypotheses being put forward.

This article has just skimmed the surface of what we can learn from complex adaptive systems theory. The seven cases the 4QP research team studied in *A Whole Greater than Its Parts* offer ideas about what promotes and what impedes adaptation in the complex systems these initiatives were targeting. The research team invites readers to continue to submit examples of emergence in complex social change here. What we said at the end of the report bears repeating: There is much more to learn . . . always.

There is much more to learn . . . always.