Interaction, Innovation, and Immunity: Enabling Agents to Play

By Luis O. Arata, Chairman
Department of Fine Arts, Languages & Philosophy
Quinnipiac University, Hamden, CT 06518
luis.arata@quinnipiac.edu

Abstract

In a dynamic environment, an autonomous system has to interact with unknowns in order to remain viable. This interaction exposes the system to problems that need solutions not derivable entirely from the system itself. The autonomous system then has to take somewhat of a leap into the unknown without breaking down. After reviewing basic concepts of play and self, I will model this situation as a process of self-repair in confronting unpredictable and previously unknowable situations. The capacity of an autonomous system to innovate appears then as an immune mechanism for interacting in a dynamic environment.

Nature is a rich source of models for constructing artificial agents, from robots to software agents. In this paper I will focus on one neglected natural feature which can enhance the autonomy of agents: that is the capacity of living beings to play. Play is a basic component that helps organisms prosper in dynamic environments. They are able to wiggle about rather safely and in the process they can develop ways to deal with changing situations that could otherwise have fatal consequences. It is as if play were a crucial part of the immune system of an organism. Without play, life may not have the capacity to evolve and adapt to new situations. It is important then to understand how to model effectively this activity. The capacity to play should help constructed agents to operate in changing, uncertain environments.

Imagine an agent facing a new situation that extends beyond whatever choices it can make with clear outcomes. In other words, the agent faces uncertainty. How would it proceed on its own? One simple solution is to make choices randomly. But each choice will have to be tested to see how it works and this could become a staggering task for even simple situations. Chances are that many outcomes will be fatal. How could an agent reduce the number of tests and survive to continue functioning autonomously in the midst of uncertainty?

In this paper I propose that to enhance its autonomy in an uncertain environment, an agent has to follow nature's example and be able to interact playfully. This in turn allows the agent to learn and modify future interactions. Learning through play augments the agent's capability to develop innovative solutions to problematic situations. In this way, the agent increases its immunity to damages from interactions with the environment.

This raises a final issue: learning is not a neutral, passive activity. It is selective and focused. Otherwise it quickly yields an overwhelming amount of raw information. To deal with this problem, an agent has to be able to select and process only what it wants. This sense of desire depends on the nature of the agent. It hinges on what we could picture as its sense of self. The issue becomes then how to design a self-aware agent, so to speak? Would this playful design strengthen a constructed agent's autonomy in uncertain environments? Can play form a loop of behaviors that helps the agent innovate, boost its immunity, and go on interacting?

To examine this question I will discuss the following topics:

- What is play?
- Where does an agent's sense of self come from?
- Why does play enhance an agent's autonomy in an uncertain environment?
- How do we build play into an agent's behavior?

Play

The nature of play appears rather obvious. The literature on play is extensive (Huizinga 1955 and Bruner 1976). Yet play tends to be underestimated. It does not help that play, by its very nature, escapes confining definitions. Clearly play is not just the opposite of work. It is not simply a frivolous activity without consequences. It cannot be reduced to games. It is important to distinguish between play as an activity, and games as structures with rules that promote certain types of focused playing. Most play happens beyond rules (Arata IEEE 2003).

Perhaps the simplest image of play is jiggling. A lose structure leaves room for play. Games provide structures to jiggle out different strategies and modes of playing. When it comes to child's play, Jean Piaget developed a very comprehensive definition that hinges on play's adaptive function. Piaget described adaptive behavior as a combination of assimilation and accommodation to the world. In their most extreme forms, accommodation is imitation and assimilation is play. In

imitation, our systems change to adjust somehow to what is perceived as external models. In play, the perceived world is dismembered and absorbed selectively to fit the systems we already have, without changing them.

Piaget developed his view of play based on early childhood studies. It needs expanding into adulthood if it is to be of help with autonomy. For adults, imitation could also be playful when it involves modeling. A model is a draft that imitates a phenomenon. If several models of the phenomenon are possible, then there is room to play when all models are considered at once. To select one model and exclude the rest does away with play. In other words, accommodation can also be playful when tried in many possible ways without settling down for just one view. As for assimilation, it can very well change gradually the systems that produce it and become a form of accommodation. Play can change the player. So, in this extension of Piaget's conception of play, we see that it does have a wide range of modes beyond pure assimilation. What varies is the level at which playing is carried out. This extension should work for all ages and be useful for developing playful agent behavior.

Lev Vygotsky touched tangentially on the cognitive function of play by suggesting where it can happen. Play and learning can happen only in a restricted zone. This is the zone of proximal development. He defined it as "the distance between the actual developmental level as determined by independent problem solving and the level of potential development as determined through problem solving under adult guidance or in collaboration with more capable peers" (Vygotsky 1978, p. 86). Cognitive play happens in a zone accessible to the player, yet at the same time the access is not causal or systematic. One can make progress in that zone with the help of instructors or other aids.

Independently of Vygotsky, Stuart Kauffman developed the more concise notion of the "adjacent possible" as the boundary, rather than the zone, beyond which one can innovate. This simplifies and focuses things more. The boundary between the actual and the possible is the most important thing. The outer boundaries of Vygotsky's zone can hardly be established except in the conservative sense of being the adult's outer edge of development and which the child would have to reach through the aid of education. But this misses the fact that children can at times innovate beyond adult capabilities, as with language for instance. Kauffman also adds to the model the crucial detail of drift as the mechanism of innovation, following the evolutionary sense of drift.

What I argue, however, is that at least for living beings, this drift is not entirely random. I suggest that play is its source. The paradox of play is that it appears frivolous yet it is a key element of innovation. In order to adapt to a changing environment, and if those changes are new, one has to be able to play to come up with adaptive behaviors that go beyond previous behaviors. In adaptive situations, play functions then as a mechanism of innovation that allows the player to find new balances. The player enters Vygotsky's zone of proximal development. It crosses Kauffman's boundary of the adjacent possible as a way to readjust to changes.

Therefore, one reason why we might want a system to play is because we give it the basis for developing some level of immunity in a changing environment. This immunity comes from the system's ability to innovate for self-repair, which hinges on the system's capacity to play.

This view of play applies to living beings. In other to extend it to constructed agents, we need to first model another component essential to play. That is the player itself. A human player acts through a sense of self that sustains the playful activities. Play does not happen automatically or randomly because it follows the inclinations, capacities, emotions, and desires of the player. And all these tendencies emerge from a sense of self. But what would constitute an agent's self? Before moving on to modeling play, we need to briefly examine what might sustain the self.

A Sense of Self

In *The Embodied Mind*, Francisco Varela proposes that there is no experience of the self as an entity. There is not even such thing as a fixed self. In *I of the Vortex*, neuroscientist Rodolfo Llinás continues the work of Varela, especially along the view that the self is enaction rather than representation. Llinás develops a fascinating model of the self, based on interactive feedback and feedforward loops. He begins with a view of the brain as a system that does isomorphic sensory-motor transformations of the outside world. This creates representations that help the body act in the outside world.

Llinás then follows the lessons of the sea squirt. This tiny sea creature has mobile state followed by a plant-like one. During the first phase, it has a brain. But when it finally attaches itself to a surface, the sea squirt digests its own brain along with the tail that provided motility. Llinás concludes that "the evolutionary development of a nervous system is an exclusive property of actively moving creatures" (Llinás 2001, p.17). The nervous system and particularly the brain, are predictive instruments that allow the organism to move more safely in search of food, often in a potentially hostile environment. The brain creates working models of the environment and interacts through those models with the world to give the body navigational capabilities. Llinás imagines that such models are very much dreams of our brain, and in the waking state those dreams are guided and shaped by the senses: "the fact is that we are basically dreaming machines that construct virtual models of the

real world" (Llinás 2001, p.94). In effect, what we perceive is a world that is both physical and virtual.

The sense of self emerges from interactions in the brain as it coordinates actions. The self could function as an avatar of sorts within the brain's representation of the world. In actual dreams, when the senses are dormant, the self moves through a recreated world made of collages of memories patched through internal logic. When awake, our actions double the displacements of the avatar in the brain's representation of the physical environment as informed by the senses.

Sensations, including the elusive self-awareness, are what reflexive loops feel like in order to help us navigate in the world. Pain and pleasure are guiding sensations. Self-awareness is perhaps the most complex manifestation of this cybernetic system. Llinás speculates that our sense of self and what could be called "intelligence" may well be an emergent property of how our brain wired itself as a navigational tool. He concludes with the realization that there are many possible architectures for cognition and sensations. Ours does not seem to be the only one.

Gerald Edelman relates the sense of self or consciousness to mental maps and reflexive loops as well. Much like Varela and especially Llinás, he views the self as a process that emerges from the coordination of cognitive aggregates in order to help us be in the world. The key mechanism that binds all our cognitive mechanisms is what he calls "reentry:" a signaling along reciprocal connections. This reflexive neural interaction works within the complex topologies of our brain to create the sense of self out of weaving memories. Such dynamic self probes as well as reacts, and is capable of learning from the process.

Edelman stresses, however, that memory is creative rather than replicative: "every act of memory is, to some degree, an act of imagination" (Edelman et al., p.101). Memory for Edelman is a process that works in the present and is also subject to uncertainties. Linking memory with the imagination opens it to changes and makes it much like dreams. Edelman's sense of self becomes quite similar to what Llinás proposed: that we are dreams guided by the senses.

Edelman agrees with Varela that the self emerges through enaction, but he thinks that we cannot avoid representation entirely. Although not mirrors of nature, our mental maps always contain some degree of representation. Our sense of self rises out of the interactions of differences among these maps, so to speak. Edelman concludes that our selves may transcend our analytic limits of representation by synthetic means that create new viable grammars of the imagination.

Taking a critical leap, biologist Lynn Margulis proposes that mind is a self-referring process present even in life forms as elementary as bacteria. This view may stretch the definition of mind, but it maintains the central feature of looping interactions, or reentry, as the most basic mechanism of any conception of mind.

All architectures of the self hinge on the function of reflexive interactions between the creature and the environment. The nervous system manages that interactivity to help the body navigate effectively in an evolutionary sense. Consciousness is an emergent property of such a system when it centralizes its control mechanisms to work more effectively in highly complex situations.

Like Margulis, Llinás cannot resist asking also if something nonbiological can have a self or a mind. This is the basic question of artificial life. His investigations from the biological side suggest an affirmative answer. And a key component of any architecture of mindful, self-aware systems is looping interaction.

For us humans, it is the reflexive interaction of memories, dreams, and the imagination, that help us create living maps tightly linked to sensations. Our minds navigate virtual models of the world to make complex choices that guide the body. Those maps come in a wealth of shapes and colors, from the geographical to the emotive (Arata M/C 2002). The arts and the sciences from all cultures contribute to the creation of our models of the world and of our selves. The question is how can constructed systems incorporate elements of this sense of self as the core for the development of agents with the capacity to play?

Toward Agents Capable of Playing

Traditionally, computer simulated play, when not part of games, tends to depend on the use of random generators. The agent randomly generates new behaviors and selects the most viable one. But this process quickly becomes too long, costly, or risky beyond the simplest situations. And besides, this procedure does not live up to one of the most important aspects of play: that it is not all that random at all. The player introduces preferences and styles, and it is within such narrow focus that some randomness comes into play.

So we need to see first how a constructed system may be given freedoms to act autonomously. Play and freedom are inseparable. And secondly, we have to explore how the system may chose from those freedoms to play within the context at hand and with its own sense of self. The choices made are based on tasks, goals, reasons, emotions, and assorted internal impulses. My task now is to develop the foundations of a rudimentary model of a constructed agent able to play.

The model begins with an ELF configuration. This is the elementary loop of functioning (ELF) architecture that Alex Meystel and James Albus presented in *Intelligent Systems*. The ELF architecture has a world model component to process incoming information from

sensors. From that information it generates appropriate behaviors based on its world model. An initial model boots the agent. From there on, the agent learns on its own and evolves under the regulation of a value judgment module that contains the agent's basic rules of operation. Sensors and behavior generators link the ELF to the world. In a self-contained ELF, the connections to the world lead instead to a simulating system. In an ELF with learning, the value judgment module is more than a static system of rules. It can evolve its rules through learning and directly affect the ELF world model. We can picture the learning module as an ensemble of neural nets capable of learning. And this learning is looped with the ELF's initial world models and value judgment rules that booted the system in the first place.

Imagine now that an agent moves in an uncertain environment and suddenly gets stuck for one reason or another. The agent does not have the capability to examine the problem and solve it from what it already knows. So what could it do without outside help?

The constructed agent must have the freedom to modify its world model or its value judgment module when it gets stuck. One route, perhaps the simplest, is to alter one of its basic judgment rules. A possible course of action would be:

1. When stuck the agent selects randomly a previously untested rule. For simplicity, the set of untested rules contains the opposite or complements of existing rules.
2. The selected rule is changed.
3. The agent slightly tests the behavior based on the changed rule.
4. If it gets stuck again, the agent returns to its original position, restores the rule and goes back to 1. Otherwise it proceeds.
5. The agent remembers the changed rule and the circumstances that required change. Then it restores its initial set of rules but saves what it just learned as new knowledge.
6. If it gets stuck again, the agent tries its new knowledge.
7. If new knowledge does not help, then the agent goes back to 1.

A more complex procedure involves a non-random way of selecting one or more rules to change, and tweak a rule rather than simply revert to its opposite or complement. For example, the agent could review how it got stuck and establish a fuzzy hierarchy of rules used at the moment. Then it can tinker with the rules following the hierarchy until something works. This enriches the way the agent is able to play with rules to innovate when faced with a problem that it has to solve on its own. An even more complex level is to combine tweaked rules and test to see what happens.

A different route, more natural but more difficult to implement, is for the agent to use errors to tinker with its world model and rules module. When the agent gets stuck, the agent looks for ways to restore its behavior. It does this by playing with its rules and models through trial and error until something works. When a solution emerges, the agent remodels its internal rules and world model. It also keeps a history that allows it to revert to previous configurations, if necessary. This process of adaptation to a changing environment is a remodeling. Previous models are extended and modified to account for new experiences. In this mode, the course of action could be:

1. When stuck the agent uses its short-term memory to review how it got stuck, and selects the rules that go it to that state.
2. From here on the agent follows step 2 of the previous procedures. The difficulty is that now we may be dealing with an ensemble of rules to tinker with. This may require testing and tentatively modifying each rule at a time as well as all combinations of rules until an optimal solution is found. But the play style of the agent may shorten the procedure by focusing otherwise random searches.

Let's imagine now that the error actually hurt the agent. If it is a virtual agent, it would simply revert to its saved initial configuration after it got unstuck. But now the agent is enabled to learn from its mistakes. To do that it explores on its own time how to self-repair to reach its initial state, as well as how to avoid similar errors in the future. In other words, the agent reflects after it restored itself. Again, the agent may play to simplify the task based on its own styles and preferences. The interesting thing here is that the agent plays internally while it continues performing its other tasks. Self-repair becomes a source of internal innovation since the agent learns new procedures that are meaningful for itself.

Physical agents would need to have a toolbox, a set of spare parts, and perhaps even repair agents. Repairs for physical agents is not as simple as reverting to the initial configuration in a virtual agent. A physical agent will have to get there using what it has at hand, and play harder to develop solutions if they require innovation.

All these procedures make the virtual or physical agent an innovator at least with respect to itself. The agent essentially discovers new rules or remodels methods on its own as it navigates uncertain environments.

Conclusion

An autonomous agent lacks omniscience. It does not have an external narrator that can forewarn the agent about changes and problems. The agent is like a first person narrative. It needs freedom to move about, bump into things, self-repair, and learn from experience when it gets stuck. The agent could evolve its internal models and this evolution would be adaptive and depend most likely

on neural nets (Fogel 1999). These neural nets would be part of the ELF and they could even be shared much like a brain, introducing yet another variation to robot architecture along the looping concept of re-entry that Edelman has suggested is basic to brain activity. It would also have the basis to develop a concept of self to launch explorations, as Llinás has shown. This is the basis of an interactive intelligence (Arata NIST 2002).

The agent would then evolve a repertoire of behaviors for different occasions it learns as it acts, as long as it has a judgment system capable of selecting appropriate fits for a current environment. This is a higher order play that would enable the agent to carry out exploration and evolutionary play.

But this architecture could not be built from pre-existing blueprints all at once. It will have to evolve through playing and learning. The agent will need an extended childhood. It will grow up gradually as it goes out to play in increasingly more complex ways and learns from it.

In more sophisticated agents, undoubtedly the judgment module and the world model will be integrated into one that can be imagined as a self. Just as in natural systems, there may not be a set of clear rules to guide behaviors. There will be instead far more complex configurations that the agent inherits and then evolves by learning. There will be rookies and experienced agents. The agent could even become self-aware as it learns more about ways to self-repair.

Agents may also network. When faced with problems, the agents can communicate. If one has already innovated when faced with a new situation, this can be transmitted to others that encounter a similar problem. We could tweak a suitable ensemble of agents to network among themselves. This could result in higher level agents. The configurations can produce a rich nesting of levels as Meystel and Albus suggested with the notion of multiresolution. It is perhaps the uncertainties and surprises of this “mise en abîme” that begin to mark the emergence of a profound sense of autonomy: one that is linked to non-trivial knowledge. And such knowledge can emerge through the interplay of levels of different networks of agents, by means of loops that Edelman saw as the foundations of the sense of self. The simplest architecture then is a multilevel configuration of ELFs linked through loops that enrich the interactions with small world properties (Watts 1999). Most likely, networking agents will yield rich emergent behaviors.

Sophisticated agents will be able to play according to their self-preferences. If we represent some of these preferences as emotions, then we can construct agents with personality. The development of emotional machines is already being explored (Norman 2003). The idea is to enable constructed agents to have emotions. Fear, for instance, would serve to limit options that feel risky or dangerous. Other emotions might help simplify searches for solutions. For example, a sense of beauty has proven fruitful in science. Solutions with more symmetries are often favored. But the problem with trying to design emotions is their intrinsically human nature. We do not grasp well enough this subjective feeling. When it comes to constructed agents, they would be derivative, more like icing on the cake. Emotions are tagged to the more fundamental sense of self that we need to model first. Emotions are subjective responses that help us make choices. Once the agent becomes self-aware, then emotions could become very helpful shortcuts to streamline response behaviors.

I finish with what I consider the sorcerer's apprentice paradox. Play hinges on having freedom of choice. But do we want systems with such freedom? Why would we want to try to construct systems that can play? It is more likely to get into trouble, disobey, and make mistakes. Play introduces imperfections into designs. It resists optimization. It would make the agent harder to control. It could lead to a sorcerer's apprentice paradox as the agent goes out of control and needs to be restrained.

Play essentially entails a tradeoff. It reduces control and certainty, yet facilitates evolution and innovation (Arata MIT 2003). In rather closed, mostly predictable systems, play might not be necessary to carry out tasks. In that ideal situation we could benefit from more goal-centered behavior and less play. Yet in certain controlled environments, agent play may be quite valuable. In the area of game-like simulations where we tweak parameters to see what develops, an agent able to play would be able to explore all sorts of possibilities on its own and analyze the results for us. Agents could even evolve more complex game-simulations on their own. We would then play with the agents to facilitate the development of new situations for problems at hand. Agents capable of playing would be most valuable helpers in virtual modeling and testing.

But it is in uncertain environments, especially in natural environments, that autonomy gains in value, at least because it enhances survivability, and such type of autonomy can’t do without play.

The issue will then be one of agent-maker interaction. Will we be able to control our own creations so that they continue to serve and obey us? Or will giving agents the freedom to play lead, for better or for worse, to something we can hardly imagine? If agents develop in such direction, perhaps their evolutionary play will also yield a form of wisdom that emerges from the close interaction of agents and makers. In other words, the agents might be able to develop an ethics of behavior based on certain principles that constitute their acquired wisdom. Perhaps we could even learn something fundamental from our creations.

References

Arata, Luis O. 2002a. Creation by Looping Interactions. *M/C: A Journal of Media and Culture* 5.4. http://www.media-culture.org.au/mc/0208/creation.html

_____. 2002b. Modeling Interactive Intelligences. In *Measuring the Performance and Intelligence of Systems*. NIST Special Publication 990.

_____. 2003a. Can Your Autonomous Robot Come Out and Play? In *Integration of Knowledge Intensive Multi-Agent Systems. KIMAS 03: Modeling, Exploration, and Engineering*. IEEE Press.

_____. 2003b. Reflections on Interactivity. In *Rethinking Media Change: The Aesthetics of Transition*. David Thorburn and Jenkins H. eds. Boston: MIT. http://media-in-transition.mit.edu/articles/index_arata.html

Bruner, J.S.; Jolly A. and Sylva K. eds. 1976. *Play: Its Role in Development and Evolution*. New York: Penguin.

Edelman, Gerald and Tononi G. 2000. *A Universe of Consciousness: How Matter Becomes Imagination*. New York: Basic Books.

Fogel, Lawrence J. 1999. *Intelligence Through Simulated Evolution*. New York: Wiley.

Huizinga, J. 1955. *Homo Ludens: A Study of the Play Element in Culture*. Boston: Beacon Press.

Kauffman, S. 2000. *Investigations*. New York: Oxford.

Llinás, Rodolfo. 2001. *I of the Vortex: From Neurons to Self.* Cambridge: MIT Press.

Margulis, Lynn and Sagan D. 1995. *What is Life?* New York: Simon & Schuster.

Meystel, A. and Albus J. 2002 *Intelligent Systems*. New York: Wiley.

Norman, D. 2003. *Emotional Design*. New York: Basic Books.

Piaget, Jean, 1962. *Play, Dreams, and Imitation in Childhood*. New York: Norton.

Varela, Francisco; Thompson E.; and Rosch E. 1991. *The Embodied Mind*. Cambridge: MIT Press.

Vygotsky, Lev S. 1978. *Mind in Society: The Development of Higher Psychological Processes*. Cambridge: Harvard University Press.

Watts, Duncan J. 1999. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton: Princeton University Press.