

Memorandum  
August 4, 1992

To: Ed Knapp  
From: Joshua Epstein

Re: Artificial Social Life and 2050

As you recall, during the Integrative Themes Meeting, July 8-15, I proposed that, as part of our theoretical work on sustainability, we attempt to apply the techniques of artificial life to the question of social evolution as a whole. Is there a small set of local rules, that is, rules governing the behavior of individual agents, that over many iterations, will generate a crude caricature of, say, the observed international system--a set of coherent societies with internal structures and dynamics, interacting with one another in various cooperative and competitive ways, in an environment that is affected by, and feeds back on, the productive activities of the agents?

The idea of approaching social evolution from the A-Life perspective certainly did not originate with me; it is mentioned by Chris Langton in the first A-Life volume and may well be under active development already. First and foremost, we should find out what, if anything, is being done along these lines specifically, and we should survey the whole A-Life field for work that could apply to our project. (This survey might be a good fit for Gottfried.) At the Integrative Themes meeting, I broached the general idea to Chris Langton who invited me to follow up, which I will do. Obviously, Chris would be an invaluable resource on all of this, and if possible, he should attend whatever part of our September Santa Fe meeting is devoted to this.

In presenting the initiative to others--and in thinking about it ourselves--it is crucial to distinguish the idea sharply from other work in socio-political systems modeling. Chris Langton, in trying to distill the "essence of A-Life," stressed:

- bottom-up rather than top-down modeling
- local rather than global specification
- simple rather than complex (local) rules
- emergent rather than pre-specified collective behavior.

Now, there are quite a few computer models of international relations that begin by positing states possessing decision-making elites and in which high-level decision rules are pre-specified, rules like:

- Ally to balance power against any potential hegemon;
- Allocate resources to domestic social programs if revolution is immanent; and
- Trade if it is profitable.

One excellent example of this type of top-down simulation model is Cusack and Stoll's 1990 EARTH (Exploring Alternative Realpolitik Theories) model. In a nutshell, a system of states--initially represented as a grid of hexagons--is assumed. The user pre-specifies the states' rules of behavior regarding war initiation, alliance formation, and the allocation of resources between domestic social needs and external military action. Empires can form (hexagons eat neighbors and grow into big irregular polygons) but can fractionate under rebellion if war costs prohibit the satisfaction of domestic needs. It is a very interesting model and we may ultimately wish to consider doing something of the sort. Initially, however, this type of top-down modeling, with pre-specified states and static (non-evolving) rules, is not what we should be about. Rather, the social organizations themselves (the "states") should emerge; any elites, divisions of labor, or classes should emerge; and so--as guided by selection pressures--should any rules or high-level regularities of behavior, like when emergent organizations compete, when they ally, when they make war, and when they trade. We want to grow all of this from the bottom up, if possible. Or, more humbly, we want to explore the possibility of doing so.

Although we will need to continue discussing the goals of the exercise ("flight simulation," humility injector, etc.), one motivation of "artificial social life" is simply scientific: to discover fundamental local, or micro, mechanisms underlying macroscopic structures and behaviors of enduring interest. Core questions might include:

- What system(s) of local rules will generate politically egalitarian societies? Totalitarian societies?
- What rule system(s) will yield social aggregates that are largely peaceful and cooperative? Warlike and competitive?
- Can internal instability--revolutions, the rise and fall of empires--be made to emerge from simple rules?
- Will ideological patches of special resilience (quasi-religions) emerge?
- Are there cyclicalities, power laws, or other regularities that seem generic in any sense?
- Is social foresight--concerning, say, environmental collapse--an emergent capacity? If not, what sort of "genetic engineering" (operation on basic rules) will make it so?

Clearly, patterns of economic growth, income distribution, and technological innovation (and diffusion) should also emerge from a respectable toy model.

Of course, most of the above language is terribly vague. We will need to think hard about the most basic definitions: What are individual agents? How do they reproduce and what is transmitted when they do? Why do they aggregate? At what point can one say that a social group (a boundary, an inside, and an outside) has emerged? To how many groups can an agent belong? In what ways can groups interact (a) with each other and (b) with an idealized environment? How to represent the environment? Some of these questions are being studied at very fundamental levels (e.g., Fontana and Buss). How do we use this work? In particular, how simple can we keep things and still get out some core behavior:

- aggregation into groups with internal structure;
- competition/cooperation within groups;
- competition/cooperation between groups;
- interaction with an environment.

I have a number of very specific thoughts on some mechanisms. But, that, generally, is a topic for our "group-grope" in September. Indeed, just to get the ball rolling, I have formulated a few questions that may focus discussion.

#### Some Candidate Questions To Focus Discussion at September Meeting

- (1) What are some target patterns and behaviors (Turing test analogues) an artificial social life system should be able to produce? For example, social hierarchy, territorial conflict?
- (2) What are individual agents and what are some candidate local rules? For example, John Holland's Echo might be exemplary. Relatedly,
- (3) What is being done in A-Life specifically that is applicable?
- (4) What is being done by people in the extended family of the sustainability project (Marc Feldman, Per Bak, ...) that could be fashioned into "modules" of this work?
- (5) What are the goals of this work in the 2050 context?

We need to think hard about whose ongoing work can fit in where. Much, I'm sure, has actually been done. Let me conclude with a few scattered thoughts.

## Scattered Thoughts on Some Specific Areas

Attitudinal Evolution. Suppose agents are n-tuples, with each slot a locus, and imagine that at each locus, a number of values--alleles--are possible. One locus might control "attitude toward the environment," with the two alleles: "don't give a hoot" (=A) and "total whale-hugger" (=a). Using methods fully developed by Marc Feldman, one can examine what mutation rate (from A to a) would be required to produce an overwhelmingly whale-hugging society in N generations.

Innovation. At the Integrative Themes meeting, Brian Arthur discussed tree-like structures of economic niches, all flowing from some basic source--a core industry, say. And he argued that technological innovations pose fundamental threats to these structures--as when the advent of railroads swept away the entire "tree" based on horses, stagecoaches, and so forth. I noted that if, in fact, innovation is the major threat to these trees (vested interests) then the survival of the tree requires the prevention of certain innovations, and, of course, we see plenty of this behavior (e.g., oil cartels opposing nuclear power, etc.). Stu Kauffman and I discussed the phenomenon at some length and the whole idea of boundaries, protective membranes, "emergent socio-economic immune systems" seems quite important to me.

Revolution. As you know, I have been using epidemic models to examine revolutions. And there must be some mechanism for generating internal frustration in these artificial societies. One process is where culturally transmitted beliefs (expectations) about society come into gross conflict with social realities. For example, one might imagine a society, call it X, in which children are routinely taught that "X is a classless society." This is the transmitted "theory." If the "data" is that an observable power elite has obvious special privileges, frustration builds up--legitimacy is undermined and the potential for instability results. Another frustration mechanism could be the observation (through media or migration processes) of neighboring societies in which higher standards of living prevail.

Conflict. As for dynamics between societies, trade must of course be possible. The co-evolution of military establishments (arms racing), the formation of alliances, and war are obvious behaviors that a respectable toy should be able to produce. I have spent much of my time studying military competition and my own Adaptive Dynamic Model of combat could be used to simulate mutual attrition and territorial conquest in war. Simpler models might suit our purposes just as well, probably better. Naive "paths to conflict" might include: population growth → growth in demand for resources → territorial expansion → conflict over turf. Obviously, we want these to emerge, not to be built in.

Cooperation. Bob Axelrod's thoughts on the emergence of cooperative behavior will certainly be invaluable in a number of contexts.

Aggregation. The aggregation of agents into groups, and of groups into bigger groups, such as alliances, is obviously an essential process. There are many candidate algorithms, including competitive neural nets (e.g., Kohonen self-organizing feature map). Another tack is to map every configuration of agents to an energy, define a Lyapunov function, and allow relaxation into groups. One fascinating example of this is Bob Axelrod's spin glass Landscape Theory of Aggregation.

#### A Preliminary Michigan Meeting in September

I have had two very useful discussions of the general idea with Bob Axelrod, who, I am delighted to report, is very interested, and has drafted some thoughtful memos of his own. When he returns from a vacation (around August 15) we will resume our discussions. He thought, and I agree, that it makes sense to have a small, informal one-day meeting in Michigan to generate some kind of straw man to focus the Santa Fe discussion. It would be great if you could make it. I will be in touch about possible dates. Bob also proposed that I consider making some sort of presentation on this effort at the November Michigan Outpost Meeting on Emergent Organizations, a thought which I will take up with John Holland.

Finally, I would add only that I had a long talk with John Steinbruner on the idea. He was very interested and had a number of extremely useful thoughts of his own. I'm feeling "up" about the whole thing.

cc: Bob Axelrod  
Rob Axtell  
Murray Gell-Mann  
John Holland  
John Steinbruner