Broken Models: Issues in Social Simulations

Jacob M. Peck
jpeck2@oswego.edu
Department of Cognitive Science
College of Liberal Arts and Sciences
State University of New York College at Oswego

April 23, 2012

Abstract

An overview of a few salient issues in the field of social simulations and modeling is presented. A brief introduction to social modeling is provided, followed by a tour of issues the author deems important in the coming years of social simulations. Extra attention is paid to the standardization of modes of communication and modes of information in such models, as well as on the philosophical aspects of abstraction, and a projection into the future, revolving around the ethics of such modeling. Some closing words reiterate and reinforce the message of the body.

1 What is modeling/simulation?

Social modeling or social simulation is a relatively new branch of the social sciences utilizing the computational power available to researchers today to simulate social situations. Such models often attempt to depict the social lives of humans (though any other social being would suffice) in a computational context, through a layer of abstraction.

Social modeling is gaining traction among select groups of social scientists, though there are several of the old guard who are questioning the validity of these models, and what their place is in the larger scientific research community.

As with any new branch of the sciences, issues arise which demand examination. This short paper will address the issues that I feel are most salient in the context of the growing use of social models.

2 Issues in social modeling

Several issues are present in social modeling, and while all of them are worth discussing, a few rise to the top as the "most pressing" issues in today's research environment. What follows is an overview of those issues I feel *need* to be addressed for the science of social modeling to be taken seriously in the coming years.

2.1 Refinement of behavior

The first major issue I see with social modeling is how to define and refine the behavior of the agents. How to clearly model an agent's behavior to reflect how it would interact with the real world is a tricky question that doesn't have a trivial, one-size-fits-all answer. There is no clear algorithm for deciding which aspects of an agent's behavior to model or how to model them. In many cases, the researcher has simply sit and think about how the model can appropriate the real world, due to computational limitations and lack of complete information. In doing this, agent behavior is defined in a sort of subjective manner, by going with whatever the researcher thinks would work best.

Oftentimes, modeling an agent's behavior comes down to modeling communication and cooperation between agents. To this respect, there are a few classes of communication and cooperation that I would like to propose.

2.1.1 Modes of communication in social simulations

Communication in a social simulation can take one of several modes on a per-agent or potentially even per-relationship basis. The modes I would propose are as follows:

Zero information Agents do not communicate with one another at all. A stricter zero information simulation might even have the constraint that no agent knows about any others.

Visible information Agents are privy to the information they can glean from visually sizing up the other agents in the simulation. This includes such data as location, gender differences, racial differences, etc.

Partial information Agents are allowed to communicate with one another to project information to other agents. This includes such data as favorite color, etc.

Global partial information This is the same as partial information, however whenever an agent communicates such information, it is globally broadcast to *every* other agent in the world, rather than just the one it is communicating with.

Complete information This is the most privileged class of information in a social simulation. An agent with complete information about another agent or another object in the world has complete and total information about that agent. In the case of a computer simulation, such an agent would theoretically be able to produce clones of the object they have complete information about, by virtue of the definition of "complete" information.

Global complete information This is the same as complete information, except that an agent with complete information about another object broadcasts this information to the entire world.

2.1.2 Modes of cooperation in social simulations

Cooperation, like communication, can take one of several modes in a social simulation. These can be on a per-agent, or per-relationship basis as well. The modes I would propose are as follows:

Altruistic (or pure) cooperation This mode of cooperation is one in which agents receive no benefits from cooperating with the other agents in the simulation, but rather the world as a whole is the only benefactor.

Parasitic cooperation This mode of cooperation is one in which agents form a parasitic relationship, whereby one agent benefits while the other is harmed.

Symbiotic cooperation This mode of cooperation is one in which agents form helpful partnerships in which every agent involved in the partnership receives a beneficial side-effect.

Competitive cooperation This mode of cooperation describes simulations where symbiotic partnerships are formed and these partnerships directly compete with other partnerships to better themselves, while not necessarily harming others.

Antagonistic cooperation This mode of cooperation describes simulations where each agent is acting independently (much like in pure cooperative simulations), however, each particle is also competing with the other particles to better themselves. Such an antagonistic cooperative system often produces highly competitive agents which are extremely well suited to their intended purposes.

Pure antagonism This mode of cooperation is the only mode in which agents aren't out to better themselves, but rather to directly harm the other agents.

2.2 Level of detail

Going hand in hand with the previous theme of behavior refinement, the next major issue is the issue of abstraction. How detailed does a researcher have to make a simulation to approximate real-world results? If too few details are implemented, then the whole system could fail in the long run to produce anything resembling human (or other creature's) behavior. But on the other side of the coin is the limitations imposed upon researchers by computational resources available. There is no supercomputer available that would be able to run even a reasonably well detailed simulation of the entirety of human social interactions in anything resembling a decent time frame–a rough estimate on my part is that anything even close to realism would take about 10 years of cpu time (accounting for the furthest advances of computing in the year 2012) for every second of simulation time... which would not provide any reasonable results. Not to mention, most of what happens in social interactions is yet to be discovered, so could not possibly be modeled.

The point here, is that simulations and modeling are trade-offs. There is no such thing as a fast, perfectly accurate simulation of social behaviors and interactions, and dare I say there never will be. But researchers are able to trim the fat, so to speak, and approximate with some discerning skill elements to simulate or avoid in their work.

Some popular examples of elements that are often left out of simulations because of computational constraints are economics and government.

2.2.1 Economics

A system of economics is certainly vital to social interactions in the real world, but in many cases modeling an entire economy is painstaking and not truly worth it. In many cases, a simple economic system of assigning agents to a social class or having a "worth" attribute is close enough.

2.2.2 Government

A system of government may well affect our communications and interactions. However, implementing an entire governmental body into a simulation would prove difficult at best, and even then would likely not embody anything that would have a noticeable effect on results. Very few social simulations that I have come across actually have a governmental body simulated within them.

2.3 Cyber-ethics

Moving on, we come to the issue of ethics in respects to artificial social agents. To many, this is a non-issue; software agents are not to be considered real creatures, so there should be no need to define their civil liberties and inherent rights. But this is

not necessarily the case. Basic rights and denial of those rights drive a large segment of our interactions with each other. We have the ability to be heard, so we speak. We have the ability to think and philosophize, so we do so. And when those rights are taken away, we fight. Why should software social agents be any different? This is again a case for abstraction, but this is a slightly different approach than previously discussed.

2.3.1 The "rights" of a software-defined social agent

So, what rights should a social agent have? Mobility seems important, as does self determinism. Basic communication rights as well, and the ability to share information with others.

But can these rights be taken away? This is a branch of social simulations that I have not discovered any work on, and with good reason: as mentioned above, this is a non-issue for many researchers. But I propose that this is an area that be focused on in future simulations studies. What happens in simulation when an agent is no longer able to move? Do they make due with what they have? Do they find a loophole, forming a chain of communication to obtain the resources they need? Or do they languish and die, if the simulation allows for death? How do these rights affect the outcomes of the simulations?

2.3.2 When does simulation become reality?

Waxing even further philosophical, we come to the question of when giving social software agents rights becomes a moral obligation. The question, rephrased, is when does a simulation stop being a model and start being a reality? There is a line somewhere between rudimentary computational simulation and full-blown reality where a simulation has enough qualities of the real world that it becomes a real world of it's own, but where that line lies is currently unknown.

For now, we can suffice it to say that a perfect one-to-one simulation of the real world would encompass this realness and be in effect a complete, real world in itself. At this point, the agents would be computationally defined humans, and have intrinsic rights that would have to be allowed. But this situation also leads us into an infinite regress, with the simulated real-world having to have a simulated real-world within it, and so on to infinity. This would require infinite computational power, which, short of highly improbable breakthroughs in quantum computing, hypercomputing, and relativity, is impossible in our current model of computation.

If the requisite level of reality for a simulation to impart rights upon it's agents lies closer to rudimentary simulations than perfect copies, then this is an issue that we will have to think about at some point in the future. Until then, though, we're just fine with the agents we have.

2.4 What models prove, if anything

The largest issue that has to be addressed in social simulation as it stands today is what scientific use these models are to us. The resounding question is "what does this prove?" In most cases, the answer to this question is an anti-climactic "nothing."

But why simulate then? If these models prove nothing, then what is the point? I say *insight*. Social simulations studies allow us to gain insight into how things could possibly work. That is how science evolves: someone has an insight, and proposes a theory. The theory is either backed up by empirical evidence (in the form of rigorous experimentation) or shown to be flawed. The theory is either revised or rejected outright. Science evolves. These simulations play a very important role in providing the catalyst for such sparks of insight.

There are several criteria, however, for determining if an insight provided by a model is one that should be further pressed, based on properties inherent to the model.

2.4.1 Soundness of the initial conditions

Assuming the model appropriates and approximates reality fairly consistently with the real world, the insights sparked by these simulations should be taken seriously. However, this is an extremely subjective criterion, and there is no set rule for determining goodness of fit between a social model and the real world, so your mileage may vary.

2.4.2 Repeatability of the experiment

A second, less subjective criterion is how repeatable the experiment is that provided the insight. Given that social simulations are in large part computer programs, this is as simple as running the program in batch for some large number of trials and collecting the data. Filtering out the outliers, this method ensures that the insight wasn't sparked by some fluke of the random number generator, but rather is an emergent property of the system that you are examining. However, this should be taken with a grain of salt, as some of the most purveying ideas in science have been influenced by a fluke in experimental conditions. Again, your mileage may vary.

3 Final words

I would like to close this paper with a quick summary of what I feel are the salient issues at hand in social simulations theory.

Perhaps the most philosophical of my concerns deals with an issue at the fringe of the theory of social modeling: the ethical issues involved in intelligent agents. While

not necessarily immediately applicable, nor for the forseeable future, it poses an interesting question that I feel is worth exploring. Agents are eventually creatures; add enough properties of life to any object and it will become life. And in many ways, with social simulations, the researcher is playing god on an incredible, omnipotent scale.

Stepping back a bit, we have the still philosophical but readily applicable issue of abstraction. How far we abstract away the details of the world surely affects the validity of the model, but at what point does the model become a reality of it's own? Should things like an economy and a governmental body be simulated? What is the scope of the investigation at hand? These are all very important issues with any simulation that must be answered with care, lest your model become useless.

With broken models, those that are overzealous in their abstractions, we have the question of validity. Do the results mean nothing? Are the initial conditions bunk? Can the results at least be repeated? Often, these models will fail based upon their model of communication and information.

With communication and information comes the great power of these simulations. But, much like Stan Lee stated (and Voltaire before him), "with great power comes great responsibility." Researchers must take great care to make sure their models aren't overstepping their bounds, while at the same time ensuring that they have enough freedom to accomplish something resembling communication or socialization. The ability for models to share and reflect upon the information they have been entrusted with is where social modeling shows it's strength, and I feel that this is the issue that needs most attention in the coming years.