SOC 527 | **Michael Macy**
Artificial Societies | **Uris 372**
Thurs. 1-3:30 | **255-4187**
Office hours: Thurs. 10-12 | **mwm14@cornell.edu**

This course introduces agent-based computer simulation as a tool of theoretical research. Agent-based models are used to explore the dynamics of interaction among interdependent adaptive decision makers. More precisely, the models impose three key assumptions:

1. *Agents interact with little or no central authority or direction.* Global patterns are generated from the "bottom up," determined not by centralized authority but by local interactions among autonomous decision-makers. Put differently, we do not program the global social system to act, adapt, or change; we program the decision-makers whose interactions comprise that system.
2. *Decision-makers are adaptive rather than optimizing,* with decisions based on heuristics, not on calculations of the most efficient action. These heuristics include norms, morals, habits, protocols, rituals, conventions, customs, routines, etc. They evolve at two levels, the individual and the population. Individual learning alters the probability distribution of rules competing for attention, through processes like reinforcement, Bayesian updating, or the back-propagation of error in artificial neural networks. Population learning alters the frequency distribution of rules competing for reproduction through processes of selection, imitation, and social influence. Genetic algorithms are often used to model adaptation at the population level.
3. *Decision-makers are strategically interdependent.* Strategic interdependence means that the consequences of each agent's decisions depend in part on the choices of others. A classic example is a "Prisoner's Dilemma." When strategically interdependent agents are also adaptive, the focal agent's decisions influence the behavior of other agents who in turn influence the focal agent, generating a complex adaptive system.

Accordingly, the course will focus primarily on the use of simulation to study the evolution of norms that affect individual behavior in social dilemma situations. The classic in this genre is Axelrod's celebrated computer tournament. From there we will look at genetic algorithms, learning models, and cellular automata. Applications will include processes of diffusion, social influence, informal control, trust, support systems, and management fads.

The course is a "hands on" lab course, not a reading course. We will have a very limited amount of reading, but students are expected to read the assignments very closely, with the aim of understanding the methodology sufficiently to explicate (if not replicate) the experimental design. The readings are to be used as templates for seminar projects.

Seminar members will develop a project at one of the following levels:

1. Criticism of an article based on agent-based computer simulation. The criticism should focus on the logic of the model and should propose (but not implement) a detailed alternative specification. The project will consist of a seminar-length paper (about 15-20 pages). The criticism should include some discussion of the source code where possible, but programming ability is not required. Thus, those without previous programming experience may prefer this approach.
2. Modification of an existing computer simulation. This is an extension of "1" but requires the ability to edit and de-bug a computer program. Depending on the extensiveness of the modification, this project might take the form of an oral presentation and demonstration of a working program, rather than a paper.

3. Creation of a new computer simulation. This is an extension of “2” but requires the ability to design and write an original computer program. This project will take the form of an oral presentation and demonstration of a working program, rather than a paper. Those electing this option are encouraged to work in two-person teams.

Seminar participants will also take turns making presentations to the group that summarize the reading (focusing on the experimental design) and suggest ways to elaborate or refine the simulations.

The readings are journal articles, not books, with one exception. Axelrod’s *Complexity of Cooperation* will be available at the Campus Store. All other readings can be downloaded from the course website. Seminar members are also encouraged to purchase Borland Delphi (version 2.0 or higher). Versions before 5.0 are available on the Web for under $20.

Feb. 1: Introduction.

We begin with an overview of agent-based computational modeling in the social sciences and how this method differs from mathematical models, and from an earlier generation of computer simulations. Our discussion will try to assess the usefulness and limitations of this method of modeling social life and grapple with the epistemological controversies. In what sense is a computer simulation an “experiment” or a “test” of a theory?

- Gilbert, “Simulation: an emergent perspective” http://www.soc.surrey.ac.uk/research/simsoc/tutorial.html
- Epstein and Axtell, *Growing Artificial Societies*, Intro and Conclusion
- Macy, “Social Simulation”
- Axelrod, “Introduction,” in *Complexity*

Feb. 8: The Ecology of Cooperation.

This week’s reading introduces computer simulation of strategically interdependent agents playing iterated Prisoner's Dilemma. The application is a computational ecology in which strategies compete for survival but do not evolve. The classic in this genre is Axelrod’s “computer tournament.” Kollock modifies Axelrod’s model by introducing “noise.” Bendor and Swistak criticize Kollock’s model and Kollock replies. These articles give examples of the three types of student projects: criticism (Bendor and Swistak), modification (Kollock), and original construction (Axelrod). Is "noise" useful in simulation models, or is it preferable to study cooperation under ideal conditions (analogous to the study of motion in a frictionless environment)? Axelrod claims that reciprocal altruism is a highly robust strategic principle. Are his results artifacts of a peculiar experimental design?

- Robert Axelrod. 1984. *The Evolution of Cooperation,* chs. 1-2.
- Peter Kollock, "An Eye for an Eye Makes Everyone Blind." *American Sociological Review*, Dec. 1993.
- Bendor, “Comment on Kollock” and Kollock’s reply, *ASR*, April, 1996, pp. 333-346.

Feb. 15: Genetic Algorithms and the Evolution of Cooperation

The agents in the ecological models from last week adapt behaviorally but not structurally. That is, they use rules that respond to the actions of others but their rules cannot change. One problem is that the results of these experiments depend entirely on the initial distribution of strategies in the population. This week, we elaborate Axelrod’s model by allowing the rules to change, not just the behaviors. Genetic algorithms provide a simple and elegant way to write game-playing strategies that can improve their performance by building on partial solutions. What is the difference between ecological and evolutionary models? What are the advantages/disadvantages of each? Can genetic algorithms be used to model learning? Imitation?

Readings:

- Axelrod, *Complexity of Cooperation*, chs 1 and 2.
- Macy and Skvoretz. 1997. "Trust and Cooperation Between Strangers" http://www.people.cornell.edu/pages/mwm14.

Feb. 22: Learning Models, Cooperation, and the "Law of Effect"

Last week, we introduced evolution at the population level, in which selection pressures operate on the frequency distribution of competing strategies. But evolution can also operate at the individual level, as *learning*. The Bush-Mosteller stochastic learning model is one of the simplest examples. Agents repeat behaviors that are satisfactory and abandon those that fail. An artificial neural network is an elaboration of this idea that allows for the evolution of *conditional* strategies, in which a response is triggered by a specific stimulus pattern.

What are the differences between the Bush-Mosteller model and an artificial neural network? What are the advantages and disadvantages of each? Compare and contrast "evolution," "learning," and "rationality." Try using Venn or tree diagrams to express your conclusions. What difference does it make if a model is deterministic or stochastic? Predict what would happen if my simple 2-person PD model had been deterministic. Would the results be the same? What is the difference between reproduction (as in a genetic algorithm) and reinforcement as adaptive mechanisms?

Readings:

- Bainbridge 1995. "Neural Network Models of Religious Belief." *Sociological Perspectives*.
- M. Macy, "Natural Selection and Social Learning in Prisoner's Dilemma: Co-adaptation with Genetic Algorithms and Artificial Neural Networks." *Sociological Methods and Research*, Aug. 1996, pp. 120-137.

Mar 1: Hopfield Nets as Models of Homophily

Hopfield Nets differ from "feed-forward" neural networks in an interesting way: Rather than being arranged hierarchally (from input to output layers), all the nodes are fully connected. Learning is based on "Hebbian" principles rather than error-correction. Hebbian learning corresponds to what sociologists call "homophily," or the tendency for likes to attract and opposites to repel. If we add to this the principle of diffusion (those who interact tend to become more alike) and the principle of negativity (influence includes differentiation as well as attraction), we have a model that captures important properties of group dynamics, including Heider's balance theory and Tajfel's social identity theory.

Readings:

- Nowak, A. and R. Vallacher 1997. "Computational Social Psychology: Cellular Automata and Neural Network Models of Interpersonal Dynamics," in S. Read and L. Miller (Eds.) *Connectionist Models of Social Reasoning and Social Behavior*.
- Kitts, Macy, and Flache 1999: Structural Learning: Attraction and Conformity in Task-Oriented Groups, *Computational and Mathematical Organization Theory*, 5:129-45. **Attractor Neural Net Program for Structural Learning Network Model**

Mar 8: Cellular Automata

Hopfield Nets are a way of modeling dynamical social networks, that is, networks whose structure changes over time. Cellular automata are an alternative network architecture that is useful for studying social distance using a two-dimensional spatial metaphor. CA models have the advantage that they allow for highly intuitive visualizations of network evolution. But while these models may be useful for studying geographical neighborhoods, critics question the spatial metaphor. Can CA results be generalized beyond geographical social networks?

Readings:

- Hegselmann and Flache. "Understanding Complex Social Dynamics: A Plea For Cellular Automata Based Modelling" JASSS, 1:3 http://www.soc.surrey.ac.uk/JASSS/1/3/1.html
- W. Clark 1991. Residential Preferences and Neighborhood Racial Segregation: A Test of the Schelling Segregation Model *Demography*, 28: 1-19.

Mar 15: Fads and Diffusion: What Goes Up Must Come Down.

Diffusion is usually modeled as a process like ink spreading in water. But many social influence processes cascade in both directions, with "bandwagons" followed by sudden collapse. Diffusion might also seem to imply a tendency toward homogeneity, but this also is too limited a view. In short, there can be diversity in both time (fads) and space (differentiation).

Readings:

- Macy and Strang, "Pluralistic Ignorance and the Top Secret Management Handbook."
- Axelrod, "Disseminating Culture," in *Complexity*

Mar 22: Spring Break

Mar 29-April 5: Programming workshop

These two weeks will be devoted to learning to read, criticize, modify, and write source code in the Delphi/Pascal language. We will practice on source code used in the readings from the first half of the course and begin organizing our seminar projects.

Reading:

- Delphi Instructional Guide

April 12-26: Presentation of Projects.

The final three weeks will be a simulation workshop in which all seminar members (students and faculty) will present their projects to the class for discussion and criticism.