Voting Processes in Complex Adaptive Systems to Combine Perspectives of Disparate Social Simulations into a Coherent Picture

Deborah V. Duong

Augustine Consulting/US Army TRAC Monterey, Monterey California
dvduong@nps.edu

Abstract

If computational social science is to find practical application in informing policy decisions and proportionately analyzing courses of action, then it will have to make progress in the area of composition of social models. Since a single simulation cannot hold a world of information, policy makers need to switch in and out modules in federations of simulations to test policies against all possible social environments. Voting processes as they occur in nature, both in the form of cognition in a human mind of disparate world views, and in the form of equilibria seeking coevolution of species, inform how to combine model results externally and deeply, respectively. These algorithms, which use the same principles of soft computation found in nature, enable any models to mesh together, even if they have different ontologies, or their data conflict, regardless of the degree they overlap. A whiteboard architecture in which models report in their own ontologies how other models may inform them and what they have to offer other models, is a framework for the arbitrary meshing of social models.

Voting Processes as Seeking Equilibria in Complex Adaptive Systems

Voting processes are critical to the coevolutionary engine of structure and of growth in Complex Adaptive Systems. Complex Adaptive Systems are living systems made up of coevolving agents, that is, entities that perceive and act in their environment based on their goals, and optimize their chances of achieving their goals by changing the way they perceive and act given the way other agents perceive and act. Social systems and biological systems are the two primary types of Complex Adaptive Systems in nature, both having their own forms of coevolution, based on social learning and biological evolution respectively. An understanding of these processes in nature informs instructs scientists how to emulate them for the computational science purpose of learning more about those systems. In particular, voting systems are an essential ingredient to bringing a federation of social simulations with disparate perspectives on the social world into a consensus and thus a coherent scenario for social analysis.

Complex adaptive systems views of evolution encompass a notion of consensus in the concept of equilibria. Biologist Maynard Smith, in his concept of Evolutionary Stable Strategies (Maynard Smith, 1982), introduced the idea that species adapt to each other, moving towards a stable point in which no species can improve itself, that is, Nash Equilibrium (Nash, 1950). This state where no species can improve is a form of consensus between species, in that it can be viewed as made up of the votes of all of the stakeholder species that have a say in the system through their behavior. These equilibria are the "order" in living systems, the dynamic structures that form the temporarily stable objects and concepts of complex adaptive system. These structures are temporary because, in reality, species always approach equilibrium, but never actually reach it because a side effect of approaching it may be that it changes. Such side effects are a property of coevolutionary systems, coming from the fact that species can be quite flexible in their ability to adapt to other species. The movement of this equilibrium point is what causes growth and complexification of biological systems.

In social coevolution, Douglas North of the New Institutional Economics (NIE) School adapts a similar explanation for economic change: that individuals, in their seeking of knowledge to reduce the uncertainties inherent in socio-economic action form expected patterns of behavior called "institutions." (North, 1981). These institutions are also examples of Nash Equilibria, points at which no individual can change behaviors in a way that will better his situation. As in biological evolution, Nash equilibria are a form of consensus, because the interests of many stakeholders have contributed to the formation of that equilibrium over time. However, as in biological evolution, as individuals learn what to expect, the new actions that result from what they

learned actually change what to expect as other individuals creatively adapt to their behavior, resulting in moving the equilibrium point, and growth in both action and perceptions that make up a social milieu.

Equilibria Seeking and the Sociology of Knowledge

It is important to understand that in the evolution of human institutions, when looked at as a compromise between stakeholders over time, the compromises include not only compromises on behaviors that could be from entirely different world views, but more importantly, compromises between the mental models that generate the behaviors. Institutional equilibria include not only expected behaviors, but also expected perceptions of behaviors, and more fundamentally, expected perceptions of how the world works and world views that determine how individuals will react to existing behavioral constraints. In interpretive social science, which the NIE is part of, the cognitions that determine institutional behaviors are important. The sociology of knowledge, the study of how world views change and evolve over time, also informs that compromise of stakeholders, that voting process of behaviors and world views that seeks elusive equilibria.

From the perspective of the sociology of knowledge, disparate world views can and do merge to form transformed world views which are more than the sum of the parts, in voting systems that cause knowledge to grow. In the sociology of knowledge, movement towards equilibria does not take place on the level of individual interest, but the level of world views and how these different models of how things work may come to make sense together. However, it is interesting to look at the evolution of world views from the NIE perspective of individual interest and cost. In the sociology of knowledge, our idea of how things work comes from many origins that are transformed by their contact with other ideas of how things work. Like the elephant of Indian legend, separate perceptions of an elephant's leg and his trunk can form a world view that takes more into account, shedding light on the fact that what one blind man perceives as a tree and another blind man perceives as a snake make something different when together. The perspective of how an individual may put different world views together to seek goals can inform how coherent narratives evolve.

Even when different world views come into conflict as one might expect in Globalization, the appropriate level of analysis is the individual to see how world views combine in the seeking of goals. Methodological individualism is important because the place where the battle between ideas takes place is the individual. It is at the level of the individual where we see the effects of cognitive dissonance created by different world views, and how they combine to

inform decisions which result in subsequent actions. According to NIE, world views, or ideas of how the world works, determine what actions we take are beneficial and what are harmful.

World views themselves evolve to serve the interests of individuals. According to the Narrative Paradigm, legends are good examples of world views that have evolved over time (Fisher, 1985). Cultural historians note recurring themes in legends, for example, a recurring "Cinderella Story" across cultures with the same theme of a poor orphaned girl becoming a princess. The reason why these similarities exist is because the legends themselves, as they are told orally over and over, evolve to better fit psychological needs, many of which are in common across cultures. In effect, the legend becomes the vote of many story tellers each trying to fulfill needs.

The needs fulfilled, however, need not be the same for all parties or even conscious to any party. Gould's theory of "Spandrels" in biological evolution (Gould, 1979), that structures start to evolve for different reasons than they continue to evolve, can be applied to social evolution as well: that institutions start in order to fulfill different needs than they continue to fill, and these needs are not necessarily conscious to the participants. One example is the Eskimo women that are leaving the large fish for their husbands, while they eat the smaller fish, that happens to have Vitamin D in it needed for health in pregnancy. They think they are doing so to honor their husbands, but the real reason exists outside of their world view.

Applying Equilibria Seeking Systems to Automating Consensus amongst Disparate Social Models

Natural processes of the combination of world views in society may be used to inform how world views in the form of simulations can be combined to form coherent pictures of the social environment. Policy makers have a need to combine the results of many simulations to make coherent wholes that inform policy decisions. In Computational Social Science, the social world is too complex to be put in a single simulation, and as social theorists disagree than policy decision makers must test their policies for robustness against many possible social environments. For this purpose the Department of Defense invests in integrative modeling frameworks for its Course of Action Analysis (Duong, 2010). These integrative frameworks need to take the conflicting data that comes from different simulations that model different schools of thought in different subfields of the social science, with concepts categorized into different perspectives and resolutions, fed with conflicting static data, and make a coherent picture out of it to inform policy decisions under uncertainty.

A simulation model of how the world works, and a men-

tal model of how the world works are both models, and the rules of model evolution and consensus that make a coherent world view that exist in the real world may be applied to simulation models to seek consensus, just as natural principles of cognition have been applied to neural networks and natural principles of evolution have applied to genetic algorithms. In natural language processing, the simplest model of consensus is direct voting system. For example, a voting system of entity extractors applies several parsers to a free text document separately, to tell which words in the document represent a person or a place. Voting systems of several extractors, in which a (possibly weighted) majority vote decides whether a word is a person or place or entity at all, tend to have greater accuracy and precision than any one extractor. The greater accuracy comes from that fact that there are more ways to be incorrect than to be correct and extractors that use different algorithms to come up with the same results are unlikely to be incorrect. It's the same principle that teachers often use to tell if their students are cheating: if they consistently give the same incorrect answers, they are likely to be sharing information, but not the same correct answers.

The same principle could be applied to voting systems of simulations. However, this voting system need not deal in exact outcomes; it could deal in types of outcomes. Further, the event in one simulation need not be compared to the same event in another simulation, but to an event that should be correlated with it. If a distance between events may be established, then a measure of how close or far events are and whether they agree or disagree can be defined.

Although simulations can hold the causal, theoretical relationships of the social sciences, theories really predict softer patterns of outcomes than single runs of simulations provide. If the simulation represents a theory well, then repeated runs will capture the patterns of outcomes predicted by the theory in proper proportion. Multiple runs create classes of outcomes, grouping outcomes according to the types of outcomes inherent in the concepts of a theory. For example, if a theory on state stability may classify a riot in the same category as a strike, in that they are both indicators of instability.

We know if a simulation agrees with the data that is expected from a theory if the output runs are properly categorized into the concepts of the theory, and the patterns between types of events exist as predicted by theory. This categorization and pattern matching is not a far cry from how theories are tested in the real world by scientific experimentation, where categories of outcomes are similarly defined beforehand that would show patterns predicted by theory. As long as we have categorized events in the same ontology as theory, as defined by the causal relations in the theory, then we can characterize how close outcomes are, and determine if they agree for a "vote."

So in seeing if simulation of the same kind of phenomena agree or disagree in voting systems, they have to be categorized into the same concepts. For example, economic simulations would have to have their concepts matched in order to be compared despite the fact that they may be from different schools of thought that have different concepts. However, what saves the day is a combination of the real world, from which all these categorizations are ultimately derived, along with the existence of another kind of social studies, scientific experiments that can find correlations between phenomena. It does not matter that concepts may be different in different simulations, only that, if there exist studies of correlations of the phenomena of their respective ontologies, then we know how often phenomena should occur together in order to say that models "agree" and have thus "voted the same."

When we bring translations across ontologies through correlation into the picture, many new possibilities for voting systems and consensus emerge. For one thing, we no longer have to have simulations about the same field of phenomena to vote, they can now be simulations about different fields that science has found to correlate a certain percentage of the time. This is fortunate because the different fields of the social sciences are highly correlated, as different perspectives of the same phenomena. For example, economic states correlate with psychological states and social stability states as well, so to make a coherent picture of an entire social environment, simulations of these separate states can "vote" on what a consistent picture is. For example, an economic downturn in an economics simulation and a psychological distress in another may offer mutually reinforcing support, while a third simulation of social prosperity may be the odd man out, depending on the findings of social correlative studies.

However, these votes are not simple, because they are not votes for the same exact events as in the voting system of entity extractors, but are votes for events that usually or often go together. More can be done with less information in this case, because even if you lack a study that tells you that two types of phenomena should occur together or in sequence, you may have a study that tells you that a third type of phenomena should appear with both, so that the phenomena are mutually reinforcing at a higher order. Correlative studies tell us which phenomena give evidence in the support of the existence (or non existence) of other phenomena, and enough of these can help us to take many different pieces of evidence and have them vote, or come to a consensus, on the most coherent picture possible.

This voting amongst many pieces of possibly conflicting evidence can be modeled in a constraint satisfaction neural network, a network where nodes represent states of the world that exist or not, and links represent evidence that one state of the world supports the existence of another. The output of simulations can be the states, and correlative

studies can be the links. If several simulations were to vote on what a plausible next state is after a policy intervention, then we could perform an optimization on their output to obtain several coherent scenarios at various degrees of likelihood. This consensus state could then be put back into the simulations that have a checkpoint-restart ability, once the consensus state is expressed in the ontological concepts of the simulation.

Does this method in fact tell us more than any individual simulation model tells us? Is the result more than the sum of its parts? Remember that each simulation represents a different world view and that their results had to be matched to other results of a different world view, either directly or indirectly through a correlation. They could be very different, and based on contradictory theories. They could have had mutually reinforcing results for entirely different reasons. However, the reason that they have results that are mutually reinforcing does not matter: because of the connection to real world experimentation, it still makes sense to have them as reinforcing. A school of thought would not exist if it was not to some extent predictive of phenomena, even if it was not completely general and always accurate. Soft constraint satisfaction allows us to take theories that are partially wrong, and make them better through consensus by seeing where they agree. After all these causal theories exist in order to predict real world results, and they would not survive if they did not do this with at least some degree of accuracy.

Further, a constraint satisfaction neural net is a model of how an individual mind puts dissonant pictures of the world together, to include what makes sense according to a person's identity and values in accordance with goals. Constraint satisfaction neural nets have been used to simulation cognitive dissonance, the cognitive properties an individual uses to resolve conflict and achieve personal goals. As in the voting systems of nature, the mind finds a Nash Equilibrium where the network converges on a satisficing, path dependent set of beliefs. Therefore, to combine world views with a CS net, we do emulate the way it is done by people in the natural world.

However, one ingredient is still missing: adaptation. The models themselves did not change and adjust. The concepts on which the models are based do not improve, even if they in combination are more than the sum of their parts. At most, a human analyst could look at the "snake" and the "tree trunk" that they separately make, and abduct that when these always occurred together that a new model of an "elephant" would be more parsimonious, general, and predictive. However, unless the concepts of the models were regrouped into different categories and relations, nothing would be there to represent the elephant. Humans have the ability to abduct, to invent new concepts.

This abduction, the ability to draw analogy across different perspectives and create new generalizations, is essential to a true combination of world views. However, if those simulations are adaptive, then it is possible to have simulations which increase in their ability to predict. After all, the enterprise of computational social science is the exploration of our assumptions about how agents are differently motivated and how they learn to adapt under different circumstances and given the adaptations of other agents. This coevolution is the method by which they approach elusive Nash Equilibria. If we do not do an actual computation of motivations under new circumstances, including the adaptations of the other learning agents, then we are not learning anything from a computer simulation run that we didn't tell the agents to do in the first place. Unfortunately, it seems that the science of agent based simulation has not advanced to the stage of being able to explore new situations that it has not programmed in the first place. The problem is at its roots, a problem of ontology. If it's new and thus foreign to the present perspectives of the agent, how do we correctly get the agent to interpret the new phenomena, even though it may come from an entirely different point of view, to navigate the world in achievement of its goals? How do we get it to actually be able to interpret and navigate a variety of different environments, in a federation of simulations that actually recombines perspectives and world views?

One answer is to again learn from nature, leveraging the nature of coevolution itself. The ecosystems we see today are actually full of animals that are imported from other ecosystems, and the importation of new concepts into an ontology is analogous to this process. To deeply combine simulation models requires a combination of ontological concepts which is similar in nature to the combination of ecosystems. In coevolution, every species adjusts and adapts to every other species. The way a species becomes a part of an ecosystem is through a process of mutual adaptation. Suppose a fast predator was introduced to a new ecosystem: it would put selective pressure on its prey to become fast as well. But that would cause the predators of the new ecosystem to become fast in order to catch the prey. Thus the seeding of the predator makes the existing predators of the ecosystem take on its traits. Similarly, adaptive agents in coevolutionary simulations can adapt to the results of other simulations and actually take on the nature of those other simulations, while retaining the concepts that make up its own point of view. By translating the results of one adaptive simulation into another, the representation of those results will become a part of the other simulation. Furthermore, according to Gould it doesn't matter why, the simulation need not have the same causal basis, only the effect, to combine and become an essential integrated part of the new simulation. For example, one simulation may be a mental model that includes an airlift of supplies to an isolated island, and the other simulation may be an islander's mental model that includes a large bird dropping

food, but the existence of the airplane/bird and the food will be occur together in both. As in nature, this coevolutionary combination is a type of voting that approaches but never reaches Nash Equilibria, but that equilibrium expresses an increasingly coherent picture of the environment.

Nature informs computational social science on how to combine the disparate world views of a federation of simulations into optimally coherent pictures that are more than the sum of their parts. These include both the cognitive combination of world views in Soft Constraint Satisfaction programs for an external consensus, and the combination of adaptive simulations that adjust to each other for a deeper internal consensus.

Bringing Together any Combination of Models in a Whiteboard Architecture.

In natural methods of combination of perspectives, whether models combine deeply by changing each other's outcomes, or shallowly in external voting processes, what one model needs to know about another is a result, not a process. What models need to know is "what" a result is, not "how" another model got the result. One model thus makes a functional specification for another, as opposed to an implementation specification. This fact, that there is an interface with which a model can express its results while hiding the process, has implications for the combining of models. It means that almost any combination of models can be combined, as long as the interface is defined functionally, in terms of what a model should be informed of in order to by in sync with the other models. Even if that interface is in the unique ontology of the model, as long as there is a direct or indirect match with the correlative phenomena of other models, then a consensus may be found. Therefore, all that is needed to mesh models is the outer functional expression of states in their own ontology, and mediation between the ontologies through simple correlation. Those requirements can use the deep coevolutionary or simple external constraint satisfaction to make sense of the disparate data. This is true whether the models are completely overlapping in the same phenomena, partially overlapping in correlated phenomena, or completely independent. The algorithms of soft computation from nature do not even care if the models contradict each other. We have simple voting process to deal with models that come up with the same phenomena, and optimization voting algorithms to deal with models having correlated phenomena. If the phenomena are correlated, model results will be matched, and if are independent, then we have orthogonal conditions with which to cover an output space. The process of meshing models is a form of limitations on input to make feasible sets of parameters for the data farming of the output space. A probabilistic ontology can hold all of the information needed, from the unique concepts of the individual models to the probabilistic matches from the data of correlative social studies (Duong, Makovoz and Singer, 2010).

The point of combining models is to switch them in and out in a feasible manner, creating a proportionate outcome space for course of action (COA) analysis. In this framework, models influence one another other by means of correlative connections to form feasible sets of parameters that express a coherent picture of the social environment, and a proportionate output space. The most general architecture that can perform this probabilistic inference on any combination of models is a whiteboard architecture. This is an architecture in which every model specifies what it can use and what it can produce functionally, and that information stays around for later data mining of possible higher order parameters. The soft computational algorithms need not be performed on simultaneously co-occurring output of models, but may be applied to sequences of states in the form of a Markov Process that characterizes the sequential dynamics of models and/or time series data. If each model specifies its parameters, the functional states that it uses to match up with other models, in a whiteboard architecture, then a coherent voting system of almost any combination of models may be created.

References

Duong, D. 2010. "Verification, Validation, and Accreditation (VV&A) of Social Simulations" Spring Simulation Interoperability Workshop, Orlando, April 12-16 2010.

http://www.scs.gmu.edu/~dduong/SIWvalidation.pdf.

Duong, D, D. Makovoz and H. Singer (Names listed in Alphabetical Order). 2010. "The Representation of Uncertainty for the Validation and Analysis of Social Simulations" Fall SISO Conference, Orlando, September 2010. http://www.scs.gmu.edu/~dduong/SIWprobOnt.pdf.

Fisher, W. 1985. "The Narrative Paradigm: An Elaboration." In *Communication Monographs* 52. December. pp. 347-367.

Gould S. and R. Lewontin. 1979. "The Spandrals of San Marco and the Panglossian Paradigm: a Critique of the Adaptionist Programme." *Proc. R. Soc. Lond.*, B, Biol. Sci. 205 (1161): 581–98.

Maynard Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge University Press.

Nash, John. 1950. "Equilibrium points in n-person games" *Proceedings of the National Academy of Sciences* 36(1):48-49.

North, Douglass. 1981. *Structure and Change in Economic History*, Norton.