Comput Math Organ Theory (2009) 15: 47–59
DOI 10.1007/s10588-008-9048-9

Computational modeling for reasoning about the social behavior of humans

Kathleen M. Carley

Published online: 1 November 2008

Abstract The number of computationally-based models of human social behavior is growing rapidly. In fact, the current ease of programming is resulting in a plethora of tools with impressive interfaces but little theoretical power under the hood. Further, the overabundance of new toolkits for building models is facilitating the excessively rapid growth of simple proof-of-concept, or intellective, models. The current state of models range from the simplistic to the elaborate, from the conceptual to the empirical, and from the purely notional to the validatable. This review briefly describes the state of human social behavioral modeling. Key issues surrounding analysis and validation are discussed.

Keywords Dynamic network analysis · Social networks · Agent based models · Multi-agent simulation · Network science

1 Introduction

Computational modeling is a growth area in the social and behavioral sciences (e.g., Harrison et al. 2007). In general this refers to any modeling effort in which a model is described within a set of computer code. This includes a computer program, or network of computers and programs, that attempt to operationalize an abstract model of the system. Such models are also referred to as computer simulation, computer models, and computational models (Law 2007). In a purely mathematical model, the relations are expressed in mathematical terms and processing is done by solving the equations. Computational modeling is a form of mathematical modeling, typically used when a closed form solution is not possible (Ross 2006). In a computational

K.M. Carley (✉)
Institute for Software Research, School of Computer Science, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, USA
e-mail: kathleen.carley@cs.cmu.edu

model the relations are expressed in mathematical or symbolic terms and processing is done by following an algorithm.

There are many types of computational models (Zacharias et al. 2008). Among the most common forms are agent based models (ABM), system dynamic models, event based models and statistical forecasting. Agent based models are also referred to as multi-agent systems. Agent based and system dynamic models are also referred to as complex system models. Reviews of computational models seek to characterize such models along a wide number of dimensions. These include, but are not limited to: intellective versus emulative, stochastic[1] versus deterministic,[2] steady state[3] versus dynamic,[4] continuous[5] versus discrete,[6] rule-based versus equation, learning versus static versus optimization, centralized multi-agent versus distributed multi-agent, local versus distributed,[7] system dynamic versus multi-agent versus multi-agent network. While modeling frameworks typically fall in to one of these categories, models developed to address real world problems that are not built in a framework are typically hybrid models crisscrossing these boundaries at will.

In comparison to traditional formal (i.e., mathematical) models, computational models have the following characteristics. They tend to be larger scale; e.g., they include more events, more actors, more entities, more time periods, etc. They tend to focus on the process and intermediate solutions and not equilibrium solutions (which are the key result of mathematical models). They tend to utilize a mix of simulated and real data as opposed to being completely algorithmic; e.g., many computational models use simulated actors using real equipment, or in real social networks. Computational models tend to handle more complexity such as a greater number of interacting parts, higher levels of non-linearity in relationships, and very non-continuous response surfaces. As such, these models are often referred to as "complex system models." Finally, due to the typically stochastic nature of the results, impossibility of calculating a complete response surface, and the attention to intermediate results statistical analysis is typically used to provide an interpretation of model outcomes for computational models.

There are a number of reasons to use computational models. These include:

- *Ethical*: You cannot test policies on real populations but can on simulated populations.
- *Preparatory*: You can use these models to create hypothetical situations with more potency than existing ones. As a result you can use the models to examine a wide range of scenarios. This enables a more systematic imaginative thinking and facilitates training.

[1]Stochastic models typically have at least one random number generation component.

[2]A special case of deterministic models are the chaotic models.

[3]Steady state models typically use a set of equations to define fixed relations.

[4]In a dynamic system, relations among variables change in response to signals.

[5]In a continuous system, periodically all equations are solved and state updated.

[6]In a discrete system, a queue of events is maintained and only items related to the queue are solved.

[7]In this case, local versus distributed refers to the hardware needed to run the computational model—a single machine (local) versus multiple machines (distributed).

- *Cost effective*: Creating new technologies, procedures and legislation for data collection is expensive. But using computational modeling you can pre-test these things for efficacy.
- *Faster*: Real time evaluation of existing systems is too time consuming; however, the in a simulation you can "speed up time" enabling rapid development and testing of alternatives.
- *Appropriate*: The world and the simulation are both complex non-linear dynamic systems. Hence the tool matches the requisite complexity and does not overly simplify the state thus affording more accurate predictions and assessments.
- *Flexible*: Response to novel situations requires rapid evaluation of previously unexamined alternatives. This can be done best in a computational framework.
- *Control*: When developing new analysis methods for social systems, computational models can provide a controlled environment for experimentation. Outside of these models, factors of interest often have high multi-co-linearity and are difficult to evaluate.

Computational models can be used for a number of purposes. Among the uses are the following:

- Test bed for new ideas
- Predict impact of technology or policy
- Develop theory
- Determine necessity of a posited mechanism
- Decision making aids
- Forecast future directions
- What if training tools
- Suggest critical experiments
- Suggest critical items for surveys
- Suggest relative impact of different variables (factors)
- Suggest limits to statistical tests for non-linear systems
- Substitute for person, group, tool, etc. in an experiment
- Hypotheses generators.

2 Veridicality and model type

One of the key issues that drive the design, assessment and validation of computational models is their level of veridicality (Carley 1996). On the one hand, many researchers would argue that Occam's razor should apply and all models should follow the KISS principle (keep it simple stupid). Examples of models that take this "proof of concept" approach are Epstein and Axtell's (1997) Sugarscape, many of the Santa-Fe institute models, many of the original "thought based" computational models such as the Cohen et al.'s (1972) garbage can model, the Schelling (1969, 1971) and later, Sakoda's (1971) segregation model, and Kauffman and Weinberger's (1989) NK model. While others argue that to have strong policy relevance and to be able to use the model to make validatable claims, a higher level of veridicality is called for. Examples of such models include Carley et al. (2004, 2006) BioWar and

Silverman et al. (2005) Athena's Prism. In general, higher levels of veridicality allow the model to be used for more types of problems. With a larger amount of code, the code is less likely to be made available. Furthermore, models with high levels of veridicality are more likely able to have aspects of the model validated, but are less likely able to have the model validated in full, since it is less likely that the entire response surface can be generated (Carley 1996). In addition, computational models with a higher level of veridicality are less likely to be built in one of the modeling frameworks available for system dynamic, agent-based, or event-based modeling as the developers will need finer control over the development environment, order or processing, and memory management.

From a human behavioral standpoint, one key issue is how sophisticated or veridical is the Model Human Agent in these computational models. In general, the higher the level of veridicality in the Model Human Agent the fewer agents are typically being modeled. Thus, multi-agent systems that have millions of agents typically have very rudimentary agents formed from only a few rules or equations that reflect very simple cognitive or social activities on the part of the agent. Models with thousands of agents tend to have fairly sophisticated and accurate models of human socio-cultural behavior. Models with less than a dozen agents are more likely to have very sophisticated cognitive and/or task models within the agents. In general, the higher the level of veridicality the fewer the agents and the longer the model processing time for determining the actions of a single agent and the greater the storage needs for a single agent. You can achieve comparable storage and speed constraints as you increase the level of agent veridicality if you reduce the number of agents. In general, the tradeoff that is made is that detailed cognitive processing and task based behavior is often less present in models with thousands of agents where as social and cultural activity and learning by being told is less present in models with a small number of agents. Epstein and Axtell's Sugarscape uses millions of simple agents, Carley's Construct (Carley 1991; Schreiber and Carley 2004) uses thousands of moderately veridical agents, and (Anderson 1993, 1996) and Soar (Newell 1990; Laird et al. 1987) models typically use a handful of highly cognitively sophisticated agents.

Carley and Newell (1994) define three dimensions along which the Model Social Agent varies: cognitive limitations, type of socio-cultural context knowledge, and amount of knowledge about the context. The amount of knowledge that the agent has might impact the speed of the computational model and the quality of the results but not the type of behaviors possible. In contrast, the other two dimensions impact the type of agent behaviors that it should be possible to generate from the computational model. The basic argument is that by placing appropriate limitations on agent cognitive activity and by placing the agents in, and giving them capability to recognize and respond to all classes of knowledge associated with a complete socio-cultural context the agent model becomes the model social agent—a highly veridical avatar of human behavior in all situations. In general, most computational models use agents in less comprehensive environments or without appropriate cognitive limitations and as a result the agents cannot truly generate all human behaviors.

Figure 1 illustrates where many current models fall on these dimensions. This figure is based on Carley and Newell's model (1994) but leaves out real-time interaction

Richness of Situation

Agent Capability

	Non-Social Tasks (Low)	Other Agents	Social Structure	Social Goals	Culture & History (High)
Omniscient (High)	uses tool & language, produces good, single goal	Model of other, takes turns, exchange	Class differences	Organizational goals	Historically situated
Rational	reasons, acquires, learns	Learns from others, education, negotiation	Promotion, social mobility	Group-level competition, cooperation, social cognition	Emergent norms *System Dynamics*
Boundedly Rational	Satisfices, plans tasks, adapts	Builds groups	Altruism, uses networks, spans boundaries *VDT, Garbage Can Model* *Sugarscape, BioWar, Belief models*	Delayed gratification, Group conflict *RTE, TAEMS*	Gate keeping, roles, moral obligation *Organizational Consultant*
Cognitive	Compulsive, lack of awareness, multi-tasking	Group think *Soar, ACT-R*	Automatic response to status cues *Construct*	Group identification	Develop language
Emotional (Low)	Habituation, variable performance	Protesting, trust *Athena Prism*	campaigning	Team bonding	Rituals, advertising **MODEL SOCIAL AGENT**

Fig. 1 Illustrative classification of activities and models

as that seems to be a separable dimension, and it adds in models mentioned herein. In Fig. 1, each computational model (in italics) is placed in the cell furthest to the right and bottom that appears possible for the model. This means that a model in a particular cell, with its current architecture by simply adding mode knowledge should be able to be used to do any and all of the behavior above and to the left of the cell. It is important to note that this breakdown of behaviors is illustrative not definitive. The key point that should be drawn from Fig. 1 is that there is NO computational model today that has a highly veridical Model Social Agent.

In addition to there not being a good candidate for a model social agent, there are a number of limitations faced by computational models in the human social behavioral area at this point in time. One key limitation is that there is no single unifying theory of human social behavior. Rather, there are a panoply of theories some of which lead to contradictory conclusions and all of which have received a limited amount of validation though often only in a specific context. Another key limitation is that there is no single data set of sufficient detail, longitudinal nature, cross-cultural and large enough size to support validation of all aspects of any of the existing models, let alone models that might be developed in the future. The higher the level of veridicality in a computational model, the more "theories" of social behavior are often embedded, at least implicitly, in the model. A third critical limitation is that these models in general are "one-off" models and take substantial work to be "re-initialized" for new scenarios, datasets, or questions.

The methodology of computational modeling is reasonably well understood. Procedures for assessment and validation exist (Yahja 2006). And, there exist specific social behavioral models that have proven to be useful in corporate and policy settings and within those settings have been validated. There are many underlying theories that again are well understood, in isolation; however, the implications for human behavior when sets of these theories are combined in to a single model are not well understood. A key focus for the future is how to make these models easily re-usable.

3 Models, metrics and social networks

Social networks, and more generally dynamic social networks are an area of computational modeling of increasing importance for assessing and understanding human social behavior. The study of social networks is a mature area. As a consequence, social network examples can be used to explore the issue of simplicity and veridicality. To begin this exploration, it is clear that in the area of social networks there is some confusion in the way the term model is used.

In general, the term model typical refers to an abstraction of reality at the system level. In other words, within a model there are numerous variables that can take on a range of values and these variables are linked together in some form of pattern of influence. The term metric typically refers to a measure with key mathematical properties such a having a true 0 point and values having the transitivity property. A variable in a model can be a metric.

In the area of social networks, or network science, these terms are sometimes used interchangeably. For example, some analysts refer to the metric, betweenness, as a social network model of power. From a computational perspective, for large networks i.e., thousands of nodes, many metrics cannot be calculated exactly in a reasonable amount of time and heuristic based computational approaches are used. In this case, the metric is being estimated by a network model. Other analysts refer to the network itself, or the graphic visualization of the empirical data on who talks to whom as a network model of the group; e.g., a network where each link is who interacts with whom among members of a small company might be referred to as the network model of that company. In this case, inherent in the "model" are a set of network properties of the nodes, i.e., their value on a set of metrics. In still other cases, an agent based model in which the agents learn from others to whom they are connected or who alter their connections to others, or a system dynamics or event-based model that uses network metrics as variables are referred to as network models. Statistical estimates of change in linkages or of the likelihood of certain patterns of linkages are also referred to as network models. Hence, when the term network model is used, it behooves the reader or listener to understand how the term model is being used.

From a social behavioral modeling perspective, the area of social networks is of critical importance. There are many reasons for this. First, of all the computational modeling areas, the area of network science is the most developed; i.e., a set of well understood, validated, documented, and meaningful metrics, toolkits, well understood procedures for data collection and analysis, easily linkable to other types of models. Network tools, and so "network models" as the metrics, are very mature technologies. Many metrics have existed since the 1950's and high speed, large scale versions exist for the most common, usable, and meaningful metrics. Many of the tools such as UCINET, Pajek, *ORA (Carley et al. 2008) have existed in some form for over ten years. Second, networks constrain and enable behavior to the extent that understanding the network in a group is critical to identifying key actors and supports course of action analysis. Third, network metrics and models have been used with demonstrable success to support real world decisions in areas such as corporate re-organization, counter-terrorism, law-enforcement, and social policy. Fourth, for each organization or group it is a matter of minutes to create a "network model" in the following sense: any set of nodes for which the links are known can be imported to a network analysis tool and the metrics calculated, the system visualized and assessments of the impact of node or link removal or addition rapidly assessed. Fifth, there is a recognizable curriculum that individuals need to know to be competent in the social network area (Tindall and Malinick 2008). In general, the most successful cases have been when a meta-network approach was taken (see below). Unfortunately, the currently popularity of network science has lead to a swell in the number of people claiming to do work in this area with most of the "practitioners" having little background. As a result, recently, there has been a tremendous amount of re-invention and re-discovery.

Networks constrain and enable behavior. In social networks, people to people, who interacts with whom impacts what information is learned and transmitted, the flow of diseases and the flow of money. However, networks impact more than just people. In many situations, it is important to think about the dynamics of meta-networks that connect the who (e.g., people and organizations), what (e.g., tasks, activities and events), where (locations either at the general level—a building or a specific latitude and longitude), why (e.g., attitudes, beliefs, norms, goals) and how (e.g., resources and expertise needed to accomplish the what and held by the who). However, most network analysis tools operate only on single networks and/or utilize standard social network metrics without re-validating or seeing if the metrics make sense for the specific type of network being examined. A key exception to this limitation is the *ORA tool which is designed to handle multi-mode multi-link dynamic networks; i.e., meta-networks, and thus intelligently contains a superset of the minimal features in the traditional software.

The importance of the meta-network should not be underestimated. If only a single social network is used to measure and inform social behavior the results are often incomplete. For example, innovators may be those on the fringe and not connected to many others, or they may be those directly and strongly connected to two or more disparate groups. For example, if the 911 hijacker network had been put together prior to 911 and the individuals detained who had the highest centrality (degree, betweenness, eigenvector or closeness) that may not have been sufficient to stop the activity; e.g., they could have been replaced, it wouldn't have stopped the financing, and so on. Whereas, combining the information contained in a social network with the information contained in relevant other networks enables: (a) better understanding and prediction of future behavior, (b) ability to predict change in or evolve the social network, (c) ability to infer missing links and possibly missing nodes in the social network, and (d) ability to design and evaluate organizational structures. Future work should concentrate on meta-network applications such as geo-temporal network analysis, or linking social and belief networks.

Network models are often touted as "data greedy" because when the model is of a group and the nodes are people the most accurate results for the metrics require knowing for each pair of individuals whether or not they are connected. However, a network science approach can be used at any level, individual, group, state, and interstate. The nodes can be practically anything and the links can represent any number of types of relations between the nodes. The links can vary in strength, directionality, and confidence. The data can even be completely hypothetical. This being said, most traditional network tools can only handle one to two types of nodes at a time, one to two types of relations, and most metrics only operate on binary data where the links have been reduced to present or not; whereas, *ORA is the only software tool that can handle dynamic multi-mode, multi-link data.

The robustness of network metrics to missing or incorrect data is only just beginning to be understood. Some studies show that, if the underlying network is random then with even 15% missing data the rank order of the nodes on different metrics can change significantly and the user should focus on the likelihood of the nodes being in the top/bottom 10% on that metric. For other network topologies, such as cellular or scale free, the metrics may be more or less robust. However, there is no comprehensive list of topologies and models of the impact of missing data on the confidence interval around the metric are not worked out. The statistical underpinning of network analysis is an area that is getting increased attention.

Due to the maturity of network modeling, models that evolve, predict, reason about actor behavior, beliefs, and so on as a function of the network position of the actors, or the network constrained activity and communication among the actors have strong utility for decision making. Moreover, if network ties are viewed as probabilistic then change detection techniques (McCulloh and Carley 2008a) can be linked with the network models to asses when the multi-agent network models are generating true shifts in social behavior. Further, multi-agent network simulations that employ and evolution of the link probabilities for the networks provide a principled and systematic approach for examining social change (McCulloh and Carley 2008b).

4 Analysis and validation

Analysis and validation of computational models in the area of human social behavior is challenging. There is no simple text book approach that can be followed in all situations. Indeed, the level and types of analysis and validation conducted depend on the purpose of the model. That being said, there are a number of principles that if followed lead to the long run success of a computational modeling effort in the human social behavior area (Maxwell and Carley in press). These principles are:

- Understand the tradeoffs in the system you are modeling
- Clearly define the purpose of the simulation and the associated users
  - If the purpose changes, revisit the assumptions
  - Decide whether model validation is warranted given the purpose
- Use good modeling practices
  - Make sure the research or analysis question is highly focused

  - If you are to validate the model, first identify the mapping between measurable data and simulated variables
  - Specify, in detail, the desired output measures
  - Define how and where uncertainty plays a role in the model and the results
  - Document all assumptions
  - Document all modeling risks
- Clearly specify all input, output and control variables
  - For human social behavioral models—agents and the environment are key
- Clearly specify all agent behaviors
  - Identify the relations among agents
  - Define how change occurs through agents and their connections
- Conduct calibration, verification and validation exercises as warranted by the model's purpose
- Assess model results by running well structured virtual experiments
  - Use good experimental design
  - Conduct rigorous statistical analysis
- Clearly present results and discuss limitations
  - Consider the audience
  - Consider the model's purpose.

It is critical to note that training in the computational modeling area, particularly for social behavioral models tends to be done using a series of case examples i.e., explorations of existing models, and then a series of programming tasks. Key issues of experimental design, analysis and validation are often overlooked or treated using a strictly engineering approach which is not appropriate for social behavioral models. This being said, it difficult to teach analysis and validation for these models as the type and level of analysis and validation must be sufficient unto the needs of the model and the way it will be used. Purpose should guide not only the design of the computational model, but the way its results are analyzed and whether or not, and if so to what level, it is validated.

Proper and careful analysis is particularly vital when the model is stochastic. In this case Monte Carlo techniques and virtual experiments are needed for assessment. A Monte Carlo approach relies on replicating the "runs" of a computational model with different values for those variables set by a random number generator. A virtual experiment is an experiment run using a computational model; standard experimental design procedures should be applied. The experiment is "virtual" as it is being done in a virtual or simulated world. Unfortunately, all too often, those using computational models to examine human social behavior just run the model a few times, generate a few figures, do not lay out an experimental design for the virtual experiment and discuss the results from a general perspective with little attention to sensitivity of the results to parameter changes, or alternative explanations.

In contrast, those analyzing a model's result should take a response surface analysis approach (Meyers and Montgomery 2002), complete with the statistical assessment of the relation of inputs to outputs of interest. It is critical to note that for human social behavioral models the models are sufficiently complex and the number of variables sufficiently large that a complete response surface analysis is not feasible in a reasonable amount of time with reasonable storage constraints. Hence, the analyst typically needs to focus in on a question to be answered and then design the virtual experiment to collect data using the computational model to answer that question. Then the virtual experiment is run and data collected. For each outcome measures of interest (typically thought of as dependent variables), often the best fit lowest order polynomial is an appropriate reduced form description of the model at that stage of processing. Then this polynomial is used to define the hypotheses and conclusions emerging from the computational model. Note, on empirical data there is a careful tradeoff between model complexity and factor significance; however, for simulated data, all factors should be significant if the model was run a sufficient number of times.

It is important to note that the role of statistics in the assessment of computational models is somewhat different than it is in the assessment of empirical data collected in the field. With computational models, if there is an observable relation then it will be significant statistically if enough simulations are run. Hence, significance is generally only used to determine if sufficient replications have been completed. This means that the focus in interpreting a statistical model derived from computational model results is on the strength of the coefficients and not on their significance. It is important to note that this is true of test statistics that are inversely proportional to the sample size. There are methods that look at estimating the distribution of a measure from two distributions and evaluating their overlap. The variance of the mean of a process is affected by sample size, however, the distribution of process observations themselves are not affected by sample size.

It is also important to note that the theory of analysis, and the types of robust experimental designs that are commonly taught in engineering simulation classes, is generally not sufficient for human social behavioral models. There are several reasons including that human social behavioral models, as compared to the engineering models, typically have a much higher number of variables, have high covariance among those variables, have discontinuities in many variables and so violate the continuity assumption, have interaction effects among variables, and have temporal variations in the relation of variables to each other due to learning and social change processes thereby violating the consistency assumption. Thus, human social behavioral models do not meet the necessary mathematical assumptions for traditional engineering analysis. Human social behavioral models mathematically present a much more challenging problem.

From a validation standpoint, the first issue is: should the model be validated? The simpler intellective models are often developed for the express purpose of telling a story or making a point. They are not meant for developing policy, guiding purchasing decisions, and designing experiments. As such, they may not need to be validated. Moreover, such models often use variables that are not unequivocally measurable and so cannot be validated. The purpose of the model should drive validation (Burton and Obel 1995).

For the more veridical models, and for models used in policy contexts validation, or at least calibration, is more of an issue. In this case, the issue is what constitutes validation and how much validation is needed. A typical approach is to build a model based on some theory, instantiate it with data from a real situation, generate a series of predictions about variable "y" and see if they are confirmed. If the model passes this test, then it is often "tested" further by conducting a new virtual experiment that includes some feature of reality not previously examined or included in the model, to see if the same results about variable "*y*" are produced. If the prior results hold, the underlying theory is strengthened. If the results do not hold, a new theory is called for. However, the model is still "valid" under the original conditions, as validity was based on the match to the original data.

In other words, validation as practiced for engineering models may be inappropriate at worst and impossible at best (given the model complexity, large number of variables and infeasibility of constructing a complete response surface). In general, the data needed to completely validate highly veridical models often does not exist in any clean single case scenario. Rather, the researcher doing the validation needs to collect disparate data from a wide variety of sources, fuse the data together, and then use it collectively validate the model. This may take as much time and resources as it takes to build the model, run the virtual experiments, and analyze the results. Validation to historic events, often leads to the model being over-fitted to a specific situation and not flexible enough to enable understanding the space of possibilities in alternative and future scenarios. Face-validation using external subject matter experts may be sufficient in many cases. Other alternatives to traditional validation as practiced on engineering models are docking (Burton 1995; Axtell et al. 1996) and validation in parts. A new science of validation for human social behavioral models is needed.

5 Conclusion

Computational models have been, and will continue to be usefully applied to solve real world issues, explore theoretical problems, and suggest possible futures. The area of human social behavior presents special challenges to the computational modeling world. Essentially, humans, particularly groups of humans, are more complex than physical or engineered systems. Human social behavior is decidedly non-ergodic. Technological and political change result in severe discontinuities. And, for any social behavioral outcome is the result of multiple causes none of which are necessary or sufficient in isolation, but collectively lead to various activities. In essence, computational models of human social behavior are operating in a quantum mechanics world where the individual particles learn and the multiple conflicting sets of rules of behavior evolve over time with large discontinuities. The level of complexity of, the lack of ergodicity in, and the discontinuities in the phenomena being modeled is unprecedented to the modeling community.

The high-complexity, non-ergodic, discontinuous, multi-causal nature of human social behavior means that many traditional approaches to analyzing and validating computational models will not work. It also means, that it is difficult and often infeasible, to capture sufficient real-world data for complete testing. Further, it means that the complete response surface of the computational model typically cannot be generated. And, it means that there is not a single unifying theory of human social behavior but a panoply of theories.

How then can we tell that we are making progress and that the computational model is providing value? First, the goal should be consistency with the real world, not validation. Consistency can be achieve by engaging in validation by parts (input is real, some output streams match, individual theories incorporated have been validated under some conditions) and by conditional validation (in a certain range of circumstances the values match the real world). Models built that rely on separately validated theories are more consistent with the real world. Second, employ theoretical integration. Computational models need to be built by integrating multiple theories, thus covering a wider range of human experience. By building a model consistent with multiple theories, the overall strength of, and confidence in the overall model is increased. Third, employ triangulation. For real world problems where policy decisions may be made on the basis of models, it is generally valuable to employ multiple models built at different levels of granularity and using different modeling approaches (such as agent based and system dynamic). When multiple models converge on a similar result, those results are more robust and a better indicator of what is likely to happen in the real world. Moreover, by using multiple models, you will get a more comprehensive understanding of the various factors that can contribute to the outcome, various stages that are likely to occur along the way, and more insight in to unintended consequences of various interventions or courses of action.

Computational modeling provides the analyst with a new symbol system for understanding, predicting, reasoning, and explain human social behavior. While tremendous strides have been made since the very first model, Cyert and March's (1963) A Behavioral Theory of the Firm; there is still a long way to go. Great advances can be made, by comparing and contrasting multiple models, and allowing them to inter-operate. The key is not to use the models to compete, but to use them in a more interoperable synergistic holistic fashion.

Acknowledgements This work was supported in part by the Office of Naval Research (N00014-06-1-0104), the National Science Foundation (SES-0452487 and IGERT 9972762), and the Air Force Office of Sponsored Research (MURI: Cultural Modeling of the Adversary, 600322) for research in the area of dynamic network analysis. Additional support was provided by CASOS—the center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the National Science Foundation, the Air Force Office of Sponsored Research, the Army Research Institute or the U.S. government. The author would like to thank Terrill Frantz and Ian McCulloh for extensive and helpful feedback on earlier drafts.

References

Anderson JR (1993) Rules of the mind. Erlbaum, Hillsdale

Anderson JR (1996) ACT: a simple theory of complex cognition. Am Psychol 51:355–365

Axtell R, Axelrod R, Epstein JM, Cohen MD (1996) Aligning simulation models: a case study and results. Comput Math Organ Theory 1(2):123–141

Burton RM, Obel B (1995) The validity of computational models in organization science: from model realism to purpose of the model. Comput Math Organ Theory 1(1):57–71

Burton RM (1995) Validation and docking: an overview, summary and challenge. In: Simulating organizations: computational models of institutions and groups. MIT Press, Cambridge, pp 215–228

Carley KM (1996) Validating computational models. Working Paper

Carley KM, Columbus D, DeReno M, Reminga J, Moon I (2008) ORA user's guide 2008. Carnegie Mellon University, School of Computer Science, Institute for Software Research, Technical Report, CMU-ISR-08-125

Carley KM, Fridsma DB, Casman E et al (2006) BioWar: scalable agent-based model of bioattacks. IEEE Trans Syst Man Cybern A 36:252–65

Carley KM, Altman N, Kaminsky B, Nave D, Yahja A (2004) BioWar: a city-scale multi-agent network model of weaponized biological attacks. CASOS Technical Report: CMU-ISRI-04-101. Carnegie Mellon University, Pittsburgh, PA

Carley KM, Newell A (1994) The nature of the social agent. J Math Sociol 19(4):221–262

Carley KM (1991) A theory of group stability. Am Sociol Rev 56(3):331–354

Cohen MD, March JG, Olsen JP (1972) A garbage can model of organizational choice. Adm Sci Q 17(1):1–25

Cyert RM, March JG (1963) A behavioral theory of the firm. Prentice-Hall, Englewood Cliffs

Epstein J, Axtell R (1997) Growing artificial societies. MIT, Boston

Harrison JR, Lin Z, Carroll GR, Carley KM (2007) Simulation modeling in organizational and management research. Acad Manag Rev 32:1229–1245

Kauffman SA, Weinberger ED (1989) The N-K model of rugged fitness landscapes and its application to maturation of the immune response. J Theor Biol 141:211–245

Laird JE, Newell A, Rosenbloom PS (1987) Soar: an architecture for general intelligence. Artif Intell 33:1–64

Law AM (2007) Simulation modeling & analysis, 4th edn. McGraw-Hill, New York

Maxwell D, Carley KM (in press) Principles for effectively representing heterogeneous populations in multi-agent simulations. In: Tolk A (ed) Complex systems in knowledge based environments. Springer, Berlin

McCulloh IA, Carley KM (2008a) Social network change and detection. Institute for Software Research, Technical Report, CMU-ISR-08-116

McCulloh IA, Carley KM (2008b) Detecting change in human social behavior simulation. Institute for Software Research, Technical Report, CMU-ISR-08-135

Ross SM (2006) Simulation, 4th edn. Elsevier, New York

Meyers RH, Montgomery DC (2002) Response surface methodology: process and product optimization using designed experiments. Wiley, New York

Newell A (1990) Unified theories of cognition. Harvard University Press, Cambridge

Tindall DB, Malinick TE (eds) (2008) Teaching about social networks. American Sociological Association, New York

Sakoda JM (1971) The checkerboard model of social interaction. J Math Sociol 1:119–132

Schelling T (1969) Models of segregation. Am Econ Rev 59:488–493

Schelling T (1971) Dynamic models of segregation. J Math Sociol 1:143–186

Silverman B, Rees RL, Toth JA, Cornwell J, O'Brien K, Johns M, Caplan M (2005) Athena's prism—a diplomatic strategy role playing simulation for generating ideas and exploring alternatives. Departmental Papers (ESE)

Schreiber C, Carley KM (2004) Construct—a multi-agent network model for the co-evolution of agents and socio-cultural environments. Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-04-109

Yahja A (2006) WIZER: a tool for validation of social simulations. Ph.D. Thesis, School of Computer Science, Carnegie Mellon University

Zacharias GL, MacMillan J, Van Hemel SB (eds) (2008) Behavioral modeling and simulation: from individuals to societies. Committee on organizational modeling: from individuals to societies. Academies Press, Washington

Kathleen M. Carley received her Ph.D. from Harvard in Sociology and is a Professor in the Institute for Software Research in the School of Computer Science at Carnegie Mellon University and the director of the center Computational Analysis of Social and Organizational Systems (CASOS). Her research combines cognitive science, social networks and computer science to address complex social and organizational problems in areas such as dynamic network analysis, computational social and organization theory, command and control, counter-terrorism and counter-narcotics. She and her lab have developed infrastructure tools for extracting networks from texts (AutoMap), for analyzing large scale dynamic networks (*ORA) and various multi-agent simulation systems including ones for assessing the evolution of networks (DyNet), the impact of interventions on belief and information diffusion (Construct), and a city-scale model for understanding the spread of disease and illness due to natural epidemics, chemical spills, and weaponized biological attacks (BioWar).