32

Social-Behavioral Simulation: Key Challenges

Kathleen M. Carley

Institute of Software Research, School of Computer Science and Engineering and Public Policy, Carnegie Institute of Technology, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Introduction

Social-behavioral simulation is often turned to as a way of reasoning about the human condition and thinking through what might happen under various interventions. Potential advantages include overcoming human decision bias, thinking through a large number of alternative situations, and reasoning about context and over periods of time that are too complex for the human unaided by the computation to reason about. While these advantages exist, simulation of complex social-behavioral systems is not routine. Nor are the simulation models, necessarily, accurate reflections of reality. There are more creditable models where the logics built into the model reflect established theory or empirical regularities. Yet despite best efforts, the logics in these models may reflect human biases. Further, in many simple models and in those where no theory exists, the model may actually reflect simply the potentially nonsensicle opinions of the creators.

Agent-based (e.g. Bonabeau 2002; Davidsson 2002; van Dam et al. 2012), event history (Box-Steffensmeier and Jones 2004), Petri net (Tabak and Levis 1985; Murata 1989), and system dynamic models (Sterman 2001; Mohaghegh and Mosleh 2009) are generally used as the simulation frameworks when modeling complex systems in the social sciences. No one methodology is proving adequate, so the field is moving toward hybrid modeling and a system of systems approach using interoperable models. These models make explicit the gaps in underlying theories; support theory comparison, integration, and development; and enable users to create a framework in which they can rapidly reason about alternative explanations (useful forensically) or alternative courses of action (useful for planning) (Gilbert and Troitzsch 2005).

Social-Behavioral Modeling for Complex Systems, First Edition.
Edited by Paul K. Davis, Angela O'Mahony, and Jonathan Pfautz.

Companion website: www.wiley.com/go/Davis_Social-Behavioralmodeling

Complex systems are generally characterized as systems composed of many interacting heterogeneous components, such that the behavior of the system is nonlinear, is not predictable from the sum of the parts, and often exhibits self-organization and such that actions at one level of granularity lead to emergence at another level of granularity (Bar-Yam 2002; Miller and Page 2009). Complex systems are difficult to understand and explain. It is difficult to predict the effects of actions on them and to predict when new events will occur within them. Social-behavioral systems are classic examples of complex systems. However, they also represent a special case of complexity because humans, who can learn, are key elements of social-behavioral systems.

The field of simulation and the validation of simulations have their roots in shop floor management and operations research. Most of the approaches to model development, the theory of validation, and the expectations for how models operate are based on an operations research conception predicated on models of physical systems. Key assumptions underlying that body of work include stationarity of process and the corollary that components do not learn. These assumptions are violated in social-behavioral systems. Consequently, the traditional science of validation does not apply, and a new science is needed (National Research Council 2006, 2008). Developing such a science will depend on meeting a number of challenges both in terms of how to communicate social-behavioral models and their results and how to develop such models to create reusable and scalable technologies.

Key Communication Challenges

Independent of the scientific challenges, there are a number of communication challenges faced by those interested in developing and using simulation to address complex social-behavioral systems. These challenges take the forms of cognitive biases that humans have when they try to understand the results of models and the models themselves. Two of these will be considered:

Consumer bias: Human consumers of simulations are themselves humans and so have a tendency to assume that they know how humans behave. In contrast, many humans have placed physics on a pedestal, have deep-seated math anxiety, and assume that mechanical and engineered systems are hard to understand. This creates a situation where the consumers of social-behavioral models think they should be able to understand and agree with models of human activities, but do not make the same assumption for physical system models. A consequence is that social-behavioral models are held to a higher standard vis-à-vis communication than other models.

Storytelling bias: Humans have a tendency to think they understand how something works when they can tell a story about it. Storytelling is a cognitive crutch by which we keep track of the order in which things occur and the relations among things. The more complex the story, the harder it is to recall.

Simple models such as the model of racial segregation (Schelling 1969) or the garbage can model (Cohen et al. 1972) or the NK model (Kauffman and Weinberger 1989) lend themselves to stories that are easy to retell even in contexts distinct from that used to justify the original model. Complex models, such as EpiSims (Mniszewski et al. 2008) or BioWar (Carley et al. 2006), do not lend themselves to such simple stories. Humans are also more likely to view as accurate models they think they understand. A result is that consumers of the NK model and the garbage can model think they understand how the models work and are more likely to view them as credible and accurate when, in fact, they really do not understand the scope conditions and the models themselves cannot be validated. In contrast, more complex models, even if they are validated, are harder to tell stories about, are less likely to be viewed as understandable, and so are less likely to be viewed as credible.

Key Scientific Challenges

Independent of the communication challenges, there are a number of scientific challenges faced by those interested in developing and using simulation to address complex social-behavioral systems. Overcoming these challenges is critical for creating scalable, reusable social-behavioral simulation systems and for reducing the high cost of developing credible simulation systems. These key challenges center around speed of development, reuse, and approaches to validation. These challenges are itemized in Table 32.1. They

Table 32.1 Summary of scientific challenges for social-behavioral modeling.

Challenge	Meaning	Example
Reuse across context	Explains/forecasts how phenomena x evolve in sociocultural context y_1 at time n_1 and in sociocultural context y_2 at time n_2	Use model to explain terrorist recruitment for IRA before 2000 in the United Kingdom and then for ISIS recruitment in Syria, 2010–2017
Reuse across level	Explains/forecasts behavior of actor and group of those actors	Use model with data set 1 to explain an individual's anger toward election fraud for a candidate and on data set 2 to explain population-level anger toward election fraud for a national election
Reuse across time	Explains/forecasts behavior of entities at time period n_1 and at time period n_2	Use model with data set 1 to predict state stability in 1910 and on data set 2 to predict state stability in 2010

(Continued)

Table 32.1 (Continued)

Challenge	Meaning	Example
Multi-scale	Explains/forecasts the behavior or change in opinions/beliefs of x, and sets of x in the same simulation run	Use model to forecast behavior of individuals moving between firms, and/or organizations moving between conglomerates or merged firms, using the same set of data and during the same set of simulation runs
Rapid instantiation	Shortens time to prepare and import data to model and adjust model parameters	Construct and use model repeatedly on different data sets, and observe a decrease in the time to prepare and ingest data
Multi-group	Generates model that remains valid as actor granularity is changed	Use model on data set 1 where agent is a human to predict consensus, on data set 2 where agent is an organization to predict consensus, and on data set 3 where agent is a nation state to predict consensus
Multi-temporal	Generates model that remains valid as time steps change	Use model to explain escalation of commitment of individuals to groups' goals using data set 1 where individual commitment is measured each day and data set 2 where individual commitment is measured each year
Multi-spatial	Generates model that remains valid as spatial region is changed	Use model to explain the formation of crime zones using data set 1 where the zones are measured at the city block level and data set 2 where the zones are measured at the city level
Interoperability	Allows model to be used with other models as part of a modeling workflow	Use model in a system of systems, demonstrating the ability to ingest data from and export data to other models
Data fusion	Allows model to use data at different temporal, spatial, and group levels with different levels of fidelity	Use model to ingest and use for instantiation, and then again for validation, with two sets of data that vary in granularity of actor, time period, or spatial region or have other collection inconsistencies
Response surface	Generates boundary conditions and the full response surface	Conduct virtual experiments to estimate the response surface via statistical calculations

overlap with the challenges discussed in overview chapter of this volume (Davis and O'Mahony 2019). Rather than providing a detailed discussion of each challenge, three issues that apply to multiple challenges will be discussed:

Model granularity: Social-behavioral systems are often reasoned about, explained, and measured at different levels of granularity. Among the levels of granularity of interest are time (minute, day, week, year), space (block, city, county, country), agent group (neuron, person, decision-making unit, organization, nation), and knowledge group (word, concept, topic, theme). Typically models that are accurate or designed for one level of granularity do not work at other levels of granularity. For example, models that are accurate at the cognitive and task level do not scale to the community or even large group level, whereas models that are informative about populations are inaccurate at the cognitive level. One reason for this is that the logics of interaction may change fundamentally with scale. Basic research on how logics change with scale is needed. Another reason for this is that the aggregation processes and disaggregation processes needed as one moves up and down the levels of granularity are not known. Two advances in this area are dynamic network modeling and social cognition modeling (Morgan et al. 2017). Dynamic network models provide a mechanism whereby the pattern of interaction among the components can be used to affect the decision logic or behavior of the component whose internal logic can also be represented as a network (networks where the nodes are networks). Social cognition is a set of mechanisms whereby individual humans make sense of the world by reasoning not about every individual but about collectives in the same way they reason about individuals. This includes social cognitive mechanisms such as recognizing and responding to a generalized other, inferring individual traits from perceptions about a group, generalizing group traits from activities of an individual, and recalling average rather than actual behavior. Focusing on these types of logics that support connecting one level of a model with another is critical for improved model accuracy and reuse.

Level of validation: Historically, the science of validation has treated validation as a single process. A model was either valid or not, and there was a single approach to validation. In contrast, for social-behavioral simulations it is generally recognized that there are multiple levels or types of validation and that the level of relevance and the extent of validation depend on the model's purpose (Burton 2003; National Research Council 2008). There are various ways of characterizing the various aspects of validation. For example, Davis and O'Mahony (2019) describe the components as being description, explanation, postdiction, coarse exploration, and prediction. An alternative approach is to think in terms of the empirical precision at which the model is validated and/or matched to the predictions of another model (i.e. docked) – including similarity in the pattern of results, the distribution, the range, or the exact value (Axtell et al. 1996). Still more detailed schemes exist that take into account the number of features on which to validate, the level of empirical match, and the extent of revalidation (Yahja and Carley 2005). To move forward with a science of validation for social-behavioral simulations, a validation framework is needed that accounts both for the various components of validation, the purpose of the model, the level of empirical match, and the nature of the underlying data to support match. See also another paper in this volume discussing validation (Grace et al. 2019).

Identifying the right level and type of validation for a model is further complicated in a reuse model–expansion situation. The crux of the problem has to do with interaction among subtheories. Most social-behavioral simulations are at their core multi-theoretical. Typically the way in which these theories linked together is not known. Model development often means suggesting a theory for these gaps. Validation then becomes a process of creating an empirically grounded mega-theory (e.g. see Schreiber and Carley 2004). This is counter to another assumption of traditional validation science, which assumes that the theory is known and well specified. This is one of the gaps through which human biases creep into these models. In a related vein, typically social-behavioral simulations are built in a building block fashion. That is, the basic model is built, possibly validated, and used to explore an issue. Then the model is reused, generally expanded on, and the process repeats. The issue here is that as models are expanded, the earlier parts of the model may no longer behave as they did previously. This is particularly true in nonlinear complex systems. This means that the results of earlier validation exercises may no longer hold, hence, leading to a situation where the model needs to be revalidated for each reuse/expansion situation.

Validation vs. tuning: Historically, the science of validation for simulation has been based around two assumptions: stationarity of process and measurability. Stationarity of process means that you can collect a data set from time period *N* or machine *y* and tune the model until you get a high level of predictive accuracy (e.g. 98%). Tuning is done by changing internal processes in the model. The resultant model is considered validated and can then be used to predict the behavior of the system at other time periods or the behavior of objects from comparable machines. Social-behavioral system violates this assumption. Consequently, tuning the model to a historic case study is actually overfitting. The resultant model is then fragile and unlikely to apply in other situations. Consider, for example, the artificial intelligence/machine learning (AI/ML) models, particularly those that require training sets. Such models typically use data from a single context/time period and split it into two sets. The model is built based on the first part of the data and then applied to the second part. Accuracy of such models can become quite high (e.g. approaching 98%). However, if moved into the wild and used in other contexts, they often will not work, and/or the accuracy can decrease dramatically due to the nonstationarity of process.

Simply tuning a model to two historical cases is not a sufficient solution. Recall that as a model is reused from one situation to the next, new modules are often added. A possible consequence is that if a model is tuned to one historical case and then reused, expanded, and tuned to a second case, the model may no longer be tuned to the first case. In complex social-behavioral simulations, the relationship between model components and validation is so complicated that special artificial intelligence (AI) tools may be needed to track and predict which parts of the model will become invalidated as the model is expanded and retuned (Yahja and Carley 2005).

Another assumption of the traditional science of validation is measurability; i.e. it is possible to delineate all aspects of the physical system and measure its properties. Social-behavioral systems violate this assumption in part because we often cannot know what people are thinking, particularly in situations where people themselves are unable to forecast their own behavior – such as life-threatening situations. For these and other reasons, such as measurement imprecision, and propagating measurement errors, predictive accuracy (i.e. predicting future behavior) of social-behavioral models is often much lower than those for physical systems. While ML/AI models often claim they are doing prediction – as was noted – the high accuracy they achieve is often using data from the same time period and context as they were tuned (i.e. trained) on and rarely have this high a level of prediction for future events.

For such reasons, simulation often plays a different role in the social-behavioral context than in the physical system context. Its value is in forecasting, not prediction. In other words, social-behavioral simulations are generally the most useful for describing the space of future possibilities and the relative likelihood of particular outcomes. Trying to express this difference to policymakers and those not steeped in simulation is difficult for a variety of reasons (see also Davis and O'Mahony 2019). First is the terminology. As was alluded to the term, prediction means, depending on the user, a priori predicting a specific event in the future (aka point prediction), explaining behavior in one part of the data given the other part of the data, and a priori estimating the relative likelihood of different outcomes (aka probabilistic forecast). Similarly the term prediction also has been used to mean extrapolation (as is often done with regression models), suggesting possibilities (aka qualitative forecast), equivalent to prediction, and a priori estimating the relative likelihood of different outcomes (aka probabilistic forecast or prediction). Using this terminology, the difference can be described as simulation of physical systems as being more capable of point prediction and simulation of social-behavioral systems as being more capable of probabilistic forecasting or qualitative forecasting.

A second difficulty is that the differences are in part a matter of degree. That is, the vast majority of physical system simulations support point prediction. Certainly the science of validation is designed for this type of prediction. In contrast, the vast majority of social-behavioral simulations support either probabilistic forecasting or qualitative forecasting. For those few social-behavioral simulations that can do point prediction, they do so only under specific scope conditions. At this point a science of validation that supports forecasting does not exist in its entirety. Further, the tactics and techniques needed to determine whether the scope conditions and the logics in social-behavioral simulations are such that they will support point prediction are understood more in the abstract than in practice. For example, stationarity of process is needed, but measuring it is often an open challenge.

Toward a New Science of Validation

A consequence of these considerations is that a new science of validation is needed for social-behavioral modeling. At the heart of such a science is the practice of validation in parts and incremental validation, instead of tuning. Validation in parts means that parts of the model (the input, the outputs, and the internal processes) are validated separately, often on different data and sometimes by different teams. Three factors are needed: (i) inputs of the model have been shown to have at least the same distributional properties as real world. In the limit the input data for the model are exactly real-world data streams and/or empirically derived from real-world sensors. Under the philosophy that the model should be able to do the task it seeks to explain, the model must take as input the exact data that the real-world social-behavioral system would. Models that fit these criteria could, in principle, be substitute for the social-behavioral system in other contexts (much as AI personal shoppers could substitute for a human personal shopper). (ii) The distribution of model outputs contains, as special cases, historic examples. However, this does not imply that the mean prediction of the model should agree with the historical data point. Indeed if the model contains discontinuities in some of the variables, the mean prediction of the model might be a situation that logically cannot exist. Logically, one would expect that as the number of historical cases increases, the distribution of historical data points should come to reflect the plausibility distribution from the simulation. (iii) Internal processes in the model have been shown to match at some level to processes observed in the real world. This match might be very qualitative as when confirming the face validity of a model by claiming that each element has an analog in the real world. Or this match might be strongly empirical and precise using exact values in equations that match the equations used to describe the real world.

Incremental validation means that the model is validated in steps, and at each step the number of aspects that is validated and the level at which the model is validated are increased. Incremental validation is an approach meant to overcome the high level of complexity and the extensive number of variables and processes needed to explain social behavior. For incremental validation a set of features to be validated are identified. This set is often modularized by granularity level, or theory being operationalized, or data context. The model is built and it is validated against the base features. Then new modules/theories are added, and the model is revalidated against the old features and also validated against the new set of features. A key advantage of this approach is that it allows theory to be developed to fill in the gaps linking different levels of granularity. Another advantage is that this approach does not require a universal perfect data set, but enables the simulationists and the consumers to make use of large diverse sets of data from multiple sources and contexts that may or may not be fusible.

A science of validation appropriate for social-behavioral simulation may be well served by the development of a model social agent (Carley and Newell 1994). The model social agent is a conceptual framework that defines the elements of a social agent and what behavior should emerge from any model containing agents composed of those elements. A first attempt in this direction was forwarded by Carley and Newell (1994). While not sufficient, the suggested framework shows the power of having such a model in providing guidance for what models should and should not be able to do and so what type of and level of validation might be needed.

Conclusion

The state of the art in simulation of complex social-behavioral systems has advanced in the past several decades. Gone are the days of all models being rational actor models. Gone are the days of model assessment being done by generating a single simulation run, viewing the results, and stating that “it looks right.” That being said, the field needs some serious advances to become more mainstream. Overcoming some of the challenges described herein will help.

Nevertheless, it is important to recognize that these are simulations of complex social behavior. It is unreasonable to expect rapid development, high levels of accuracy, and high levels of reuse when the subject being modeled is still so poorly understood. The *laws* governing human interaction, learning, social engagement, and cognition are still being discovered. Each new discovery increases the value of and the ease of building simulations and validating them. In that sense, advancing the field of simulation is codependent with advancing our understanding of human social behavior.

It is also important to recognize that humans and the groups or communities they form are an evolving system. This suggests that the simulations themselves will need to evolve. With complex and evolving systems, ensemble approaches are often of value. For example, major advances in weather forecasting occurred once sets of models were used in an ensemble fashion to provide estimates. Standardization in output formats and levels of granularity for outputs enabled cross-fertilization. Such an approach may do well in the social-behavioral context as well.

A final approach is that social-behavioral modeling is *big science*. Strong, reusable systems require substantial resources to be developed. Large teams are needed to develop and fuse the data needed for model development and validation, for building the infrastructure and tool chains for using models in an interoperable fashion, to run virtual experiments using the model to examine results, and to describe and document the model. The infrastructure around the model needed to ensure its use requires more people and more time than the model design and development itself. If we continue to think of social-behavioral simulation as being comparable with running a statistical analysis of a moderate data set, progress in this area will be slow.

References

Axtell, R., Axelrod, R., Epstein, J.M., and Cohen, M.D. (1996). Aligning simulation models: a case study and results. *Computational and Mathematical Organization Theory* 1 (2): 123–142.

Bar-Yam, Y. (2002). General features of complex systems. In: *Encyclopedia of Life Support Systems*. Oxford: EOLSS UNESCO Publishers. Retrieved 16 September 2014.

Bonabeau, E. (2002). Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America* 99 (suppl 3): 7280–7287.

Box-Steffensmeier, J.M. and Jones, B.S. (2004). *Event History Modeling: A Guide for Social Scientists*. Cambridge University Press.

Burton, R.M. (2003). Computational laboratories for organization science: questions, validity and docking. *Computational and Mathematical Organization Theory* 9 (2): 91–108.

Carley, K.M. and Newell, A. (1994). The nature of the social agent. *Journal of Mathematical Sociology* 19 (4): 221–262.

Carley, K.M., Fridsma, D.B., Casman, E. et al. (2006). BioWar: scalable agent-based model of bioattacks. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans* 36 (2): 252–265.

Cohen, M.D., March, J.G., and Olsen, J.P. (1972). A garbage can model of organizational choice. *Administrative Science Quarterly* 17 (1): 1–25.

van Dam, K.H., Nikolic, I., and Lukszo, Z. (eds.) (2012). *Agent-based modelling of socio-technical systems*, vol. 9. Springer Science & Business Media.

Davidsson, P. (2002). Agent based social simulation: a computer science view. *Journal of Artificial Societies and Social Simulation* 5 (1): 1–5.

Davis, P.K. and O'Mahony, A. (2019, this volume). Improving social-behavioral modeling. In: *Social-Behavioral Modeling for Complex Systems* (ed. P.K. Davis, A. O'Mahony and J. Pfautz). Hoboken, NJ: Wiley.

Gilbert, N. and Troitzsch, K. (2005). *Simulation for the Social Scientist*. McGraw-Hill Education (UK).

Grace, E., Blaha, L.M., Sathanur, A.V. et al. (2019, this volume). Evaluation and validation approaches for simulation of social behavior: challenges and opportunities. In: *Social-Behavioral Modeling for Complex Systems* (ed. P.K. Davis, A. O'Mahony and J. Pfautz). Hoboken, NJ: Wiley.

Kauffman, S.A. and Weinberger, E.D. (1989). The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology* 141 (2): 211–245.

Miller, J.H. and Page, S.E. (2009). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life: An Introduction to Computational Models of Social Life*. Princeton University press.

Mniszewski, S.M., Del Valle, S.Y., Stroud, P.D. et al. (2008). EpiSimS simulation of a multi-component strategy for pandemic influenza. In: *Proceedings of the 2008 Spring Simulation Multiconference*, 556–563. Society for Computer Simulation International.

Mohaghegh, Z. and Mosleh, A. (2009). Incorporating organizational factors into probabilistic risk assessment of complex socio-technical systems: principles and theoretical foundations. *Safety Science* 47 (8): 1139–1158.

Morgan, G.P., Joseph, K., and Carley, K.M. (2017). The Power of Social Cognition. *Journal of Social Structure* 18 (1): 0_1–0_22.

Murata, T. (1989). Petri nets: properties, analysis and applications. *Proceedings of the IEEE* 77 (4): 541–580.

National Research Council (2006). *Defense Modeling, Simulation and Analysis: Meeting the Challenge, National Research Council*. Washington, DC: The National Academy Press.

National Research Council (2008). *Behavioral Modeling and Simulation, from Individuals to Societies, National Research Council*. Washington, DC: The National Academies Press.

Schelling, T. (1969). Models of segregation. *American Economic Review* 59: 488–493.

Schreiber, C. and Carley, K.M. (2004). Going beyond the data: empirical validation leading to grounded theory. *Computational and Mathematical Organization Theory* 10 (2): 155–164.

Sterman, J.D. (2001). System dynamics modeling: tools for learning in a complex world. *California Management Review* 43 (4): 8–25.

Tabak, D. and Levis, A.H. (1985). Petri net representation of decision models. *IEEE Transactions on Systems, Man, and Cybernetics* 6: 812–818.

Yahja, A. and Carley, K.M. (2005). WIZER: an automated intelligent tool for model improvement of multi-agent social-network systems. In: *Proceedings of the Eighteenth International Florida Artificial Intelligence Research Society Conference (FLAIRS 2005)* (15–17 May 2005), 44–50. Clearwater Beach, FL: AAAI Press.