

Agent-based Modelling

Elizabeth M. Gallagher and Joanna J. Bryson

December 3, 2017

Synonyms

Individual-Based Modelling. Individual-Orientated Modelling. Multi-Agent Systems (has other meanings, occasionally used as synonym.)

Definition

Agent-based models are a type of model based on computer simulation, where the behaviour of a system is determined by the activities of autonomous individuals and their interaction with and through an environment.

1 Introduction

Agent-Based Modelling (ABM) is a research method for understanding the collective effects of individual action selection. More generally, ABM allows the examination of macro-level effects from micro-level behaviour. Science requires understanding how an observed characteristic of a system (e.g. a solid) can be accounted for by its components (e.g. molecules). In ABM we build models of both the components and the environment in which they exist, and then observe whether the over-all system-level behaviour of the model matches that of the target (or *subject*) system. Constructing agent-based models (ABMs) can be seen as a form of theory building. The consequence of expressing a theory as a simulation is that many aspects of the coherence and explanatory power of the theory can be checked by examining the simulation outputs.

Typical elements of an agent-based model are the attributes and behaviours of agents, the relationships between agents and their interactions, the environment, and how the agents interact with with and through the environment (Macal and North, 2010).

ABM is still a sufficiently new research methodology that there is still some controversy in its use, and still some unevenness in its application and description in scientific papers. ABMs can be very sensitive to initial conditions, and can produce complex behaviour which can be difficult and computationally

intensive to analyse and understand. Without sufficiently-established methodological practice, it can sometimes be difficult to incorporate ABM results into true scientific discourse. However, the practice is becoming more widespread and ABMs are frequently included as one form of evidence in many high-profile scientific articles.

In this article we review ABM and the techniques for its analysis. In Section 2 we discuss the history of ABM and give some examples of where it has been used. In Section 3 we discuss the considerations to be taken into account when designing and running ABMs, and popular platforms for constructing them. In Section 4 we discuss the analysis of ABMs and how this can be used in theory building. Finally in Section 5 a simple and well-known ABM is described — that of the flocking behaviour of birds by Reynolds (1987).

2 History

Agent-based models emerged from work on cellular automata (CA), the first of which was created by John von Neumann and Stanislaw Ulam in the 1940s. Their idea was to try to build a machine that could autonomously reproduce itself and the solution was a complicated set of rules on a grid. In 1970 John Conway simplified this idea in his Game of Life (1970). This game involves a grid of cells in which every cell interacts with its eight neighbouring cells according to four basic rules which decide whether the cell will live or die. These basic rules create many different emergent and complex patterns including one which copies itself in the process of destroying itself.

In the 1970s one of the first ABMs was developed. This model was developed by Schelling and was used to study segregation (Schelling, 1971). This model was implemented by physically placing dimes and pennies on a piece of graph paper, where a coin moves location if it has a majority of immediate neighbours of the other coin type. After a number of these moves the emergence of segregation of pennies and dimes can be seen. Thus, in this model agents (the coins) make decisions and interact with one another in an environment (the grid) — contrasting with CA models where it is only the environment that is considered.

After Schelling's physical implementation of an agent-based model ABMs began to be programmed using computers. This meant they could be more complicated and run for many more iterations. Thus, ABM became a more useful tool and work using them truly began. Robert Axelrod was a key player in the first work using computer based agent-based models. Axelrod's early work focused on using ABM to study the best solutions to the iterated prisoner's dilemma — a game theoretical model of cooperative behaviours (Axelrod, 1984) Axelrod remains one of the area's main advocates.

Other early applications of ABM were in biology. Hamilton (1971) and Hogeweg and Hesper (1979) used early modelling techniques to establish basic principles of animal sociality. In the 1980s ABMs also started to be used to study specific animal behaviours, such as the social interaction structure and ontogeny of bumble bees (Hogeweg and Hesper, 1983), and the flocking behaviour of birds

Reynolds (1987). After these initial developments and progressively powerful computational power, ABMs have become an increasingly popular tool in animal behaviour, ecology, the social sciences, the life sciences and economics.

A notably large early agent-based model was the simulation of whole artificial societies in the Sugarscape model by Epstein and Axtell (Epstein and Axtell, 1996). This model looks at a simple grid environment where sugar is distributed unevenly, agents move from cell to cell and use the sugar as a resource (whether it be for food or a trading material). By adding more rules, the model can be interpreted to show complex social behaviour including disease transmission, inheritance, trade networks, economic inequality and combat.

Research using ABMs has also included how culture changes and spreads (Axelrod, 1997a), the evolution of cooperation (Axelrod, 1997b), the societal collapse of the Anasazi people (Dean et al., 2000), social behaviour in primates (Hemelrijk, 2000), the study of agricultural economics (Berger, 2001), the spread of cancer (Preziosi, 2003), land use (Brown et al., 2005), predator-prey relationships between killer whales and other marine mammals (Mock and Testa, 2007), sharing information (Čače and Bryson, 2007), crowd behaviour during emergency evacuation (Pan et al., 2007), and the adaptive immune system (Folcik et al., 2007). They have become an established method used broadly in high-impact publications in the social and biological sciences, including studies of political party policy dynamics in Ireland (Laver, 2005), schisms in religion (Whitehouse et al., 2012), the impact of naive individuals on consensus formation (Couzin et al., 2011), the appearance of modern human behaviour (Powell et al., 2009), and the origins of war (Choi and Bowles, 2007).

Agent-based models can therefore now be accepted as widely used. They are able to demonstrate interesting, unexpected and complex behaviours even from very basic rules, sometimes of extreme scientific importance. However, without properly analysing and validating the results found from an agent-based model the work can be as ungrounded as a computer game, rather than demonstrating real scientific progress. We address this issue in Section 4.

3 Designing and Running ABMs

Designing a model involves knowledge and assumptions about the system being modelled. Models are abstractions of reality, and therefore they do not have to contain every single factor to do with the system. Because of this, decisions on which important components to include in the model, and those which do not need to be included, need to be made. These decisions can be based on which dynamics are of particular interest, the aspects of the system for which data is available, and more subjective or literature-based decisions about which factors are believed to be theoretically important.

There are several platforms for implementing agent-based models; including NetLogo (Wilensky, 1999), Repast (North et al., 2013), MASON and Swarm. A comparison of these is given by Railsback et al. (2006). However, the use of a platform is not mandatory and an agent-based model can be coded relatively

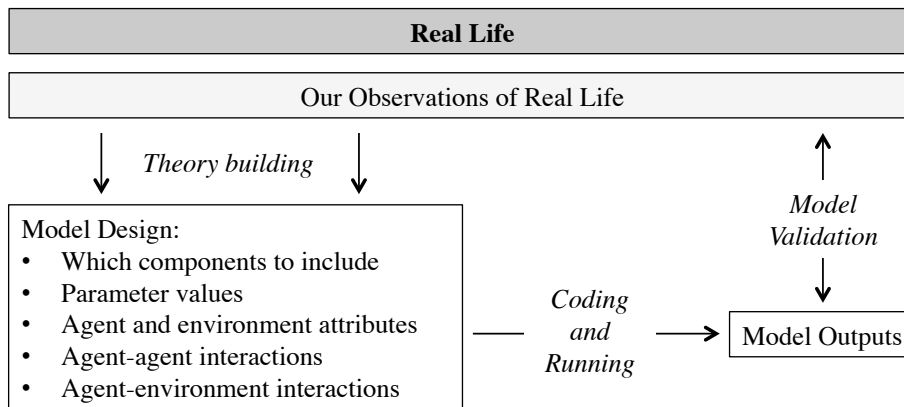


Figure 1: A flow diagram of key steps in the modelling process. It is important to note that the observed data which is used to inform on the model’s design should not also be used in the model-validation step.

easily in an ordinary programming language, particularly languages that are object oriented. However, many platforms provide tools for visualisation and analysis that model authors find useful.

Often there will be stochastic elements to ABMs. Due to this every time a model is run the results will be somewhat different, just as in many experiments with the physical world or living animals. In these cases, models should be run multiple times in order to know that the results found are representative. By running the ABM many times the user can get an idea for how variable the results are, and what kind of general model results occur. If the results of different runs of the model vary more largely than is predicted by theory, it could be a sign that the model has not been implemented properly and there may be a bug in the code, or equally that the theory is incorrect. Determining which is the problem is also a matter of analysis.

After designing, building and running the model it should then be verified and validated. These will be discussed in the next section. Key steps in the modelling process are shown in Figure 1.

4 Analysing ABMs

After running an ABM the next step is to analyse the resulting outcomes. Analysing results is done by observing what happens in the model after a certain number of iterations, or over a number of runs. Different experimental runs may use either the same parameters, in order to discover the range of possible results due only to the effects of random variation; or use systematically varying parameter values, to test the significance of each parameter set or *condition* (parameter sensitivity analysis). In-depth analysis of these observations can aid in understanding how the ABM is working and how sensitive it may be to different

parameter values.

In order for ABM (or any model) to be useful to science, it must provide both a means of explanation and a mechanism for improving that explanation. The explanatory force of the model is the extent to which an observed meta-level phenomena can be accounted for by the behaviour of its micro-level actors (Bryson et al., 2007). Thus, in model validation the explanatory force of the model is analysed. Models that do not perfectly match real world data can be improved by postulating additional characteristics of the model; those that do meet requirements may possibly be improved or made more general by simplification — that is by discovering whether any aspects of the model are not necessary to the outcome and can be removed.

4.1 Parameter Sensitivity Analysis

Sensitivity analysis is used to investigate how input parameters affect the outcomes of the model. If there is uncertainty in the value for the input parameters then this analysis will give some idea of the degree to which this impacts the results from the model. Unless a particular range of parameter values is specified in the theory behind the model, over-dependence on specific values can be seen as fragility in the model, though they may also serve as a prediction coming from the model concerning what values the parameter correlates may be expected to hold in the target system the model describes.

A commonly used type of parameter sensitivity analysis is investigating the effect of varying one parameter whilst keeping all the others at their default values (the ‘fix-all-but-one’ method). With increasing model complexity the effects parameters have on the outcomes of a model can become exponentially more difficult to analyse. Varying one parameter at a time reduces this complexity, but relies on having a default value to set the other parameters to. Where such values are not known, e.g. when there is insufficient data available, an ill-considered default value may be chosen and an important range of model outcomes overlooked. However, if there is data which suggests that only a small range of values are realistic, then the complexity of the search may be reduced even where no exact value is known.

A more thorough type of parameter sensitivity analysis is to vary all parameters randomly within a predefined range of realistic values. By doing this for many runs of the model the parameter set space can be investigated more comprehensively and there is less chance of overlooking important model outcomes because of ill-defined default parameter values. However, since this approach investigates multiple parameter changes at the same time the results are more challenging to analyse than those using the fix-all-but-one approach. Cluster analysis can be used to distinguish the types of model output that occur when the model is run many times with random parameters, and then the parameter value distributions for each cluster can be determined (Gallagher, 2017, pg. 162)

4.2 Model Verification and Validation

As ABM has become more prevalent there has been an increasing emphasis on developing methods of verification and validation (Balci, 1998; Kennedy et al., 2005). *Verification* is the process of making certain a model runs as designed. In science, this is roughly equivalent to ensuring that good experimental practice has been followed. This can be done by debugging and peer reviewing the code to make sure that all components of the system act as expected. *Validation* is the process of making certain the model actually models the target system. When ABMs are used for sciences such as biology, validation is equivalent to hypothesis testing (Bryson et al., 2007). Model validation may be seen as one mechanism of model verification, provided that successfully replicating the behaviour of the target system is unlikely to occur by chance.

There are only two important criteria for validating an ABM. These are the same as for validating any behavioural model:

1. Does the behaviour of the ABM match the behaviour of the target system within the standard metrics of hypothesis evaluation?
2. Do all the attributes of the agents and their environment have plausible correlates in the target system being modelled?

Regarding the ‘standard metrics’, these depend largely on the success of previous explanatory efforts. If the literature contains no prior explanation or model, then it may be sufficient to show a qualitative similarity between the model and the target system. The model is now a theory explaining the data, and as the first one it is necessarily the best. However, if there is another competing model, then we need either to use standard statistical hypothesis testing to decide which will be the better match, or possibly to argue one model is more parsimonious, though such arguments can be problematic (Myung et al., 2000). For the second criterion, the issue is whether the modeller has given the artificial agents any capacities that real subjects could not or arguably would not possess.

There is a common perception that ABMs are so complex (that is, have so many parameters) that they can be made to easily match any data or predict any outcome, but that having done so the system will have no capacity for generalisation, and therefore no predictive power. In practice, however, building and debugging ABMs is a difficult skill, and matching datasets is not easy. While it is true that ultimately most datasets can be matched, the principle value of the model is expressing what aspects need to be changed in order to generate these various outcomes. Again, a working model is best understood as a theory of the system, and any parameters and parameter values required for that model to work that were not a part of the original theory are predictions derived from that theory.

If a model is built first to a set of justified assumptions, and *subsequently* matches a dataset with minimal adjustment, then it can be considered to be at least partially validated. Of course, the more datasets it matches, the better-

validated the model becomes. As this notion of *better*-validated implies, validation is not simply a state that either holds or does not hold for a model. Rather a model, like any scientific hypothesis, becomes more *probable* (given the data) the more it is validated. But a model never becomes perfectly certain (Box, 1979). Indeed, the purpose of any theory is to abstract from the real world, and that process of abstraction necessarily entails losing detail and precision. The only exception is if a model becomes understood to such an extent that it can be *proven* correct in a logical or formal analytic sense.

This observation returns us to the matter of verification. Verification is most pernicious in purely formal systems, where validation cannot be grounded in real-world data. A proof can only be judged by human reasoning, and is only as good as its premises (Bundy et al., 2005; Edmonds and Bryson, 2004). Formal systems are used in mathematics and similar disciplines as a mechanism of knowledge *discovery*, and therefore verification is both more critical and more difficult. Again, when validation is performed via hypothesis testing against real-world data, validation itself serves as a form of verification. To the extent a computational model reliably matches and predicts a target system’s performance, then it *is* a model, in the formal sense of the word.

4.3 ABMs as Scientific Hypotheses

The scientific method is a process of “systematic observation and experimentation, inductive and deductive reasoning, and the formation and testing of hypotheses and theories” (Andersen and Hepburn, 2016). In this way hypotheses can be continuously improved by refinement or rejection. As stated above, ABMs may be viewed as a scientific theory about how the real-life system behaves, and as such model validation is a form of hypothesis testing (Bryson et al., 2007). Hence, every time we validate a model we can learn something new about how the real-life system might work, and then refine hypotheses about this system accordingly. Thus, in the next version of modelling a particular system we may have even more robust ideas about parameter values or relationships within the model and its target system.

The analysis of an existing agent-based model can be done as a three phase process (Bryson et al., 2007). The first phase is a *replication* of the model. This may not seem (or even be) strictly necessary in the case where the model is publicly available — the results in that case can be checked just by rerunning the original model on another computer. However, reimplementing the model from its description in the literature can be a valuable exercise. Reimplementation may uncover important aspects of the model that the model’s original authors either took for granted, overlooked, or even forgot about during the course of their research (King, 1995; Axtell et al., 1996). ABMs may be valid without actually having been fully verified or understood. This is true of any scientific hypothesis; part of the scientific method is improving this understanding of a theory as a community.

Once the critical attributes of the model are well-understood, we can enter the second phase of the analysis, *model understanding*. Here, we carefully con-

sider what the implied or the explicit correlates of those attributes are. Again, just as in any science, we go through a process of finding testable predictions and implications that result from our hypothesis. The third and final phase is *testing* these predictions and implications, looking first into the extant literature, and then (if necessary) to proposing and executing new experiments or measures on the target system.

With regard to the second phase, two similar methods can be used to predict what parameter values in the model give the most realistic results, and thus help us to make inferences about the real-life system. These methods are fitting to idealised outcomes (FIO) (Gallagher et al., 2015) and approximate Bayesian computation (ABC) (Beaumont, 2010). Both approaches have three steps:

1. The model is run many times with random parameter values.
2. Observed data, whether rich in detail (in the case of ABC) or imprecise (in the case of FIO), is compared to every run of the model. Model runs which give the closest match to the observed data are then separated out.
3. The parameter value combinations which were used in these closest-match model runs are interpreted as potentially being most realistic. Where the model includes stochastic elements, it is wise to check that the fit holds for repeated runs.

As mentioned in Section 4.1 step 1 allows trends between the parameters and the outcomes to be seen, as well as studying parameter interdependencies and sensitivity.

Step 2 is where ABC and FIO differ. In ABC summary statistic(s) from observed data are compared to the same summary statistic(s) from the modelled data. Thus in ABC the comparison of the simulation and observed data is much more robust, and thus ABC can be used to validate complex models. However, to do this ABC relies on having detailed observed data which is not always possible. Thus, if observed data is not rich enough to allow ABC, but either some general outcome is known or is of interest, FIO can still be used in model validation.

In step 3 the two methods allow a prediction of what parameter values are likely to cause observed data, and which parameter values are not. And thus using these methods we can make new and informed predictions about the real-life system, which should then be tested in the third phase of analysis.

5 Example – Flocking Behaviour

A popular and early example of ABM is the model of coordinated animal movement posited by Craig Reynolds (Reynolds, 1987). This model, called Boids, has proven extremely fertile in biology (Couzin et al., 2011; Hemelrijk, 2000).

The Boids model has three basic rules which determine how each agent (or ‘boid’) moves:




Figure 2: Positions of agents in the Boid agent-based model by Reynolds (1987). Implemented in the ‘Flocking’ model example in NetLogo 5.1.0 (Wilensky, 1999). (a) shows the starting positions of agents, (b) shows the positions of agents after 100 iterations of the model when agents can see moderately far around them, and (c) shows the positions of agents after 100 iterations of the model when agents can see very far around them.

1. Separation: if another agent is too close, then steer to avoid it.
2. Alignment: calculate the average direction of the surrounding agents and steer to move in this same direction.
3. Cohesion: steer to move toward the average position of nearby agents (but do not get too close).

Depending on various parameter values chosen in this model (e.g. the minimum distance apart and the maximum distance an agent can observe around itself) and how long the model is run, different macro-level behaviour can emerge. For example, if the agents can be quite close together and see moderately far around them then here might be multiple flocks of agents, but if they can see very far around them then a single flock of agents can emerge. Examples of these flocking behaviours can be seen in Figure 2.

The initial locations and directions of agents are randomly chosen every time the model is run, and thus even with the same parameters each run of the model can look very different.

As an entirely fictitious example for purpose of illustrating the process of making predictions from this model, let’s assume that the average number of birds in a real-life flock is 20, for some species and set of circumstances. We can run the model many times whilst varying the parameter which accounts for how far around it each agent can observe other agents. Every time the model is run we record the mean number of agents in each flock. By selecting the runs of the model which gave average flock sizes of 20 we can then find which parameter value was used in these runs. Thus, from this we can make a prediction of the distance birds can observe the direction of travel and positions of other birds.

A real-world example of an analysis of a very similar model to boids posited to explain primate social hierarchies is presented in Bryson et al. (2007).

6 Conclusion

Since the 1980s agent-based modelling has become an increasingly popular tool in animal behaviour studies, social sciences, economics, and life sciences. ABM focuses on modelling the micro-level behaviours of the system, and examining the macro-level effects these can have. To be useful to theory building ABMs need to be validated by comparison to the real-life system which is being modelled. In this way new insights, sharper intuitions and predictions can be made.

7 CrossReferences

artificial categories, exemplar theory, feature learning (maybe), theory of categorization (maybe), model fitting, behavioral variation, invasive research, Turing test, Proximate Causation, linkage analysis, parsimony, neural network, virtual reality experiments, intervening variables, motivational state, effect size, Bayesian causality, teleology, threshold, social network analysis, intervention (if that's about experiments)

References

- Andersen, H. and Hepburn, B. (2016). Scientific Method. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2016 edition.
- Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books.
- Axelrod, R. (1997a). The dissemination of culture a model with local convergence and global polarization. *Journal of Conflict Resolution*, 41(2):203–226.
- Axelrod, R. M. (1997b). *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton University Press.
- Axtell, R., Axelrod, R., Epstein, J. M., and Cohen, M. D. (1996). Aligning simulation models: A case study and results. *Computational & Mathematical Organization Theory*, 1(2):123–141.
- Balci, O. (1998). Verification, validation, and testing. volume 10, pages 335–393. New York, NY: John Wiley and Sons.
- Beaumont, M. A. (2010). Approximate Bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics*, 41:379–406.

- Berger, T. (2001). Agent-based spatial models applied to agriculture: A simulation tool for technology diffusion, resource use changes and policy analysis. *Agricultural Economics*, 25(2-3):245–260.
- Box, G. E. (1979). Robustness in the strategy of scientific model building. *Robustness in statistics*, 1:201–236.
- Brown, D. G., Page, S., Riolo, R., Zellner, M., and Rand, W. (2005). Path dependence and the validation of agent-based spatial models of land use. *International Journal of Geographical Information Science*, 19(2):153–174.
- Bryson, J. J., Ando, Y., and Lehmann, H. (2007). Agent-based modelling as scientific method: A case study analysing primate social behaviour. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1485):1685–1699.
- Bundy, A., Jamnik, M., and Fugard, A. (2005). What is a proof? *Philosophical Transactions A: Mathematical, Physical and Engineering Sciences*, 363(1835):2377–2391.
- Čáče, I. and Bryson, J. J. (2007). Agent based modelling of communication costs: Why information can be free. In *Emergence of Communication and Language*, pages 305–321. Springer.
- Choi, J.-K. and Bowles, S. (2007). The coevolution of parochial altruism and war. *Science*, 318(5850):636–640.
- Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., Conradt, L., Levin, S. A., and Leonard, N. E. (2011). Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062):1578–1580.
- Dean, J. S., Gumerman, G. J., Epstein, J. M., Axtell, R. L., Swedlund, A. C., Parker, M. T., and McCarroll, S. (2000). Understanding Anasazi culture change through agent-based modeling. *Dynamics in human and primate societies: Agent-based modeling of social and spatial processes*, pages 179–205.
- Edmonds, B. and Bryson, J. J. (2004). The insufficiency of formal design methods — the necessity of an experimental approach for the understanding and control of complex systems. In Jennings, N. R., Sierra, C., Sonenberg, L., and Tambe, M., editors, *The 3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS 2004)*, pages 936–943. ACM Press.
- Epstein, J. M. and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Brookings Institution Press.
- Folcik, V. A., An, G. C., and Orosz, C. G. (2007). The Basic Immune Simulator: an agent-based model to study the interactions between innate and adaptive immunity. *Theoretical Biology and Medical Modelling*, 4(1):1.

- Gallagher, E. M. (2017). Evolutionary Models for the Origins of Agriculture. Unpublished Doctoral Thesis, University College London.
- Gallagher, E. M., Shennan, S. J., and Thomas, M. G. (2015). Transition to farming more likely for small, conservative groups with property rights, but increased productivity is not essential. *Proceedings of the National Academy of Sciences*, 112(46):14218–14223.
- Gardner, M. (1970). Mathematical games: The fantastic combinations of John Conway’s new solitaire game “life”. *Scientific American*, 223(4):120–123.
- Hamilton, W. D. (1971). Geometry for the selfish herd. *Journal of Theoretical Biology*, 31:295–311.
- Hemehrijck, C. K. (2000). Towards the integration of social dominance and spatial structure. *Animal Behaviour*, 59(5):1035–1048.
- Hogeweg, P. and Hesper, B. (1979). Heterarchical selfstructuring simulation systems: Concepts and applications in biology. In Zeigler, B. P., Ezas, M. S., Klir, G. J., and Ören, T. I., editors, *Methodologies in Systems Modelling and Simulation*, pages 221–231. North-Holland, Amsterdam.
- Hogeweg, P. and Hesper, B. (1983). The ontogeny of the interaction structure in bumble bee colonies: a MIRROR model. *Behavioral Ecology and Sociobiology*, 12(4):271–283.
- Kennedy, R., Xiang, X., Madey, G., and Cosimano, T. (2005). Verification and validation of scientific and economic models. In North, M., Sallach, D., and Macal, C., editors, *Proc. Agent 2005: Generative Social Processes, Models, and Mechanisms*, pages 177–192. Chicago, IL: Argonne National Laboratory.
- King, G. (1995). Replication, replication. With comments from nineteen authors and a response, ‘A revised proposal, proposal’. *PS: Political Science & Politics*, 28(3):444–452.
- Laver, M. J. (2005). Policy and the dynamics of political competition. *American Political Science Review*, 99(2):263–281.
- Macal, C. M. and North, M. J. (2010). Tutorial on agent-based modelling and simulation. *Journal of Simulation*, 4(3):151–162.
- Mock, K. and Testa, J. (2007). An agent-based model of predator-prey relationships between transient killer whales and other marine mammals. *University of Alaska Anchorage, Anchorage, AK, Tech. Rep.*
- Myung, J., Forster, M. R., and Browne, M. W. (2000). Special issue on model selection. *Journal of Mathematical Psychology*, 44(1).
- North, M. J., Collier, N. T., Ozik, J., Tatara, E. R., Macal, C. M., Bragen, M., and Sydelko, P. (2013). Complex adaptive systems modeling with repast simphony. *Complex adaptive systems modeling*, 1(1):3.

- Pan, X., Han, C. S., Dauber, K., and Law, K. H. (2007). A multi-agent based framework for the simulation of human and social behaviors during emergency evacuations. *Ai & Society*, 22(2):113–132.
- Powell, A., Shennan, S., and Thomas, M. G. (2009). Late Pleistocene demography and the appearance of modern human behavior. *Science*, 324(5932):1298–1301.
- Preziosi, L. (2003). *Cancer Modelling and Simulation*. CRC Press.
- Railsback, S. F., Lytinen, S. L., and Jackson, S. K. (2006). Agent-based simulation platforms: Review and development recommendations. *Simulation*, 82(9):609–623.
- Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH computer graphics*, 21(4):25–34.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186.
- Whitehouse, H., Kahn, K., Hochberg, M. E., and Bryson, J. J. (2012). The role for simulations in theory construction for the social sciences: Case studies concerning Divergent Modes of Religiosity. *Religion, Brain & Behavior*, 2(3):182–224.
- Wilensky, U. (1999). Netlogo. <http://ccl.northwestern.edu/netlogo/>, Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.