

WIZER: What-If Analyzer for Automated Social Model Space Exploration and Validation

Alex Yahja and Kathleen M. Carley

Institute for Software Research, International

Center for the Computational Analysis of Social and Organizational Systems

Carnegie Mellon University

Pittsburgh, Pennsylvania 15213

Abstract

Complex social problems modeled by multi-agent systems have very large parameter and model space. The problem of how to model, validate, detect, and plan for the event of bioterrorism is one of these, as it requires faithful modeling of dynamic signal (bioattack event) from complex dynamic noise (normal disease outbreaks and people activities). Indeed, the dynamic and very large space – numeric or symbolic or both – nature of the problem makes manual exploration spotty, cumbersome, implicitly-biased, and thus incomplete. Scaling up multi-agent systems exacerbates these and makes the automation of exploration, modeling, and validation more critical. WIZER – a social inference engine and simulation combination capable of principled exploration through meta-models and parameters based on empirical data and knowledge – addresses the above problems by knowledge-guided & simulation-guided search. This paper describes the design of WIZER and presents a preliminary result.

Contact:

Alex Yahja

Institute for Software Research, International

Center for the Computational Analysis of Social and Organizational Systems

Carnegie Mellon University

Pittsburgh, PA 15213

Tel: 1-412-268-5866

Fax: 1-412-268-1744

Email: ay@cmu.edu

Key Words: multi-agent systems, validation automation, social inference engine, model space search, knowledge-guided search, dynamic behavior, simulation, modeling

Acknowledgement: This research was supported, in part, by DARPA for work on Scalable Biosurveillance Systems, the NSF IGERT9972762 in CASOS, and by the Carnegie Mellon Center on Computational Analysis of Social and Organizational Systems.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA, the National Science Foundation, or the US Government.

WIZER: What-If Analyzer for Automated Social Model Space Exploration and Validation

Alex Yahja

Computational analysis [Axelrod 1997][Carley 1995][Carley & Lee 1998][Burton & Obel 1995] is re-shaping how we think about how to model individuals, teams, groups, markets, organizations, institutions, and societies. Statistical methods on their own are inadequate as they could not capture the highly dynamic, heterogeneous, symbolic-numeric, cause-effect nature of the domain. Outside randomized experiments, statistical methods can capture correlation but not causality. Computational models can provide detailed insights otherwise impossible on complex and realistic social systems through faithful modeling and model space exploration [Prietula, Carley & Gasser 1998]. One of the computational models is the multi-agent model.

The reality of social science is that there are voluminous findings in publication about diverse things [Bednarz & Bednarz 2000][General Social Surveys at <http://www.icpsr.umich.edu:8080/GSS/homepage.htm>][Christopherson 1987][Oliver 1986]. With the exception of social network analysis [Wasserman & Faust, 1994], most of this huge knowledge is not – unfortunately – operational, meaning that it is not part of any computerized experimental model. Interviews and real world experiments commonly performed in social science have the limit on how many experiments one could perform, not to mention the ethical constraints limit. While virtual experiments are not the same as real experiments, virtual experiments by the virtue of fast execution over many scenarios can provide insights that real experiments cannot. Moreover, by validating virtual experimental models with empirical data & knowledge, we can be assured that they closely match the reality with certain confidence level. While experiment design is mostly based on statistics [Box, Hunter & Hunter 1978][Cornell 1990][John 1998][Keppel 1991], recent advances allow the explicit modeling and operationalization of causality [Pearl 2000] [Glymour, Scheines & Spirtes 2001] thus enabling the automation of experiment design.

As most multi-agent systems [Sycara 1998][Payne et. al. 2002][Sycara et. al. 2002][Jensen & Lesser 2002] are complex, finding out what really happens in their virtual experiments is hard, especially if we want to have a complete picture. It also follows that validating them is hard. One solution is to have computers do brute-force search over the space of ensembles interactively guided by human beings. This is known as “Exploratory Modeling”, in which the virtual experiments are meant to be used in arguments that do not depend on the predictive accuracy of the models [Banks 2002]. Exploratory modeling gives us the capability to identify an ensemble of plausible models and modeling assumptions, identify the range of outputs predicted by plausible models under plausible assumptions, identify the relationship between modeling assumptions and model outputs, find assumptions that have a large impact on model outputs, and identify predictions that are robust across different modeling assumptions [Banks 1994][Banks & Gillogly 1994][Banks & Gillogly 1994b][Banks 1993][Dewar et. al. 1993].

A better method would be to let the computer do the search automatically. As advances on causal calculus enable the modeling and inference of causality [Pearl 2000][Tian & Pearl 2002][Tian & Pearl 2001], it is now becoming possible to create a social inference engine to do automated experiment design and evaluation. This engine while complex – and because of its ability to handle complex social situations and processes – fits the social problems and supports the scaling up of multi-agent models. The combination of the engine with multi-agent virtual experiments allows the computer to gain insight to what is happening in the experiments.

Problem Statement

Principled automated exploration of response surface and problem space for complex non-Markovian systems and validation of their simulation systems have not been adequately solved. The problem lies in disconnect between search algorithm and hypothesis testing via real & virtual experimentation, in the underutilization of compiled domain knowledge, and in the lack of automation especially experiment design & evaluation automation. We propose a remedy that addresses these and apply it to the problem of scalable modeling the events of bioterrorism in a dynamic society.

Related Work

Representation and search [Peshkin 2000][Bryson 2001][Neller 2002][Craven et. al. 1999] are active research areas in Artificial Intelligence [Mitchell 1997]. However, most work on representation and search ignore the fact that experimentation may be needed and that representation may need to be dynamic. Experimentation on models and real world data is usually carried out by human researchers. Moreover, compiled knowledge is often underutilized – it is reduced to simple heuristics which trivializes important subtleties of search. There has also been

a wall of separation between Artificial Intelligence and soft sciences such as social sciences. Below we describe relevant search methods in the context of Scalable Biosurveillance Systems, or BioWar, project.

Genetic algorithm, genetic programming, and evolutionary modeling [Bentley & Corne 2001][Koza et. al. 1999][Werfel 2002][Pereira et. al. 1999][Shaw & Fleming 1996] use crossover & mutation with fitness function as search. Evolutionary modeling [Gilbert, et. al. 2000] has been successfully applied to some domains, including an automated design of some electronic circuits [Zebulum et. al. 2002]. While these techniques can achieve adequate solutions, they usually waste much effort and resources in weeding out unwanted genes. Furthermore, they use small steps of crossover & mutation, when large inferential steps may be possible. These result in the inefficiency of search. Science and technology would not progress as fast if they are done solely by crossover & mutation – absent the inference and the scientific method of experimentation. For BioWar, crossover & mutation do not correspond to the underlying social and epidemiological mechanisms. Crossover & mutation do not recognize the networked reality of social systems. They ignore empirical findings from social sciences and social network analysis.

Simulated annealing [Spears 1996] uses temperature and surface gradient as search. While this technique can find answer to optimization problems, it has no way of knowing whether it is trapped in the local minima. While for some applications, local minima are acceptable, for critical problems such as BioWar, local minima are not acceptable. This technique under-utilizes the empirical knowledge of domain fields, resulting in the less efficient jumps of search probe.

Neural networks [Schneider & Riesenhuber 2002][Kilmer et. al. 1994] use weight gradient as search. While neural networks have been successfully applied to silicon retina, the problem with neural networks lies in the interpretation of the weights. There is also a mapping problem for real world variables and models to the weights, the input variables, and the output variables of the neural networks. Neural networks are ill-suited for the purpose of mapping symbolic variables, which BioWar contains.

Work in expert/knowledge systems [Jackson 1999] has progressed to the extent that they are used in specialized areas of engineering, of medicine, and of organizational science to augment professional judgment. Expert systems usually take a form of “if-then” rule-based system, of constraint-satisfaction system, and of Bayesian networks. Examples of successful expert systems include MYCIN bacterial infections of the blood and meningitis expert, SPEED solutions pricing & end-to-end design for bandwidth trading for the British Telecom, CARMA grasshopper infestations expert for Wyoming ranchers, and OrgCon expert system for designing organizations. They operate successfully within each of their specialized domains, but one of problems is that there is less success in combining the judgments of various expert systems together. Complex problems such as bioterrorism require multiple experts and faithful combination of their judgments. Work in risk ranking and management of engineering and public policy field [Morgan & Henrion 1990][Fischhoff et. al. 1984] addresses the combination of risk judgments, but ways for combining not just risk judgments are needed for complex problems. While expert system techniques such as forward & backward chaining, Bayesian Network Inference (BNI) [Heckerman 1995][Murphy 2002], and Monte-Carlo Markov-Chains (MCMC) [Gamerman 1997][Neal 1993][Liu & Rubin 1996] produce inference, they are never meant for use in search. Additionally, both BNI and MCMC use Bayesian correlation rule – insufficient for handling causality in inference. BioSTORM (Biological Spatio-Temporal Outbreak Reasoning Module) uses Bayesian correlation in its Bayesian reasoning, in addition to Kalman filter, of its RASTA engine [Buckeridge et. al. 2002].

Work on software engineering [Bachmann et. al. 2002][Linger et. al. 2002][Hudak et. al. 2002][Brown, Carney, & Clements 1995], especially the work on code certification and automatic program synthesis [Whalen, Schumann, Fischer, 2002] allows automated code generation and certification generation in a specialized domain. Related work allows automated generation of statistical data analysis programs [Fischer & Schumann, 2002]. While these are significant, it does not take into account the future potential changes to the specifications nor the structure of the social & physical world. Complex and changing problems such as bioterrorism often require frequent updates to the specifications due to changing real world.

Work on organizations and artificial societies is beginning to shed light into the complex workings of social systems [Carley & Prietula, 1999][Epstein & Axtell, 1997][Lomi & Larsen, 2001][Carley 2001]. Some of these works involved extensive simulations and validation to show the results of the systems. However, due to the limitation of computer hardware and software technologies, most validations are done manually.

The state-of-the art of cognitive modeling is that of Soar and Act-R, while both models have successfully modeled how humans solve problems, they have difficulty of scaling up to large problems. Part of the problem is that they focus solely on **inference on rules**, either pure symbolic (Soar, Multi-agent Soar, and Social Soar) or with an additional numerical underpinning (Act-R). We envision that simulation is a necessary part of intelligence. In other words, **inference on rules & simulations** is critical. Simulations are needed to resolve conflicts between rules better, to provide constraints & contexts, and to enable the inference engine to do virtual experiments. Humans think

partly by seeing things through, not necessarily logically (but maybe socially, habitually, culturally, economically, organizationally, politically, etc.). Humans often ask “what-if” questions naturally and perform virtual experiments (a.k.a. mind experiments). Humans also use various sources of knowledge: it would not be wise not to use things we know.

What-If Analyzer

We have taken an approach of exploring the model space by using a combination of real-world simulation and causal inference engine that captures societal norms & behaviors. Domain knowledge is used as constraints to the causal inference engine in generating a new search step and to the real-world simulation. We call the complete system WIZER (for **What-If-AnalyZER**) with the Causal Inference Engine & Meta-Modeler and the Real-World Simulation as its primary components, as described below.

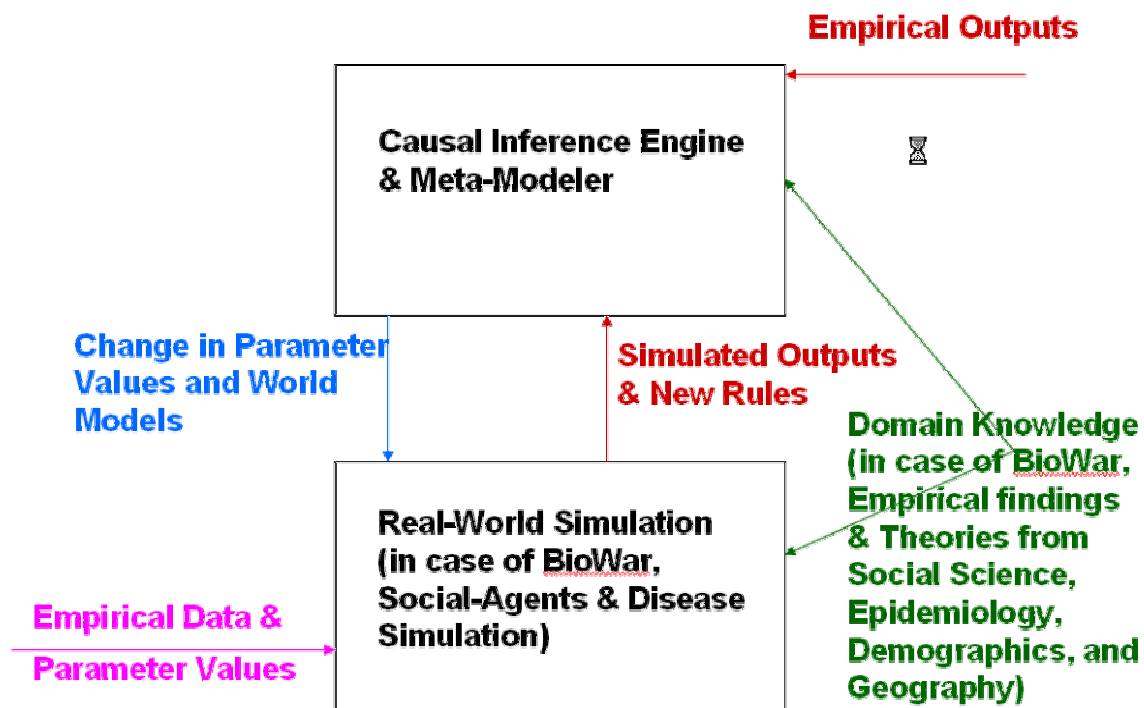


Figure 1. Components of What-If Analyzer

Causal Inference Engine & Meta-Modeler (CIEMM) takes simulation outputs, empirical outputs (e.g., in the form of target response surface), and domain knowledge (serving as constraints and rules), and tries to determine the causes of discrepancy between the simulation outputs & empirical ones and then designs a new simulation and a new virtual experiment. This new simulation could be viewed as an intelligent search in the dynamic multi-dimensional space. Not only can the parameter values change, but also the simulation models themselves. CIEMM combines causal calculus, various kinds of rules (logical inference rules, social rules, social networks rules, etc.), simulation, and multiple domain knowledge to arrive at a better simulation and/or virtual experiment for the next step in search. As CIEMM operates on social knowledge, we also call it **Social Inference & Experiment Engine**. This engine acts as an automatic experiment designer and combines the intermediate steps and judgments from various experts in the form of expert systems of social science, networks, epidemiology, organization science, psychology, software engineering, etc.

Real-World Simulation (RWS) takes empirical constraints, new parameter values, domain knowledge, and new models from CIEMM, and then reconstructs the simulation models and rerun the simulation. The outputs of a simulation run are in many formats: numerical & symbolic response surface, performance data (e.g., how fast a

simulation runs), and new rule/knowledge. Note that these simulated findings are checked by CIEMM against the empirical findings to see if they are reasonable before they are incorporated into CIEMM's knowledge base.

Domain knowledge is represented in a novel way. Not only would the knowledge be represented symbolically and/or numerically, it would be also be represented by a snapshot(s) of simulations. For example, if the empirical knowledge says “kids go to schools within the school district near their homes”, there would be (1) a symbolic causal graph representation of this knowledge in CIEMM, (2) a direct association to the parameter values and models of RWS (in the case of BioWar, RWS is in the form of social network multi-agent system), and (3) a link to the semantic web and/or to the ontology of existing publications providing the context for this knowledge. We can view this as knowledge representation as simulation.

During the course of the cycles of simulation and inference, both RWS and CIEMM are re-grounded on empirical data or ground truth every few cycle. As the models and parameters evolve, the grounding causes the models more closely mimic the reality, and this is done automatically. This represents a way to address the research question of how to faithfully match simulation to reality and how to make multi-agent system scalable. In the absence of the abundance of empirical data, exploring the space is still possible provided we have good enough domain knowledge.

Social Inference & Experiment Engine

Social Inference and Experiment Engine contains components as shown in Figure 2 below.




Figure 2. Components and Data Streams of Social Inference & Experiment Engine

As shown, the knowledge miner & causal relation extractor extracts both knowledge nuggets (hard knowledge such as causal relations, rules, facts, hard constraints, etc.) and soft knowledge (patterns, norms, beliefs, culture, trends, differentials, comparisons, soft constraints, etc.) from existing empirical literature and data of diverse domain fields. The causal detector uncovers causal relations from simulation happenings and put the results for the causal inference engine to turn into knowledge nuggets to store in the knowledge base. Note that the simulation derived rules are tagged as such to distinguish them from empirically derived ones. The trend inference engine examines the differentials and trends of simulation outputs based on empirical data, resulting in new soft knowledge. The knowledge base has a correspondence with simulation model & parameters, as we represent knowledge partly as simulation. The experiment designer takes simulation histories, performance measures, old multi-agent model, old experiment specification & results, along with simulation nuggets, soft knowledge, and knowledge nuggets to generate a new experiment specification. This specification is fed into the meta-modeler to generate new multi-agent model/submodel with the help of software engineering knowledge. The automated code generator then generates a new partial or whole multi-agent model for the simulator. The experiment executor controls the simulation runs based on commands given by the automated code generator.

Application

WIZER would be used to improve the fidelity and the scalability of BioWar social multi-agents model simulation and to significantly increase the speed at which BioWar could be reconstructed and revalidated.

WIZER Version 0

Work is in progress in the implementation of WIZER Version 0. Due to the complex scope of the task, we have taken an approach of incremental improvement. The implementation is done in C++ on top of the existing BioWar simulation. We are transitioning to using more flexible languages such as CLOS and C++ code generator. This version of WIZER only changes input parameters and leaves unchanged the models.

WIZER Version 0 takes the empirical outputs from various sources below and tries to change the input parameters and the multi-agent model parameters to fit the simulated outputs to the empirical as closely as possible. It also generates alarms if certain simulated outputs go beyond certain bounds.

Empirical output data sources:

- Doctor and emergency room visits
- Workplace visits and absenteeism
- School visits and absenteeism
- Pharmacy visits

The heart of WIZER version 0 is a routine(s) that compares the empirical output curve with the simulated one to determine which part fits, is outside the boundaries, or needs increasing/decreasing. Once determined, the routine(s) examines which input and model parameters need adjustments (for version 0, causal links between input and model parameters to outputs are determined manually). The simulation is then re-run and the process repeats till sufficient fidelity is achieved. Curve fitting is done either sequentially or in a batch of multiple runs.

Algorithm WIZER0_Curve_Fit

```

Inputs: simulation parameters SP, model parameters MP,
        number of batch simulations N
Constraints: social rules and boundaries RB (for setting off alarms)
Outputs: simulated data curves SC, empirical data curves EC
do while (the fidelity of the model is insufficient)
  n=0
  do while (n<the number of batch simulations N)
    do a simulation run
    collect input, model, output, performance statistics, causal relations
    n=n+1
  end
  take the statistics (mean, variance, etc.) and do curve fitting of SC against EC
  by adjusting SP and MP for the next batch of simulations,
  constrained by their causal relations
  if some of the statistics are outside normal bounds RB,
    set off specific alarms
  check if the fidelity of the model is sufficient
end
end

```

An example of the run: BioWar outputs of the number of people going to work, school, pharmacy, doctor, and ER are compared to the user-defined ranges. If they are not within the allowable ranges, the thresholds that govern agent's decision to go to places are automatically increased or decreased. A typical cycle of runs is shown below.

```
original thresholds: 5 20 130 260 0
Work: mean 24880.8, std 15741.7
School: mean 6499.44, std 4115.47
Pharmacy: mean 1427.91, std 762.439
Doctor: mean 199.369, std 146.913
ER: mean 30.1935, std 24.1886
work is outside bound
threshold t0 is the actual cause of work being too low
work is too low, decrease th0
school is within bound
pharmacy is within bound
doctor is within bound
er is outside bound
thresholds th2 and th3 are the actual causes of er being too high
er is too high, increase th2, increase th3
modified thresholds: 3 20 132 262 0
```

Future Directions

This paper describes a knowledge-based and simulation-based approach to validate and scale up multi-agent systems. Scalability to complex reality is critical. Indeed, an improved scalable model that can handle local variations would provide higher fidelity modeling of signal and noise, resulting in more precise detection of the events of bioterrorism. Precise detection allows precise planning, which would save lives.

Future iterations of WIZER will facilitate more complex social constraints and inferences. Evaluation & optimization of WIZER performance will be done as well. It is hoped that WIZER would become a general system, allowing the automation of the construction and validation of complex multi-agent systems and – in the more distant future – of software systems. It is envisioned that people will no longer just write books or journal papers when they publish their ideas, but also “write” working simulations. A big project to build an automated integrated gigantic simulation of the whole Earth including all its systems – natural and otherwise – may also happen in the distant future. As the fidelity of this simulation gets very high due to frequent automated updates, social & natural scientists will be able to perform experiments on this simulated world otherwise impossible in the real world.

References

- [Axelrod 1997] Axelrod, Robert, "Advancing the Art of Simulation", *International Conference on Computer Simulation and the Social Sciences*, Cortona, Italy, 1887.
- [Banks 2002] Banks, Steven, "Reasoning with Complex Models Using Compound Computational Experiments and Derived Experimental Contexts", CASOS Conference, 2002.
- [Banks 1994] Banks, Steve, "Computational Experiments and Exploratory Modeling", RAND Corp., RP-273, 1994.
- [Banks & Gillogly 1994] Banks, S. and Gillogly, J., "Exploratory Modeling: Search Through Spaces of Computational Experiments", RAND Corp., RP-345, 1994.
- [Banks & Gillogly 1994b] Banks, S. and Gillogly, J., "Validation of Exploratory Modeling", RAND Corp., RP-298, 1994.
- [Banks 1993] Banks, S., "Exploratory Modeling for Policy Analysis", RAND Corp., RP-211, 1993.
- [Bachmann et. al. 2002] Bachmann, F., et. al., "Illuminating the Fundamental Contributors to Software Architecture Quality", CMU Software Engineering Institute, CMU/SEI-2002-TR-025, 2002.
- [Bryson 2001] Bryson, Joanna J., "Intelligence by Design: Principles of Modularity and Coordination for Engineering", MIT AI Lab, AITR-2002-003, September 2001.
- [Bednarz & Bednarz 2000] Bednarz, S.W. and Bednarz, R., "Social Science on the Frontier", SSEC publication, 2000.
- [Bentley & Corne 2001] Bentley, P.J. & Corne, D.W., "Creative evolutionary systems", Morgan Kaufmann, 2001.
- [Box, Hunter & Hunter 1978] Box, G. E. P., Hunter, W. G. and Hunter, J. S., "Statistics for Experimenters", John Wiley, 1978.
- [Brown, Carney, & Clements 1995] Brown, A.W., Carney, D.J., and Clements, P.C., "Assessing the Quality of Large, Software-Intensive Systems", 1995.
- [Buckeridge et. al. 2002] Buckeridge, DL, Graham, J, O'Connor, MJ, Choy, MK, Tu, SW, Musen, MA., "Knowledge-Based Bioterrorism Surveillance", American Medical Informatics Association Symposium, San Antonio, TX, November 9-13, 2002.
- [Burton & Obel 1995] Burton, Richard and Obel, Borge, "The Validity of Computational Models in Organization Science: From Model Realism to Purpose of the Model", *Computational and Mathematical Organization Theory*, 1(1): 57-72.
- [Carley 1995] Carley, Kathleen, "Computational and Mathematical Organization Theory: Perspective and Directions", *Computational and Mathematical Organization Theory*, 1(1): 39-56.
- [Carley 2001] Carley, Kathleen, "On the Evolution of Social and Organizational Networks", 2001.
- [Carley & Lee 1998] Carley, K.M. & Lee, J., "Dynamic organizations: Organizational adaptation in a changing environment", In J. Baum (Ed.), *Advances in strategic management*, Vol 15 (pp. 267-295). Disciplinary Roots of Strategic Management Research. JAI Press.
- [Carley & Prietula, 1999] Carley, Kathleen & M. Prietula (Eds.), "Computational Organizational Theory", Lawrence Erlbaum Associates, 1999.
- [Christopherson 1987] Christopherson, Susan, "Workforce Flexibility: Implications for Women Workers", Institute for Social Science Research, 1987.
- [Cornell 1990] Cornell, J.A., "Experiments with Mixtures", John Wiley, 1990.
- [Craven et. al. 1999] Craven, M., et. al., "Learning to Construct Knowledge Bases from the World Wide Web", *Artificial Intelligence*, Elsevier, 1999.

- [Dewar et. al. 1993] Dewar, J.A., Builder, C.H., Hix, W.M., Levin, M.H., "Assumption-Based Planning: A Planning Tool for Very Uncertain Times", RAND Corp., MR-114-A, 1993.
- [Epstein & Axtell, 1997] Epstein, Joshua & R. Axtell, "Growing Artificial Societies", 1997.
- [Fischer & Schumann, 2002] Fischer, Bernd and Schumann, Johann, "Automated Generation of Statistical Data Analysis Programs", RIACS & NASA Ames, 2002.
- [Fischhoff et. al. 1984] Fischhoff, B., et. al., "Acceptable Risk: A Critical Guide", Cambridge University Press, 1984.
- [Gamerman 1997] Gamerman, Dani, "Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference", CRC Press, 1997.
- [Gilbert et. al. 2000] Gilbert, R. J., et. al., "Genomic computing: explanatory modelling for functional genomics", in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)* (ed. D. Whitley, D. Goldberg, E. Cantú-Paz, L. Spector, I. Parmee and H.-G. Beyer), pp. 551-557. Morgan Kaufmann, San Francisco.
- [Glymour, Scheines & Spirtes 2001] Glymour, C., Scheines, R., & Spirtes, P, "Causation, Prediction, and Search", MIT Press, 2001.
- [Heckerman 1995] Heckerman, David, "A Tutorial on Learning with Bayesian Networks," Technical Report MSR-TR-95-06, Microsoft Research, 1995.
- [Hudak et. al. 2002] Hudak, J., et. al., "Model-Based Verification: Abstraction Guideline", CMU Software Engineering Institute, CMU/SEI-2002-TN-011, 2002.
- [Jackson 1999] Jackson, Peter, "Introduction to Expert Systems, 3e", Addison Wesley Longman, 1999.
- [Jensen & Lesser 2002] Jensen, David and Lesser, Victor, "Social Pathologies of Adaptive Agents", in *Safe Learning Agents: Papers from the 2002 AAAI Spring Symposium.*, Volume TR SS-02-07. August, 2002.
- [John 1998] John, P.W.M., "Statistical Design and Analysis of Experiments", SIAM, 1998.
- [Keppel 1991] Keppel, G., "Design and Analysis: A Researcher's Handbook", Prentice-Hall, 1991.
- [Kilmer et. al. 1994] Kilmer, R.A., Smith, A.E., and Shuman, L.J., "Neural networks as a metamodeling technique for discrete event stochastic simulation," in *Intelligent Engineering Systems Through Artificial Neural Networks*, Volume 4, ASME Press, 1994, 1141-1146.
- [Koza et. al. 1999] Koza, J.R., Bennett F.H., Andre, D., and Keane, M.A., "Genetic Programming III: Darwinian Invention and Problem-Solving", Morgan Kaufmann Publishers, 1999.
- [Linger et. al. 2002] Linger, R.C., et. al., "Life-Cycle Models for Survivable Systems", CMU Software Engineering Institute, CMU/SEI-2002-TR-026, 2002.
- [Liu & Rubin 1996] Liu, C. and Rubin, D.B., "Markov-Normal Analysis of Iterative Simulations before Their Convergence", *Journal of Econometrics*, 1996, **75**, 69-78.
- [Lomi & Larsen, 2001] Lomi A. & Larsen E., "Dynamics of Organizations: Computational Modeling and Organization Theories", MIT Press, 2001.
- [Mitchell 1997] Mitchell, T.M., "Machine Learning", McGraw Hill, 1997.
- [Morgan & Henrion 1990] Morgan, M.G and Henrion, Max (Eds.), "Uncertainty : A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis", Cambridge University Press, 1990.
- [Murphy 2002] Murphy, Kevin, "Dynamic Bayesian Networks: Representation, Inference and Learning", PhD Thesis, UC Berkeley, Computer Science Division. July 2002.
- [Neal 1993] Neal, R.M., "Probabilistic Inference using Markov Chain Monte Carlo Methods", Technical Report CRG-TR-93-1, Dept. of Computer Science, University of Toronto, 1993.

- [Neller 2002] Neller, Todd, "Action-Based Discretization for AI Search", in S.Marshall, ed., *Proceedings of the Game Developers Conference 2002* (GDC 2002, San Jose Convention Center, San Jose, California, USA, 2002), CMP United Business Media LLC, 2002.
- [Oliver 1986] Oliver, Melvin L., "Beyond the Neighborhood: The Spatial Distribution of Social Ties in Three Urban Black Communities", Institute for Social Science Research, 1986.
- [Payne et. at. 2002] Payne, T.R., Singh, R., and Sycara, Katia, "Calendar Agents on the Semantic Web." *IEEE Intelligent Systems*, Vol. 17(3), pp. 84-86, May/June 2002.
- [Pearl 2000] Pearl, Judea, "Causality: Models, Reasoning, and Inference", Cambridge University Press, 2000.
- [Pereira et. al. 1999] Pereira, F.B., et. al., "Graph Based Crossover - A Case Study with the Busy Beaver Problem", in W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela and R. E. Smith (Eds.), *Procs. of the Genetic and Evolutionary Computation Conference, GECCO-99*, Orlando, Florida, USA, pp. 1149-1155, Morgan Kaufmann, July, 1999.
- [Peshkin et. al. 2000] Peshkin, L., Kim, K.-E., Meuleau, N, and Kaelbling, L.P, "Learning to Cooperate via Policy Search", *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 2000.
- [Prietula, Carley & Gasser 1998] Prietula, M.J., Carley, K.M. & Gasser, L. (Eds), "Simulating organizations: Computational models of institutions and groups", Menlo Park, CA: AAAI Press / The MIT Press.
- [Schneider & Riesenhuber 2002] Schneider, R., and Riesenhuber, M., "A Detailed Look at Scale and Translation Invariance in a Hierarchical Neural Model of Visual Object Recognition", MIT AI Lab, AIM-2002-011, August 2002.
- [Shaw & Fleming 1996] Shaw, K.J. and Fleming, P.J., "An Initial Study of Practical Multi-Objective Production Scheduling Using Genetic Algorithms", in *Proceedings of International Conference on Control '96*, University of Exeter, September 2nd - 5th, 1996.
- [Spears 1996] Spears, William M., "Simulated Annealing for Hard Satisfiability Problems", in *Cliques, Coloring, and Satisfiability: Second DIMACS Implementation Challenge*, David S. Johnson and Michael A. Trick (eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Volume 26, American Mathematical Society, 533-558.
- [Sycara 1998] Sycara, K., "Multiagent Systems," in *AI Magazine* Vol. 19, No. 2. Intelligent Agents, 1998: 79-92.
- [Sycara et. at. 2002] Sycara, K., Widoff, S., Klusch, M., and Lu, J., "LARKS: Dynamic Matchmaking Among Heterogeneous Software Agents in Cyberspace." *Autonomous Agents and Multi-Agent Systems*, 5, 173-203, 2002.
- [Tian & Pearl 2001] Tian, J. and Pearl, J., "A general identification condition for causal effects", in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2002.
- [Tian & Pearl 2002] Tian, J. and Pearl, J., "Causal Discovery from Changes", in *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2001.
- [Wasserman & Faust, 1994] Wasserman, S. and Faust, K., "Social Network Analysis: Methods and Applications", Cambridge University Press, 1994.
- [Werfel 2002] Werfel, J., "Implementing Universal Computation in an Evolutionary System", MIT AI Lab, AIM-2002-010, July 2002.
- [Whalen, Schumann, Fisher, 2002] Whalen, Michael, Schumann, Johann, and Fischer, Bernd, "Synthesizing Certified Code", RIACS Technical Report, March 2002.
- [Zebulum et. al. 2002] Zebulum, R.S., Pacheco, M.A.C., and Vellasco, M.M.B.R., "Evolutionary electronics: automatic design of electronic circuits and systems by genetic algorithms", CRC Press, 2002.