The Design of Automated Validation and Explanation for Large-Scale Social Agent Systems

Alex Yahja and Kathleen M. Carley

Institute for Software Research International, School of Computer Science
Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
ay@cmu.edu, kathleen.carley@cmu.edu

Abstract

Socio-technical problems, such as how smallpox outbreaks would spread in and affect modern societies, often have complex interrelated parts that defy simple mathematical analyses. A promising toolkit to solve these problems is large-scale multi-agent models, whose subsets with stochastic and knowledge-intensive networked interactions are social agent models. The value of these models and their simulations increases significantly if they can effectively exploit existing data-streams and knowledge for validation and explain emergent behaviors. Most of the existing technology for validating computational models is designed for deterministic and/or small scale systems where it is often possible to obtain validation manually or semi-automatically by brute-force. Large-scale social agent systems pose an entirely different set of challenges. Given the size of such systems, the vast quantities and variable quality of empirical data involved, automated validation and explanation approaches are crucial. In this paper, such an approach is described in the design of an automated validation and explanation tool called WIZER that utilizes knowledge-intensive simulation-aided search and inference techniques -- and knowledge-based control of simulation -- capable of principled exploration of the parameter and model space, constrained by empirical data and knowledge. WIZER inference engine is built upon our novel Probabilistic Argumentation Causal System, derived from Probabilistic Argumentation Systems and Causal Analysis.

Contact:
Alex Yahja
Institute for Software Research International
Carnegie Mellon University, Pittsburgh, PA 15213
Tel: 1-412-268-7527
Fax: 1-412-268-6938
email: ay@cmu.edu

Key Words: validation, explanation, social agents, probabilistic argumentation, causal analysis, simulation.

Acknowledgement: This research was supported, in part, by DARPA for work on Scalable Biosurveillance Systems, the NSF IGERT9972762 in CASOS, the MacArthur Foundation, and by the Carnegie Mellon Center on Computational Analysis of Social and Organizational Systems. The computations were performed on the National Science Foundation Terascale Computing System at the Pittsburgh Supercomputing Center. Any opinions, findings, 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, the Pittsburgh Supercomputing Center, the MacArthur Foundation, or the US Government.

The Design of Automated Validation and Explanation for Large-Scale Social Agent Systems

Alex Yahja and Kathleen M. Carley

"For we are lovers of the beautiful in our tastes and our strength lies, in our opinion, not in deliberation and discussion, but that knowledge which is gained by discussion preparatory to action. For we have a peculiar power of thinking before we act, and of acting, too, whereas other men are courageous from ignorance but hesitate upon reflection. And they are surely to be esteemed the bravest spirits who, having the clearest sense both of the pains and pleasures of life, do not on that account shrink from danger." From Pericle's Funeral Oration in Thucydides' History of the Peloponnesian War). The Athenians however did not pursue vigorous quantitative approaches to properly debate and assess the matters at hand before action. There was a lack of probability theory, causal analysis, computational modeling & analysis, and proper validation of models.

More than two millennia later, we are experiencing a major change in how we think about individuals, networks, organizations, and other societal systems due to the developments in computational analysis [Axelrod 1997][Carley and Prietula 1999][Epstein and Axtell 1996][Lin and Carley 2003][Prietula 1998]. Recent years have seen a rapid increase of the use of multi-agent models to address complex socio-technical problems of societal systems and model assessment becomes a major concern. The assessment includes validation and explanation. Large-scale social agent systems pose a challenge for validation technologies. Given the size of such systems and the vast diversity of empirical data and knowledge involved, automated validation approaches are crucial.

This paper describes the design of such an approach by tightly integrating simulation with inference engine, along with the simulation explanation facility, called WIZER, for ***W**hat-**I**f Analy**zer***. This tool will be described in the context of BioWar, a bioterrorism impact simulator on a city [Carley et. al. 2003].

Validation Experience

In validating BioWar outputs with hand-crafted parametric studies, the complexity of ensuring "correct" results of agent-based simulation models was evident. Behaviors in BioWar are driven by stochastic processes whose parameters are given as input to each simulation. These parameters number in the dozens, hence varying all of them through the entire parameter space is all but impossible. We have thus consistently chosen just three or four variable parameters which have the largest impact on simulation results. Even with these few variables, each parametric study needs careful planning and execution. The lessons learned of our validation experience include:

1. We need sophisticated analysis and response techniques to optimize the space over which parameters must be varied for correctness, and thus increase the number of parameters which can be studied.
2. We need new approaches to simulation scaling so as to reduce the size of the simulations which produce validated output streams.
3. We need tools to semi-automatically create and execute parametric studies to minimize the manual intervention currently required to perform these studies.

WIZER will address points (1) and (3).

WIZER

WIZER is a tightly-coupled inference and simulation engine that extends response surface methodology [Myers and Montgomery 2002][Neddermeijer 2000][Gaston and Walton 1994] to deal with high dimensional, symbolic, stochastic, emergent [Rasmussen and Barrett 1995], and dynamic nature of large-scale social agent systems. It extends response surface methodology by performing knowledge-intensive data-driven search steps via an inference engine constrained by simulation, instead of just doing statistical and mathematical calculations. WIZER forms both knowledge-based control of simulation and simulation-assisted inferences, enabling reasoning about simulations and simulation-assisted reasoning. It facilitates the management of model assumptions, contradictory or incomplete data, and increases the speed & accuracy of model validation and analysis. WIZER dataflow is shown in Figure 1.

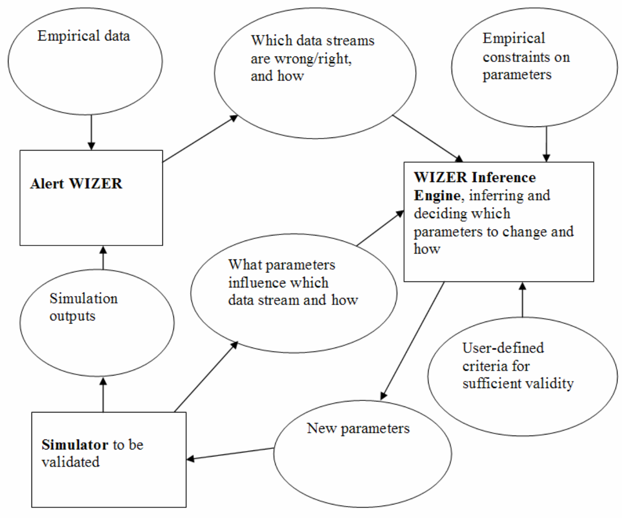

Figure 1. WIZER Dataflow

As shown, WIZER includes Alert WIZER and WIZER Inference Engine. Alert WIZER determines which data streams are wrong and how they are wrong. WIZER Inference Engine takes the simulator's influence diagram of what parameter influences which data and the empirical constraints and confidences on parameters to make a judgment on which parameters to change and how. This results in new parameters for the next simulation. This simulation in turn yields new outputs which are fed into Alert WIZER. A local response surface analysis is embedded, allowing simple virtual experiments for parametric study. This cycle repeats until sufficient validity is achieved based on user-defined criteria. Figure 2 shows the coupling between the expert system and the simulation system of WIZER.

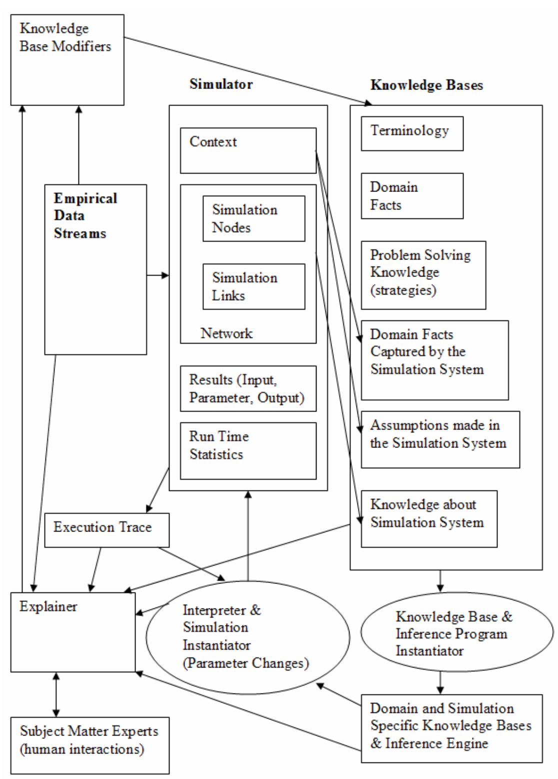

Figure 2. Components of the Coupled Simulation and Expert System

As shown, the knowledge bases will be populated with the knowledge about the simulator, simulation outcomes, domain facts, assumptions, and problem solving strategies, among others. The knowledge bases will contain both knowledge (hard or certain rules and facts) and assumptions (soft or uncertain rules and facts). The assumption part is connected to the simulation so that its outcome will provide justification for the degree-of-support the assumption has. The explainer, as shown in the figure, takes the execution trace, the knowledge base, derived domain & simulation knowledge, and simulation instantiations (and possibly the empirical data streams) to provide explanations for subject matter experts. Subject matter experts could modify the knowledge bases assisted by the explainer and knowledge base modifiers. WIZER conducts inferences from the results of the application of statistical tests, which it has knowledge about and control.

WIZER Inference Engine

The basic design of WIZER Inference Engine is based upon a rule-based Probabilistic Argumentation Systems [Haenni et. al. 1999] for handling assumptions, simulation experiments for estimating and/or testing probability distribution of an assumption, and knowledge-base modifier & explainer for generating & modifying knowledge bases and explaining inferences. While the basic design is sufficient if knowledge engineers are able to check the causal relations [Pearl 2003][Pearl 2000] inherent in some rules, for big knowledge bases manual checks are cumbersome and prone to errors. Thus there is a need for automated and formal causal checking. This is addressed by our novel probabilistic argumentation causal system (PACS), which uses the probabilistic argumentation in causal analysis. Users of WIZER will be able to specify which rules are causal in nature and WIZER will also be able to suggest and do causal checks on some rules based on domain knowledge -- this can also be based on empirical data, but causal discovery from data [Heckerman et. al. 1997] is not the focus of this work. The advanced design of WIZER addresses this causality issue and includes both the rule-based and causality inferences.

Basic Design

While classical logic as it is cannot handle, represent and compute numerical uncertainty (as it represents Boolean logic only), it can be extended to handle the uncertainty by adding a certain type of propositional symbols called assumptions. With these assumptions, uncertain facts and rules can be modeled by allowing facts and rules to be true if specific assumptions are true. Letting P*n* be proposition *n*, Table 1 shows the extension.

Table 1. Representing Uncertainty using Assumptions

Type of Knowledge	Logical Representation	Meaning
A fact	P1	P1 is true
A rule	P1 => P2	P1 implies P2
An uncertain fact	a1 => P1	If assumption a1 is true, then P1 is true
An uncertain rule	a2 => (P1 => P2) or equivalently P1 $\wedge$ a2 => P2	If assumption a2 is true, then P1 implies P2

The probability is simply assigned to an assumption. For example, if for the uncertain rule P1 $\wedge$ a1 => P2, the assumption a1 is known to hold with probability 0.4, then we may write *prob*(a1) = 0.4. Note that this is conceptually different from *prob*(P1 => P2) = 0.4.

The symbolic support for hypothesis *h*, denoted *sp(h, K)*, can be computed, which contains the disjunction of all symbolic arguments which allow to derive *h* if added to the knowledge base *K*. The reliability or degree of the support given by arguments is computed using the probabilities assigned to assumptions: *dsp(h, K) = prob(sp(h, K))*. Evidence against a hypothesis *h*, or reasons to doubt about *h*, is defined as the disjunction of all arguments supporting *~h* and not supporting *h*, and is denoted by *db(h, K)*. The inverse of doubt, the plausibility of *h* is defined as *pl(h, K)=~db(h, K)*. The degree of plausibility is thus *dpl(h, K) = prob(pl(h, K))*.

The above described Probabilistic Argumentation Systems (PAS), which subsumes Assumption-based Truth Maintenance Systems, Bayesian Networks, classical probability theory, and Dempster-Shafer theory of evidence [Haenni 2001][Haenni et. al. 1999]. The inference computation in Probabilistic Argumentation Systems is mainly the computation of quasi-supports based on resolution and variable elimination [Haenni 2001][Haenni et. al. 1999].

Note that we only have the probability number, not the probability distribution nor the procedure to arrive at the probability distribution. WIZER extends the Probability Argumentation Systems by:

- Providing a mechanism to arrive at probability distributions using simulation experiments driven by empirical data.

- Allowing an enhanced inference by letting the inference engine run simulations (e.g., response surface analyses) in the middle of its inferences if needed.
- Allowing what-if tests using simulation experimentations
- Bounding the inferences (the number of arguments) by simulation system, real data, domain knowledge, and problem-solving knowledge.
- Elucidating types of uncertainty affecting resulting probability distributions
- Integrating knowledge-base modifier and inference explainer

while preserving the classical logic of PAS.

Advanced Design

The basic design as described earlier is based on rule based systems. It hence has an unfortunate artifact of producing incorrect inferences if knowledge engineers do not take special precautions in encoding the rules. This artifact is easily shown by the following incorrect inference from two correct rules using a correct inference mechanism (chaining):

Rule 1: *If the lawn is wet, then it rained*
Rule 2: *If we break the water main, then the lawn gets wet*
Inference: *If we break the water main, then it rained*

Thus there is a need to explicitly represent causality – which includes representing actions instead of just observations and formally addressing confounding -- which has fortunately been formalized and mathematized [Pearl 2003][Jewell 2003]. Incorporating causality would enable users to make adjustment to the above rules:

Cause 1: *Raining caused the lawn to be wet*
Cause 2: *Breaking the water main causes the lawn to be wet*
Inference: *None*

As can be seen, Rule 1 was encoded erroneously if causal relations are to be taken into account. While erroneous in its cause-effect relation, Rule 1 can still be useful as a suggestion. Thus we needed both the rule-based and the causal inferences, so this advanced design covers the basic design described earlier.

Causal analysis involves causation (encoding of behavior under interventions), interventions (surgeries on mechanisms), and mechanism (stable functional relationships). Causal model consists of actions (*B* will be true if we do *A*), counterfactuals (*B* would be different if *A* were true), and explanation (*B* because of *A*). Counterfactuals do not address necessity (ignoring aspects of sufficiency and failing in presence of other causes) and coarseness (ignoring structure of intervening mechanisms and failing when other causes are preempted), but they are remedied with sustenance [Pearl 2003][Pearl 2000]. Sustenance methods include causal beam, temporal preemption & dynamic beams. All this machinery allows us to specify *A* ***caused*** *B*.

If it were to account for probability of causation, this causal model as it is [Pearl 2003][Pearl 2000] specifies Bayesian priors to encode the probability of an event given another event. It does not distinguish between different kinds of uncertainty. It is unable to model ignorance, ignores contradictions, and is inflexible in expressing evidential knowledge without the use of probability distribution format. With the intended use of WIZER to do validation with often incomplete, contradictory, and uncertain environments & knowledge and with the need to clearly delineate between assumptions and facts, it needs a novel causal model borrowing concepts from Probabilistic Argumentation Systems (PAS). Table 2 shows the encoding of the certain facts and causations for causal analysis augmented with PAS-like assumption management. Let P*n* be proposition *n*.

Table 2. Assumptions Encoding for Causality

Type of Knowledge	Logical Representation	Meaning
A fact	P1	P1 is true
A causation	P1 caused P2	P1 caused P2
An uncertain fact	a1 => P1	If assumption a1 is true, then P1 is true
An uncertain causation	a2 => (P1 caused P2)	If assumption a2 is true, then P1 caused P2

We call the above the probabilistic argumentation causal systems (PACS). The causal inference formalism is based upon sustenance's causal and dynamic beams with temporal preemption in causal graph, with the ability to specify ignorance, degree of support, degree of plausibility, and hints. PACS algorithmic details are derived from both PAS

[Haenni et. al. 1999] and causal analysis [Pearl 2003]. The PACS algorithm is basically a construction of causal (& dynamic) beams in causal graph followed by the application of PAS, but more advanced version is possible, e.g., the degree that ignorance, hints, degree of support, and degree of plausibility may affect the construction of causal & dynamic beams. Simulation experiments can be seen as a proxy of doing interventions and counterfactual/sustenance reasoning to derive causal relations, when doing real world interventions are unrealistic or unethical.

In addition to creating our novel probabilistic argumentation causal systems (PACS) just described, WIZER enhances PACS by

- Providing a mechanism to arrive at probability distributions or profiles for assumptions related to causations using simulation experiments driven by empirical data.
- Automating causal analysis for simulations and enhancing it with virtual experiments. In particular, WIZER improves upon dynamic beams of causal analysis by doing virtual experiments, and allows the estimation of sufficiency by virtual experiments.
- Utilizing simulation experiments as a proxy of the real world for doing interventions & causal beam calculations to uncover true causal relations driven by empirical and simulated data. Simulation model is analogous to imperfect electronic diagrams. Validated experiments allow us to defend and "validate" modeling assumptions of causal analysis.
- Allowing better inference by letting the inference engine run simulations (e.g., supports for dynamic beam) in the midst of causal inferences as needed. This is not just a nice feature, but serves a function of checking the empirical claims of causal relations, as all causal claims are empirical.
- Bounding the causal inferences by simulations, real data, domain knowledge, and problem-solving knowledge.
- Elucidating types of uncertainty affecting resulting probability profiles (points, sets, or distributions)
- Integrating causal knowledge-base modifier and causal inference explainer

Preliminary Results

We have implemented the alert part of WIZER, the Alert WIZER, which performs checks on the ranges and the means of simulated output data streams for BioWar. It was run on "Challenge 4" data of BioWar, which has twelve data streams, including school absenteeism, work absenteeism, doctor visits, emergency room visits, emergency room visits using the Surveillance Data Inc. data, and seven drug type purchase data streams. The table below shows the percentage of validated data streams for six cities for no attack case.

Table 3. Percentage of Streams Validated

City	Streams Validated
San Francisco	5/12 = 41.67%
San Diego	7/12 = 58.33%
Pittsburgh	7/12 = 58.33%
Norfolk	6/12 = 50.00%
Hampton	4/12 = 33.33%
Washington DC	4/12 = 33.33%

Conclusions

The paper described an automated tool to enable the human analyst to have more trust in a simulation system which WIZER assists to automatically validate and provide explanations. The human analyst can see his or her domain knowledge in action in simulations, thus freeing the analyst to think about problems and be innovative. WIZER will assist human experts by doing inferences of the simulation systems based on ranges of values and assumptions provided by the human experts. The human experts would be able to specify both rules and causations. As WIZER can integrate varied knowledge-based, the expertise resided in many human expert domains and silos can be readily employed into one single simulation system and WIZER would show the consequences and explanations. WIZER would be able to search for appropriate values, assumptions, rules, and causations, by doing what-if analyses in a simulation-assisted inference engine bounded by empirical data. Automating inferences, simulation experimentation and explanation, WIZER would cut the time and resource needed to do analyses using simulation systems.

References

Axelrod, R. (1997), Advancing the Art of Simulation, International Conference on Computer Simulation and the Social Sciences, Cortona, Italy.

Carley, K. M. and M. Prietula (eds.) (1999), Computational Organizational Theory, Lawrence Erlbaum Associates.

Carley, K. M., D. Fridsma, E. Casman, N. Altman, J. Chang, B. Kaminsky, D. Nave, and A. Yahja (2003), BioWar: Scalable Multi-Agent Social and Epidemiological Simulation of Bioterrorism Events, 2003 NAACSOS conference proceedings, Pittsburgh, PA

Epstein, J. M. and R. L. Axtell (1996), Growing Artificial Societies, MIT Press.

Gaston, G. J. and A. J. Walton (1994), The Integration of Simulation and Response Surface Methodology for the Optimization of IC Processes, IEEE Transactions on Semiconductor Manufacturing, Vol. 7, Issue 1, page 22-33

Haenni, R. J. Kohlas, N. Lehmann (1999), Probabilistic Argumentation Systems, Technical Report 99-9, Institute of Informatics, University of Fribourg.

Haenni, R. (2001), Cost-bounded Argumentation, International Journal of Approximate Reasoning, 26 (2), 101-127.

Heckerman, D., C. Meek, and G. Cooper (1997), A Bayesian Approach to Causal Discovery, in C. Glymour and G. Cooper, editors, Computation, Causation, and Discovery, pages 141-165, MIT Press, Cambridge, MA, 1999. Also appears as Technical Report MSR-TR-97-05, Microsoft Research, February 1997.

Lin, Z. and K. M. Carley (2003), Designing Stress Resistant Organizations: Computational Theorizing and Crisis Applications, Kluwer Academic Publishers.

Neddermeijer, H. G., G. J. van Oortmarssen, N. Piersma, D. Dekker (2000), A Framework for Response Surface Methodology for Simulation Optimization, Erasmus University Rotterdam, Econometric Institute, Econometric Institute Report No. 192.

Myers, R. H. and D. C. Montgomery (2002), Response Surface Methodology: Process and Product Optimization using Designed Experiments, 2nd ed., John Wiley & Sons, Inc.

Pearl, J. (2000), Causality: Models, Reasoning, and Inference, Cambridge University Press.

Pearl, J. (2003), Statistics and Causal Inference: A Review, in Test Journal, Vol. 12(2), pp. 281-345, December 2003.

Prietula, M. J., K. M. Carley, and L. Gasser (1998), Simulating Organizations, AAAI Press/The MIT Press

Rasmussen, S. and C. L. Barrett (1995), Elements of a Theory of Simulation, ECAL 95, Lecture Notes in Computer Science, Springer Verlag.