

Aligning Simulation Models of Biological Attacks

Li-Chiou Chen¹, Kathleen M. Carley^{1,2}, Boris Kaminsky¹, Tiffany Tummino²,
Elizabeth Casman², Douglas Fridsma³, Alex Yahja^{1,2}

¹Institute for Software Research International, School of Computer Science

²Department of Engineering and Public Policy

Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

{lichiou, borisk, ttummino, carley, casman}@andrew.cmu.edu

³Center for Biomedical Informatics, School of Medicine, University of Pittsburgh

Pittsburgh, Pennsylvania 15213

fridsma@cbmi.pitt.edu

Abstract

We aligned two fundamentally different models of disease transmission after a bioterrorist attack: a multi-agent model (BioWar) and the conventional Susceptible-Infected-Recovered (SIR) model. The purpose of this alignment is part of a greater validation process for BioWar. We conducted two model alignment studies based on smallpox and anthrax attack simulations. From these two studies we were able to show that, at the minimum, the epidemiological curves produced by the two models were approximately equivalent, both in overall and the time course of mortality. Subtle differences on the model results revealed the impact of heterogeneous mixing in the spread and the progression of a disease. Based on this foundation, we will be able to further investigate the policy responses against the biological attacks by improving heterogeneous properties of agents, which cannot be simulated in a SIR model.

Contact:

Li-Chiou Chen

Institute for Software Research International

Carnegie Mellon University

Pittsburgh, PA 15213

Tel: 1-412-268-7527

Fax: 1-412-268-6938

Email: lichiou@andrew.cmu.edu

Key Words: docking, model alignment, biological surveillance, biological attacks, computer 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.

Aligning Simulation Models of Biological Attacks

Li-Chiou Chen, Kathleen M. Carley, Boris Kaminsky, Tiffany Tummino, Elizabeth Casman, Douglas Fridsma, Alex Yahja

Numerical simulation models can be used to estimate the impact of large-scale biological attacks and to design or select appropriate response strategies. The “correctness” of the model is critical since the “wrong” model may lead to “wrong” decisions, but no model is perfect and few models can ever be considered thoroughly validated. Studies [Sargent 1984 & 2003] have agreed that it is often too costly and time-consuming to determine if a model is absolutely valid. Instead, evaluations are conducted until sufficient confidence is obtained that a model is valid for its intended application. We developed a methodology to align an agent-based model of biological attack simulations (BioWar) against the classical susceptible-infected-recovered (SIR) box model as part of the validation process. Our purpose is to verify whether the agent-based model can produce results that closely resemble those of the well accepted and venerable SIR model, thus giving BioWar a sort of reflected credibility from the SIR model. This is not sufficient validation, but it is a confidence building step in the much larger task of validating BioWar [Carley et. al 2003].

Model alignment [Axtell et. al 1996], also referred to as “docking,” is the comparison of two computational models to see if they can produce equivalent results. Properly done, model alignment can uncover the differences and similarities between models and reveal the relationships between the different models’ parameters, structures, and assumptions. By aligning a complex new model with a simpler and well-understood model, one can obtain a sense of validity needed to develop the new model. The same technique has been used previously to validate a model of organization performance [Louie et. al 2003]. We conducted two studies to align revised SIR models and BioWar based on smallpox attack simulations [Chen et. al 2004] and anthrax attack simulations [Chen et. al 2003]. This paper summarizes our method, results and lessons learned from the two studies.

The Two Models

BioWar is a simulation tool that combines computational models of social networks, communication media, disease models, demographically resolved agent models, spatial models, wind dispersion models, and a diagnostic model into a single integrated system that can simulate the impact of a bioterrorist attack on a city. In BioWar analysts can model real cities using census, school district demographics, and other publicly available information. When a biological attack occurs, those in the vicinity of the release may become infected, following probabilistic rules based on received dose and age of the agent. The infected agents modify their behaviors as their disease progresses and they become unable to perform their normal functions as the disease worsens. A detailed description of the model along with a plan for validation can be found in [Carley et. al 2003].

Disease processes and response strategies are traditionally modeled by the susceptible-infected-recovered (SIR) model. The SIR model and its variations have been widely used to model the spread of epidemics and to study immunization strategies [Anderson and May 1992]. The SIR model is a “population-based” description of disease progression processes that assume homogeneous mixing of individuals. The model categorizes the entire population into three states: susceptible (S), infected (I) and recovered (R). All members of a particular state are identical and have predefined transition probabilities of moving to another state in the model. Although variations in the way in which the disease is manifested and symptom based behaviors can be tracked using Monte Carlo simulation methods, the interaction among population members is often lost.

In contrast, the agent-based BioWar takes a different approach thus allowing us to model the complex social interactions absent in most SIR models. However, in order to understand the benefits and limitations of using BioWar to model biological attacks, we aligned BioWar with population-based models revised from the SIR model.

In principle, agent-based models have the advantage that the heterogeneity of individual response can be accounted for, thus enabling a finer grained analysis and allowing the tools to be used for training and intelligence purposes. In BioWar, a further advantage is that the diseases are modeled at the symptom level thus enabling the model to contribute to our understanding of the ways in which early symptomatic based behaviors, such as the purchase of the over-the-counter-drugs are likely to emerge after a biological attack. Further, by using a general symptom based framework, new diseases and even “unheard of” diseases can be rapidly modeled in BioWar. Additionally, in BioWar, multiple diseases are simultaneously tracked so that disease interactions can be examined.

The Process of Model Alignment

The detail steps for aligning two models based on smallpox simulations and anthrax simulations are slightly different because smallpox is contagious and anthrax is not. The processes of the model alignment in the two studies are similar. In general, our model alignment approach includes four processes: qualitative comparison, parameter alignment, design of simulation scenarios and comparison of population level results.

Qualitative comparison

Qualitatively speaking, the differences between SIR and BioWar can be summarized as follows:

- Population assumptions: SIR models population cohorts as they transition through different disease states. BioWar models heterogeneous individuals and their interaction in social settings.
- Disease model design: SIR simulates the disease progression from a macro point of view. That is, the model uses a state machine to describe the state changes among sub-populations and uses proportional state transition probabilities to describe the migration of sub-populations. BioWar simulates the emergent properties of individual agents from a micro point of view. That is, to describe the population level disease status, BioWar models and summarizes the disease state of the individual agents. The macro behavior of the population emerges from the outcomes for the individual agents.
- Computational process: To generate the prevalence of a disease over time, the BioWar model requires more computational power than does SIR. In addition to tracking the maliciously introduced infection in exposed agents, BioWar models behaviors and information used in early detection algorithms as well as health status information.
- Initialization: BioWar is initialized with information that describes individual differences. SIR requires initial state characterization and state transition probabilities of the population. The entire population is divided into several sub-populations according to the disease stages.
- Parameterization: While SIR takes both the exposed population and infected population as inputs, BioWar can calculate them as emergent properties from simulating parameterized attacks.

Parameter alignment

Both BioWar and SIR simulate disease progression in terms of the transition of infected individuals between disease stages, but with different stochastic framing. In BioWar, we assume a probability distribution for the duration of a disease stage for each agent. In SIR, a fraction of agents move from one disease stage to another governing by state transition probabilities. Due to this reason, we can only compare the average case from the two models by aligning the mean values of the probability distributions in BioWar with the state transition probabilities in SIR. In the smallpox attack simulations, we have to align the reproduction rate since smallpox is contagious. The reproduction rate is an emergent property in BioWar but an input parameter in SIR.

Design of simulation scenarios

In the anthrax attack study, we simulated a mass anthrax attack scenario over the municipal stadium in Hampton Roads, VA. The total population is around 142,000. In the smallpox attack study, we simulated three smallpox scenarios in Washington, DC, scaling down to 10% of its original population. The total population after scaling is around 55,900. The three scenarios include a base scenario with average 7 initial infections, a vaccination scenario and a quarantine scenario.

Comparison of population level results

For both studies, we compared the over time infections and mortality after an attack. For each scenario, we presented the results as averages of 100 runs because the fluctuation of disease reproductive rates is negligible in around 100 runs.

Results and Discussions

From these two studies, we found that only certain aspects of the models could be compared. Because of the different ways the models account for parameter uncertainty, it is necessary to compare average results over numerous runs. Based on these average results, we found that BioWar can generate population level results that are close to a revised SIR model. In the case of anthrax attack simulations, BioWar results are comparable to the two empirical data sets from US mail attack in 2001 and Sverdlovsk outbreak in 1979. Subtle differences exist because of the differences in mixing assumptions.

When the level of detail in a simulation increases, the number of model parameters needed increases. For example, in smallpox simulations, the transmission probability may vary by age group or occupation (such as medical workers, family members of an infected person, or general public). BioWar provides a way to manage these

model parameters in order to represent the heterogeneous properties of individuals. Although we can revise SIR model to simulate the same level of fidelity by dividing the population into several categories, it is not advisable because the number of model parameters would increase nonlinearly to an unmanageable level. In addition, revising SIR to have finer population categories overlooks an important aspect of disease transmission: the fact that the population reproductive rate is actually partly the result of interactions between individuals and these interactions are emergent properties of agent-based models which cannot be generated from the SIR model.

Conclusions

We aligned a multi-agent model of weaponized biological attacks, BioWar, with the classical susceptible-infected-recovered (SIR) model. Using both smallpox attack and anthrax attack simulations, we showed that the average results from BioWar are comparable to the SIR model, when the models are properly parameterized. The key parameters include the average disease-stage durations, the reproductive rate (for smallpox only), the initial infection and the probability of death following infection.

The successful docking of the two radically different models provided a degree of confidence in the agent-based model, showing that its results are not far from those of the established SIR model. This work is our first step of the larger task on validating BioWar. Tools for finer-granularity validation of agent-based models are underway [Yahja 2004]. Based on this foundation, we will further investigate the policy responses against the outbreaks of contagious diseases by changing heterogeneous properties of agents (such as social networks, daily activities, and reactions to an attack), which cannot be simulated in a SIR model.

References

- Anderson, R. M. and May, R. M., *Infectious Diseases in Humans*: Oxford University Press, 1992.
- Axtell, R., Axelrod, R., Epstein, J. M., and Cohon, M. D., "Aligning simulation models: a case study and results," *Computational and Mathematical Organization Theory*, vol. 1, 1996.
- Carley, K.M., Altman, N., Kaminsky, B., Nave, D., and Yahja, A., *BioWar: A City-Scale Multi-Agent Network Model of Weaponized Biological Attacks*, Technical Report (CMU-ISRI-04-101). Pittsburgh, PA: CASOS, Carnegie Mellon University, 2003
- Chen, L-C, Carley, K.M., Fridsma, D., Kaminsky, B., and Yahja, A., *Model Alignment of Anthrax Attack Simulations*, accepted for publication in *Decision Support Systems*, 2003.
- Chen, L-C, Kaminsky, B., Tummino, T., Carley, K.M., Casman, E., Fridsma, D., and Yahja, A., *Aligning Simulation Models of Smallpox Outbreaks*, 2004, *Proceedings of the Second Symposium on Intelligence and Security Informatics*, Tucson, Arizona, June 10-11, 2004. Also to appear in *Springer-Verlag Lecture Notes in Computer Science* as volume 2665.
- Louie, M. A., Carley, K. M., Haghshenass, L., Kunz, J. C., and Levitt, R. E., "Model Comparisons: Docking ORGAHEAD and SimVision," presented at *Proceedings of NAACSOS conference*, Pittsburgh, PA, 2003.
- Sargent, R.G.: *Simulation Model Validation*. in *Simulation and Model-Based Methodologies: An Integrative View*, T.I. Oren, B.P. Zeigler, and M.S. Elzas: Editors. Springer- Verlag: Heidelberg, Germany, 1984, 537-555.
- Sargent, R.G.: *Verification and validation of simulation models*. *Proceedings of the 2003 Winter Simulation Conference*.
- Yahja, A.: *WIZER: Automated Validation of Large-Scale Multi-Agent Systems*. CASOS working paper, Carnegie Mellon University, 2004.