

THE FUTURE OF COMPUTERIZED DECISION MAKING

Bruce G. Elmegreen

IBM Research Division
T.J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598 USA

Susan M. Sanchez

Naval Postgraduate School
Operations Research Department
Monterey, CA 93943-5219 USA

Alexander S. Szalay
Department of Physics and Astronomy
The Johns Hopkins University
3701 San Martin Drive
Baltimore, MD 21210, USA

ABSTRACT

Computerized decision making is becoming a reality with exponentially growing data and machine capabilities. Some decision making is extremely complex, historically reserved for governing bodies or market places where the collective human experience and intelligence come to play. Other decision making can be trusted to computers that are on a path now into the future through novel software development and technological improvements in data access. In all cases, we should think about this carefully first: what data are really important for our goals and what data should be ignored or not even stored? The answer to these questions involves human intelligence and understanding before the data-to-decision process begins.

1 INTRODUCTION

Computers are handling more and more of our everyday data, making numerous tiny decisions behind the scenes that usually help us without our knowing it (e.g., finding credit card fraud), but which occasionally confound and surprise us, reinforcing the old adage that “to err is human but to really foul things up requires a computer” (attributed to William E. Vaughan (1969) in <http://quoteinvestigator.com/2010/12/07/foul-computer/>). With more and more data stored in digital form, some of which is totally irrelevant, or worse, erroneous, and with computers seemingly everywhere churning away on these data whether we want them to or not, we are entering a world where computerized decision making becomes an important reality. How far will this go? How far can it go? How far should it go? These and other questions are addressed in what follows by the three members of our panel.

2 BRUCE ELMEGREEN: WHAT’S TOO HARD TO DO AND HOW DO WE MANAGE NOW?

Decisions about what actions to take in order to achieve a goal require information about the consequences of many possible actions, an evaluation of the good and bad aspects of these consequences and their relative weights, a tally of the total value of each action in terms of a weighted sum of the good consequences minus the bad, and then a judgment about which action to take based on the relative values,

or whether to wait until more information becomes available, the evaluations are more accurate, or the goals become more clear. We take these steps every day using our brains for small and large decisions, although sometimes we have to wait for months or even years for the whole sequence to play out. We also have to act occasionally without going through these steps, and trust that we will be able to adjust our actions or know more about them later.

The first step of acquiring information about the consequences of various actions involves fetching historical data on previous similar actions, or statements and threats made by people explaining how they will react to certain things. This step is well suited to modern computers if the relevant data are available in digital form -- data such as newspaper articles, video, and books -- and can be searched according to topics and keywords, or sounds and images of people and objects (Roy, Faulkner and Finlay 2007).

The second step of evaluation requires some judgment on the consequences of these actions, i.e., good or bad relative to the goal, and this can come again from archival records of similar events, but it may also require knowledge about the trustworthiness of people making statements based on their personal histories. These are again somewhat amenable to computerized data research. However, if the actions or the environment for possible actions are unprecedented, then evaluation may require detailed computer simulations with world-scale ecosystems involving all related parties, markets, and natural processes (e.g., Grabianowski 2012). Because of the complexity of many human events and markets, simulation outcomes could be highly sensitive to the assumptions, input data, and scale of the computation, and therefore inconclusive or time variable.

The sequence of steps in decision making continues to get ever more complex (Krause 1993, Heingartner 2006, Maule 2009). The assignment of weights to the outcomes of actions, whether they are strongly good in favor of the desired goal, or just a little bit good, and how two weakly good consequences might balance one strongly bad consequence, is often a highly complex task for humans and may be even more so for machines. Numerical weights can be assigned for the computation, of course, but humans viewing the outcome of the computerized decision may disagree with those weights or their implications because of some intangible feeling. "Yes, but it's not that simple," might be a common response to an attempt at computerized decision making. Also, people and parliaments are somewhat unpredictable. New political parties could develop on short timescales, or persuasive books, leading people to shift their own evaluations and weights for activities and consequences. The result would seem to be a fundamental inability to make a single best decision by either human or computer means, and indeed governing bodies usually reserve an option to change their collective minds.

By the time we get to the end of the sequence, there could easily be an enormous jumble of facts and precedents, possibly too many for our minds to contain and evaluate all at once, combined with some intuitive feeling about how certain people or groups will react to the events as they unfold, and even a lingering question about whether we have understood the best possible goal to begin with. How can computers help us put all of this together? One possibility was shown to be effective by the IBM Watson computer that played the television game Jeopardy ([http://en.wikipedia.org/wiki/Watson_\(computer\)](http://en.wikipedia.org/wiki/Watson_(computer))). This is also a strategy planned for future uses of Watson-type question and answer systems, namely to have the computer give probabilities for success against a goal, with explanations available for the facts, weights, and strategies that went into these probabilities. Then humans would view these probabilities and explanations and either make the final decision or decide to wait.

A problem's amenability to computation would seem to depend on its degree of isolation. The computerized reaction of stepping on the brakes in a car when some obstacle is in the way involves a relatively isolated event; at present only the obstacle and the road conditions for braking might matter (Howard 2013), although perhaps more sophisticated models in the future will also consider other nearby cars and their computer reactions to the first car. Decisions that involve many separate parts -- uncountably many in the case of some human endeavors -- are often made today by committees, i.e., using many brains, or ballot, i.e., using a culture's collective intelligence and experience, or market forces, i.e., using many independent binary actions between individual entities. A computer-aided ranking of the

consequences of many possible actions might eventually be evaluated by some combination of these same human assemblies.

3 SUSAN SANCHEZ: STEPS TOWARD A COMPUTER-CENTRIC FUTURE

3.1 Introduction

I believe several aspects of the intersection between big data, simulation, and decision making will be of increasing interest in the near future. Here is a quick summary.

3.2 Future Simulation Clients

Complex Problems: All too often, the elegance and rigor of having a closed-form or mathematically tractable solution have been touted as advantages over a simulation modeling approach. This ignores the introduction of “type III errors” (Mitroff and Featheringham 1974) that occur when we solve the wrong problem. Perhaps this becomes harder to justify in the face of readily available big data. For example, when it's not necessary to task someone to go and collect a lot of information, because that information is already available, it is harder to justify assuming i.i.d. exponential random variates. Increasingly, the lack of a closed-form solution is not an issue when our software is capable of computing results to a desired level of accuracy in a small amount of time; this is called *computational tractability* (Lucas et al. 2014). Climate change, economics, transportation, combat, and social dynamics are just a few of the areas where closed-form analytic models will not suffice---computational models are better at capturing the complexity of the underlying systems.

Complex questions: When clients have complex problems and are studying complex systems, they are not likely to be interested in answers to simple questions. Just as “having” big data from the internet meant that companies found new and exciting things to do with it, we've seen that having big data from simulation experiments offers the opportunity for new and interesting ways of looking at the results. “How should I set up my transportation network?” and “What are the impacts of the affordable care act on health costs and health outcomes?” are much more complicated (and interesting) questions than “What is the expected time-in-system for a customer in an M/M/1 queue with no balking and unlimited buffering?”

Comfort with computerized and computer-based decisions: The current fascination with big data has several secondary effects. The rapid evolution of data science means that a greater number of simulation and non-simulation professionals will be becoming more adept at scripting, modeling, graphical and statistical displays. Decision makers may, similarly, be less likely to shy away from using observational or model-driven data to inform their decisions. At the same time, enhancements to methods for rapidly creating, merging, searching, displaying, and analyzing data from large repositories may prove to be useful new tools for the simulation community. Comfort with computerized and computer-based decisions is also increasing in other ways. If we trust a well-written computer program to drive a car (Jaffe 2014), why not trust a well-written computer program for other types of decisions?

3.3 Future Simulation Methods

Continual processing: Most often, we see examples of simulation studies defined to address a specific question. I believe it is time to view simulation-based decision making as a process, not an end state. Why do we churn up the CPU cycles when we're in the midst of an analysis activity, and then let our computer sit idle for the rest of the time? One intriguing idea is that of going back and forth between models of different types or different fidelities, as we seek to learn more about these systems: this is being done in some scientific computing communities, such as computational physics, and it may have interesting parallels in the discrete-event simulation community. Another approach, even if we begin with a specific question, is to generate more output (in a structured manner) so that we are prepared for

the next round of questions from our model, and we have been able to identify interesting features in the response surface metamodels that might not have initially been apparent.

Changing areas of research emphasis: There has been a great deal of good work in our community on simulation optimization, ranking and selection, and response surface modeling. But these presuppose that the decision maker knows what question they want to ask. I believe that more research is needed on multi-objective procedures, exploitation of parallel computing, adaptive methods, and the design and analysis of large-scale simulation experiments. At the same time, there appear to be new opportunities for some established research areas. Importance sampling becomes potentially more of interest as we pull in real-time data. Can we easily update the state of our simulation, identify branching opportunities, and move forward quickly in parallel? Regarding simulation optimization and other adaptive search techniques, it may be that we should be doing optimization on metamodels, rather than on the simulations themselves---and that we need automated ways of reoptimizing as these metamodels evolve over time.

Causal computerized decision making: As I discuss elsewhere in this proceedings (Sanchez 2014) simulation can be the core for model-driven big data and inferential decision-making. We need to stake this area out. One of the criticisms rightly trumpeted over and over for big data is that “correlation is not causation” and that the results are “descriptive, not prescriptive.” In our field, we deal with prospective decision making. We have an advantage in this area: since our output data are generated from models, we do not have many of the issues of data quality and availability that occur in many real world situations. Unfortunately, simulation is still often viewed as a second-class field within operations research. We have the opportunity of becoming recognized as the gold standard for model-based decision making within the big data analytics community.

3.4 Future Simulation Software

More automation, broader interfaces: Data capture is being automated at an incredible rate from real-world systems, ranging from satellite imagery, to web site navigation, to social network analysis, to engine systems. In the future, I see a growth in automating linkages between real-world data and simulation modeling environments. This increases the potential for using simulation as a real-time decision support and control system, as it has in recent biopharmaceutical applications (Johnston et al. 2008). Major simulation packages may adopt an “app” approach, and use common data exchange protocols and interface protocols to link simulation models with external data sets. If so, the same protocols might also allow the practitioner to easily find or create suitable apps for analysis (e.g., simulation experiments, simulation optimization, ranking & selection, importance sampling) as well as for big data visualization. I expect an increase in the use of adaptive, automated analysis methods.

Simulation as a service: Simulation software developers should start taking more advantage of cloud computing, coupled with the ability to run models remotely via a web interface. Software developers might consider whether there's an analog to a subscription service for running simulation models, rather than licensing software for individual machines or users. At the server side, intelligent resource allocation (“automated data farming”) can take advantage of parallel processor capabilities in stand-alone clusters or in clouds.

Smarter computational agents: This an area that is ripe for improvement. The medical field has a few applications where intelligent software agents search through large data sets and find correlations. These have led to theories (e.g., on environmental or genealogical links to certain diseases later in life) that can then be examined more thoroughly and tested by medical researchers. Can we do the same with simulation? One way is to construct intelligent agents to search through model-driven data sets, identifying important factors and interesting features in the responses. Another way is to embed some of this capability into our models themselves; for example, rather than relying on calls to random variate generators with fixed parameters, we might allow intelligent agents within our simulation model to access near-real-time big data and determine whether or not these distributional models still appear to be valid.

3.5 But Wait, There's More!

I have listed a few changes that I feel are on the horizon, but if there's one thing that the past few decades have taught us, it's that we never know exactly what the future will hold. When the internet got started, we viewed e-mail as a faster alternative to letters, and word-processing as a potential way of cutting down on waste paper---in other words, incremental change instead of revolutionary change. Similarly, when the web got started, we did not envision how this connectedness would change our society. So whatever the future holds, the simulation community should be poised to identify, respond to, and ideally blaze a trail that leverages emerging technologies.

Our simulation community has an important role to play. We have been interested in many of these ideas before they captured the public's attention. Because we have been wrestling with them for years, we already have a rich literature of effective ways to deal with complex problems. If we take steps to push this work out to broader communities, we will help those unfamiliar with the current state-of-the art in simulation avoid reinventing the wheel (or worse, repeating the mistakes of the past). More importantly, we will help jump start the process of improving decisions that may affect our businesses, our lives, and our planet.

4 ALEXANDER SZALAY: FINDING THE DIMENSIONS OF SPARCITY

In science a few years ago we realized the impact of the emerging huge volumes of data, leading to the "Fourth Paradigm" (Hay, Tansley, and Tolle 2009). After empirical, theoretical and computational approaches, we see Data-Intensive Science to appear in every discipline of science. We also see a convergence of physical and life sciences through the same computational technologies. Introduction to data science is rapidly becoming the most enrolled class on many university campuses.

How does this revolution in data analytics impact decision making in science? When we had just a small amount of data, and it took a huge effort to collect even that, it was clearly a human decision on what experiments to do next, what tradeoffs and compromises to make between, cost, complexity, time and increased scientific insight. As we have more and more complex data in our repositories, spanning a large number of dimensions, it is even hard to visualize their relationship, not to mention asking what data to collect next. Yet, given the sophistication of our computerized experiments, like large robotic telescopes, genomic sequencers, supercomputers running large simulations, it is relatively easy to think of new experiments to do. In fact, it is all too easy to collect large amounts of new data. Of course, scientists, when faced with the question: "do I have enough data, or would I like to have more?" have rarely opted to stop acquiring new data! This is the phenomenon that is leading to an exponential growth of scientific data – it is doubling every year. Continuing this trend is rapidly becoming untenable, we cannot keep buying more disk drives and more computers to analyze them.

The hard question we are faced with, how can we collect more *relevant data*! A lot of scientific phenomena are based upon rather simple underlying causal relationships that we seek to find from the complex correlations detected by our data analytics tools. These simple relationships mean that the underlying models for the phenomena are rather *sparse*, in the right (but unknown) space they can be described with a small number of parameters. Finding the ideal transformation of the data into this simplest representation is an NP-hard problem. However, over the last decade many of the world's top mathematicians discovered that approximate but fast solutions to this problem are possible, and this led to compressive sensing, the mathematic theory of explicitly including our knowledge about the simplicity, or sparseness, of the phenomenon.

Following this train of thought, one can also ask that given these assumptions, and all of our existing data, which one of the many possible experiments we can perform will lead to the greatest incremental growth in our knowledge, or the ability to reconstruct the underlying signal?

Today, at best humans make computer aided decisions. But, this is turning upside down rapidly. The future is clearly human aided decision making by computers. In 2004, Ross King of Manchester (King et

al. 2004) has published the first successful demonstration of this idea, and applied it to drug design, and has carried it to build Adam, the Robot Scientist. It is obvious that we can find many areas of science today where such principles can be applied. In materials science, it is impossible even to simulate all the possible combinations of elements to build alloys with certain specific targeted properties. Using a machine learning approach to make conscious tradeoffs and decide which experiments (or computations) should be performed is clearly the only way to go.

In astronomy, it is much easier to collect photometric information about distant galaxies (multicolor imaging). But, if we want to infer the detailed physical properties of these objects we need to take high resolution spectra. This latter process is many orders of magnitude less efficient. Taking spectra of faint galaxies requires a large telescope and a long integration. Based upon the observed photometric properties of the object in our sample we need to decide which objects to observe spectroscopically, given the enormous cost of each spectrum.

In medical practice soon computers will make the decisions what diagnostic tests to perform on a patient to determine the effectiveness of a treatment, given all the information about the patient, but also given the information about a much larger background population.

These days people often ask what comes after the Data Driven Discoveries of the Fourth Paradigm, what is the next step? Today, computers aid us in making detections, but discoveries are done by humans. It is clear that soon computers will also make the decision about which path to take to augment our existing detections with new information that leads to genuine new discoveries. Humans will still be in the loop, by specifying the broad context (the “dimensions of sparsity”) of the problem. We are looking at a formidable decade ahead of us, in which human aided machine learning decision making will produce major discoveries and surprises worthy of science fiction stories today.

5 REFERENCES

- Grabianowski, E. 2012. “3 Computer Simulations that Changed The World (And 2 That Are on the Verge).” *io9: We come From the Future*, July 6, 2012
- Hay, T., S. Tansley and K. Tolle. 2009. “The Fourth Paradigm, Data-Intensive Scientific Discovery”. *Microsoft Research*, Redmond, Washington.
- Heingartner, D. 2006. “Maybe We Should Leave That Up to the Computer.” *The New York Times Technology Section*, July 18, 2006.
- Howard, B. 2013. “What is adaptive cruise control, and how does it work?” *Extreme Tech, Electronics Section*, June 4, 2013
- Jaffe, E. 2014. “The First Look at how Google’s Self-Driving Car Handles City Streets”. *The Atlantic* April 28.
- Johnston, L., L. Schruben, A. Yang, and D. Zhang. 2008. “Establishing the Credibility of a Biotech Simulation Model”. *Proceedings of the 2008 Winter Simulation Conference*, edited by S. J. Mason, R. R. Hill, L. Monch, O. Rose, T. Jefferson, and J. W. Fowler, 822–826. Piscataway, New Jersey: Institute of Electrical and Electronic Engineers, Inc.
- King, R.D., K.E. Whelan, F.M. Jones, P.G.K. Reiser, C.H. Bryant, S.H. Muggleton, D.B. Kell, and S.G. Oliver 2004. "Functional genomic hypothesis generation and experimentation by a robot scientist". *Nature*, 427, 247-252
- Krause, C., editor. 1993. “Take a number: On letting computers make decisions.” *The Oak Ridge National Laboratory Review*, 26(2).
- Lucas, T. W., S. M. Sanchez, P. J. Sanchez, and W. D. Kelton. 2014. “Changing the paradigm: When your method of choice should be simulation”. *Working paper*, Naval Postgraduate School.
- Maule, A.J. 2009. “Can computers help overcome limitations in human decision making?” *Proceedings of NDM9, the 9th International Conference on Naturalistic Decision Making*, British Computer Society London, UK.

- Mitroff, I. I., and T. R. Featheringham. 1974. November. "On Systemic Problem Solving and the Error of the Third Kind". *Behavioral Science*, 19 (6): 383–393.
- Roy, S.C., G. Faulkner and S.-J. Finlay 2007. "Hard or Soft Searching? Electronic Database Versus Hand Searching in Media Research." *Forum: Qualitative Social Research*, 8(3): 20.
- Sanchez, S. M. 2014. "Simulation experiments: Better data, not just big data". *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk, S. Y. Diallo, O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller. Piscataway, New Jersey: Institute of Electrical and Electronic Engineers, Inc.

AUTHOR BIOGRAPHIES

BRUCE G. ELMEGREEN is a Research Staff Member in the Research Division of IBM. He has a B.S. in Physics and Astronomy from the University of Wisconsin and a PhD in Astrophysics from Princeton University. He has published over 500 papers on star formation and galactic structure. Other research includes designs and simulations of computer memory and switches using phase change material and piezoelectronic activation. His email is bge@us.ibm.com.

SUSAN M. SANCHEZ is a Professor in Operations Research at the Naval Postgraduate School, and Co-Director of the Simulation Experiments & Efficient Design (SEED) Center for Data Farming. She also holds a joint appointment in the Graduate School of Business & Public Policy. She has a B.S. in Industrial & Operations Engineering from the University of Michigan, and a Ph.D. in Operations Research from Cornell. She has long been active within the simulation community, including the Winter Simulation Conference Board of Directors, the American Statistical Association's Section on Statistics in Defense and National Security, and NATO's *MSG-088: Data Farming in Support of NATO*. Her web page is <http://faculty.nps.edu/smsanche/> and her email is ssanchez@nps.edu.

ALEXANDER SZALAY is the Alumni Centennial Professor of Astronomy at the Johns Hopkins University, and Professor in the Department of Computer Science. He is the Director of the Institute for Data Intensive Science. He is a cosmologist, working on the statistical measures of the spatial distribution of galaxies and galaxy formation. He is a Corresponding Member of the Hungarian Academy of Sciences, and a Fellow of the American Academy of Arts and Sciences. In 2004 he received an Alexander Von Humboldt Award in Physical Sciences, in 2007 the Microsoft Jim Gray Award. In 2008 he became Doctor Honoris Causa of the Eotvos University, Budapest. He enjoys playing with Big Data. He can be reached at szalay@jhu.edu.