

THE ROLE OF LEARNING ON INDUSTRIAL SIMULATION DESIGN AND ANALYSIS

Bahar Biller
Stephan R. Biller
Onur Dulgeroglu

Software Sciences and Analytics
General Electric Global Research
Niskayuna, NY 12309, USA

Canan G. Corlu

Metropolitan College
Boston University
Boston, MA 02215, USA

ABSTRACT

The capability of modeling real-world system operations has turned simulation into an indispensable problem-solving methodology for business system design and analysis. Today, simulation supports decisions ranging from sourcing to operations to finance, starting at the strategic level and proceeding towards tactical and operational levels of decision-making. In such a dynamic setting, the practice of simulation goes beyond being a static problem-solving exercise and requires integration with learning. This article discusses the role of learning in simulation design and analysis motivated by the needs of industrial problems and describes how selected tools of statistical learning can be utilized for this purpose.

1 INTRODUCTION

In today's uncertain world, simulation, with its capability of modeling real-world system operations, has become an indispensable problem-solving methodology for business system design and analysis. Whether decisions concern forecasting of sales volumes, supply chain risk management, or financial asset allocation, simulation of business systems has become the decision aid of choice for executives when making decisions under uncertainty. However, the large scale of business systems, coupled with the need to make data-driven decisions, have led to data challenges for simulation to support executive decision making and to ensure success of implementation over time. This temporal aspect of an industrial simulation project necessitates learning to occur, originating at the time of defining the problem and extending beyond calibration and validation of the developed model. In this article, we discuss the learning challenges that may be observed in practice and how the state-of-the-art methods of the discrete-event stochastic simulation literature can be used for guiding the industry practitioners in search of the ways to overcome those challenges in their development of simulation-based decision-support systems.

Specifically, we restrict our focus to a decision-support framework building on a transfer function which combines inputs with operational policies through a stochastic model to be calibrated, validated and used for system performance prediction. The analysis of the outputs is expected to generate insights to inform the decision maker about the pinch points of the solution recommended (e.g., system bottlenecks) and the uncertainty surrounding their prediction. Such a transfer function can be simply represented by

$$\mathbf{Y} = f(\mathbf{X}, \mathbf{Z}), \quad (1)$$

where $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_k)'$ is the k -dimensional vector of (possibly dependent) stochastic input processes (i.e., the temporal sources of uncertainty typically indexed to time), \mathbf{Z} is the collection of the operational policies that form the logic of the simulation whose execution provides an approximation to the transfer function f and \mathbf{Y} is the vector of the performance measures we aim to first accurately estimate and then robust optimize on the design parameters to improve system performance. If the transfer function f is

simple enough to be characterized by traditional mathematical analysis, exact results can be obtained from models built by queuing theory, differential equations or linear programming without any need for stochastic simulation (Kelton, Sadowski, and Zupick 2014). However, we often find ourselves faced with complex systems that cannot be validly represented by simple analytical models and that require the use of a stochastic simulation to approximate the transfer function. Without loss of generality, we choose a discrete-part production line as an example that is representative of a complex system and discuss the role of learning in simulation design and analysis as we experience in industrial research. In particular, we consider a process flow of ten steps with sharing of equipment and operator between the selected process steps for the purpose of numerical experimentation in this paper. A detailed description of the numerical experiment is presented in Section 3.

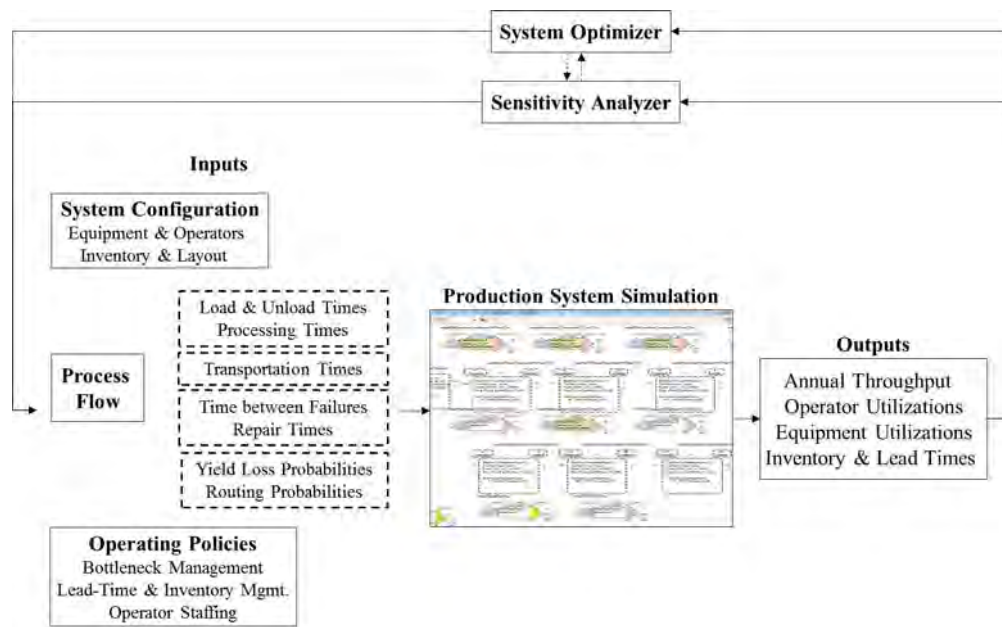


Figure 1: Illustration of a production operations transfer function.

Figure 1 illustrates the components of a simulation study to approximate the production system transfer function to accompany our discussion. The relation between this illustration and the algebraic characterization in (1) is as follows. The input random variables representing load, process, unload, transport and repair times, times between failures, yield loss and routing probabilities constitute the vector \mathbf{X} whose modeling builds on the state-of-the-art in simulation input modeling. Operating policies ranging from bottleneck management to lead-time and inventory control to operator staffing are summarized in the vector \mathbf{Z} while the output vector \mathbf{Y} consists of the output performance measures listed as annual throughput, operator and equipment utilization, inventory and lead-time in Figure 1. The production system simulation to be built depends on the type of decisions this simulation is intended to support. Thus, it is important to recall the well-known classification of the management decisions into the following three categories: Strategic planning, tactical planning and operational planning, which differ from each other by the duration of time to the implementation. If the production system is not built yet, then one of the key strategic decisions will be the identification of an optimal equipment portfolio to meet future demand while satisfying CAPEX budgetary constraints. The development of a simulation to inform us of a robust optimal equipment portfolio also requires the identification of operational strategies to support portfolio selection; e.g., tactical-level joint lead-time and inventory management policies and optimal staffing plans. This is the stage of the project with the highest level of uncertainty and often the input modeling for \mathbf{X} is based on experts' opinions

(and on the historical data collected for similar products and processes when available). It is, therefore, critical for the simulation practitioner to quantify the level of input uncertainty (i.e., the uncertainty due to the absence of complete information about the distribution of the input random variables) in the production system performance predictions and guide the learning of the simulation inputs in the direction to improve the accuracy of the predictions. In the stage of strategic planning, however, learning is not limited to the characterization of the stochastic inputs of the system. Sensitivity analysis and system optimization provide us valuable insights into how the simulation outputs respond to departures from the assumptions of system configuration and operating policies. This would enable the development of effective heuristics customized to solve the underlying optimization problem under uncertainty, which is the equipment portfolio selection in our example setting.

As the type of decision support expected from the execution of the project changes from strategic planning to tactical planning (for example, with the arrival of the purchased equipment), we find ourselves in the position of collecting operational data from part and equipment qualifications. It becomes critical to combine the newly collected data with the experts' opinions and analyze the simulation outputs in a way to better understand the sensitivities of the operations to many sources of uncertainty. Finally, the model is expected to be equipped with the capability to present support in real time and to enable continuous improvement via dynamic operational policies. Thus, model calibration and validation naturally arise as the key challenges to overcome. We provide examples of successful industrial applications in Section 4.

Motivated by the strategic, tactical and operational levels of decision making, we focus on the following three examples cases of learning: (1) Representing the uncertainty in the estimation of the mean performance measures due to the absence of full knowledge of the distribution of the input random variables; (2) improving the simulation input models with learning; and (3) operating policy learning via simulation optimization. In Section 2, we present a review of the simulation methodology work related to these three cases of learning. It is important to note that these three example cases are not meant to represent all learning related activities in simulation research. For example, we do not discuss the neural network that is well known to be applied to a wide range of learning problems with a high degree of accuracy. Instead, we restrict our focus to three examples of learning for clarity in the exposition of the paper. In Section 3, we consider ourselves in a situation of having developed a discrete-event stochastic simulation of a facility under design with the purpose of supporting equipment portfolio selection. Moving the clock forward, we also consider the case of having purchased equipment and making tactical delivery prediction constrained by the purchase. We conduct a numerical experiment for the discussion of these two example cases (1) and (2). In Section 4, we switch our focus to short-term delivery prediction and discuss how statistical learning can be further utilized for fast (and ideally automated) simulation building, analysis and optimization. We conclude with a brief of summary of discussion in Section 5.

2 RELATED WORK

A close look at the existing literature reveals that simulation input modeling research with focus on input uncertainty representation presents solutions to the learning cases (1) and (2). It is important to note the setting description as running a single simulation experiment and learning about the impact of input risk for given system configuration. The simulation methodology research building on the concept of exploration and exploitation, on the other hand, provides good solutions to the learning case (3). It is important to emphasize that calibration and validation of simulation models also fall into this learning category. It is because of the challenge created by the real-time aspect of dynamic problems in shorter horizons that industry practitioners continue to be in need of fast and robust solutions that will well integrate with the complexity of their industrial settings. Finally, if many simulation experiments have already been conducted across many scenarios (e.g., when it is not expensive to perform function evaluation using simulation), then it becomes a viable option to resort to the (scalable and supervised) machine learning algorithms with the purpose of learning from the large set of simulation experiment data. In particular, the decision tree

learning may provide a practical method for learning of discrete-valued functions, and thus, be utilized to support sensitivity analysis for discrete-event stochastic simulations.

The common statistical approach to each research stream described in the previous paragraph is the Bayesian method. Bayesian model development starts with the selection of a joint prior density function which quantifies the prior information about the inputs' distributions and parameters. The prior density function is then updated with the input data (when available) to obtain the posterior density function from which the inference is made. In the absence of a well-defined form for the joint posterior density function, it is common practice to resort to the Markov Chain Monte Carlo (MCMC) method for sampling from the posterior density function and obtaining estimates of the unknown distribution parameters. The idea behind any MCMC method is to simulate a random walk in the entire parameter space that converges to the joint posterior density function of the parameters (Gilks, Richardson, and Spiegelhalter 1996). As will be discussed in Section 3, this approach will be utilized when additional loops are incorporated into a simulation experiment with the purpose of representing additional layer of input uncertainty in the simulation output data. However, when the focus switches to learning from the simulation outputs in multiple stages across different design scenarios, there may arise situations when MCMC may not be the preferred method; e.g., when the prior beliefs are assumed to exist across different design scenarios with unknown correlations among them. The method of choice in this case would be Bayesian approximation which would approximate the posterior density function to be conjugate in the absence of a conjugate prior.

When the input distributions and their parameters are unknown and the historical input data available for their estimation are limited, there are three main sources of uncertainty to represent in the output analysis: stochastic uncertainty (i.e., the uncertainty that is due to the dependence of the simulation output on random input processes) (Helton 1997), model uncertainty (i.e., the uncertainty that is due to the selection of a single input model from a set of alternative models), and parameter uncertainty (i.e., the uncertainty that is due to the estimation of the input-model parameters from limited data) (Raftery, Madigan, and Volinsky 1996). Stochastic uncertainty is inherent in every simulation and controlled by the number of simulation replications. The model and parameter uncertainties are, on the other hand, often ignored as a result of driving the simulations with the probability distributions estimated from input data of finite length. This practice of simulation design and analysis leads to not only inconsistent estimates for the mean performance measures but also inconsistent coverage of the confidence intervals (Barton 2012). A close look at the stochastic simulation literature on accounting for parameter uncertainty shows that Cheng and Holland (1997) are the first to show the dependence of the simulation output on stochastic and parameter uncertainties using the delta and parametric bootstrap methods. These solution approaches are followed by the two-point method in Cheng and Holland (1998) as well as an alternative approach that particularly considers the simulation bias in Cheng and Holland (2004). Barton and Schruben (2001) characterize the parameter uncertainty in the simulation output by using the non-parametric bootstrap method. Barton, Nelson, and Xie (2014) introduce the metamodel-assisted bootstrapping and illustrate its performance in dealing with parameter uncertainty in the simulation of queuing systems. Following a parametric approach, Chick (2001) and Zouaoui and Wilson (2003) develop Bayesian models to represent the parameter uncertainty in the simulation output. Similarly, we show how to account for the parameter uncertainty with the parametric method in the numerical experiment presented in Section 3. As is the common practice, we do this by building on a random-effects model, in particular, the hierarchical normal model. However, the better approach that we also utilize in our industrial simulation research is to replace the hierarchical normal model with a response-surface model that will be robust to any deviation from the output response assumptions.

The simulation research with the capability to address the issues of calibration and validation (and adaptive improvement of operational policies by learning) originates from the study of sequential decision problems. It is important to note that learning is not directly on the simulation inputs in this case. Certain beliefs are assumed to exist among the alternative system designs while the data to be blended with the prior density function comes from the simulation itself. In particular, the knowledge gradient

policy – originally proposed for off-line ranking and selection problems – has been adapted to be used for online decision-making (and the study of multi-armed bandit problems); see Ryzhov, Powell, and Frazier (2012) for an example study of the theoretical foundation and Frazier, Powell, and Simao (2009) for an example calibration study via simulation for the transportation industry. Recently, Edwards, Fearnhead, and Glazebrook (2017) identify weaknesses of the knowledge gradient policy for online decision making and propose variants of the policy to overcome the weaknesses.

3 IMPROVEMENT OF SIMULATION INPUT MODELS WITH LEARNING

In the stage of strategic planning, it is often the case that there is no available historical information about an input random variable. The input model development depends on the elicitation of the information from the subject matter experts. Because a beta-PERT (or triangular) distribution is completely constructed with minimum, mean (or mode), and maximum values a random variable can assume, it is a particularly suitable candidate to model an input random variable in the absence of data. Nevertheless, it is natural to expect considerable uncertainty around the three parameters, and therefore, to ask the question of the level of confidence we really have around our simulation-based predictions. In this particular case, it is important to impose probability distributions on the input distribution parameters by assuming varying levels of uncertainty around them. More specifically, the input parameter values should be sampled from a region of uncertainty (characterized by the imposed probability distribution) prior to the execution of a simulation replication. Specifically, the uncertainty region of a beta-PERT distribution may be simply defined by an identical percentage deviation around each of the minimum, mode and maximum parameters. In the case of approximating the three-estimate distributions with a highly flexible system distributions known as the Johnson translation system (Johnson 1949), the uncertainty region could be built on an identical percentage deviation around each of the first four moments of the input random variable of interest because of the ability of the Johnson distribution system to capture any first four moments any (continuous) random variable can have.

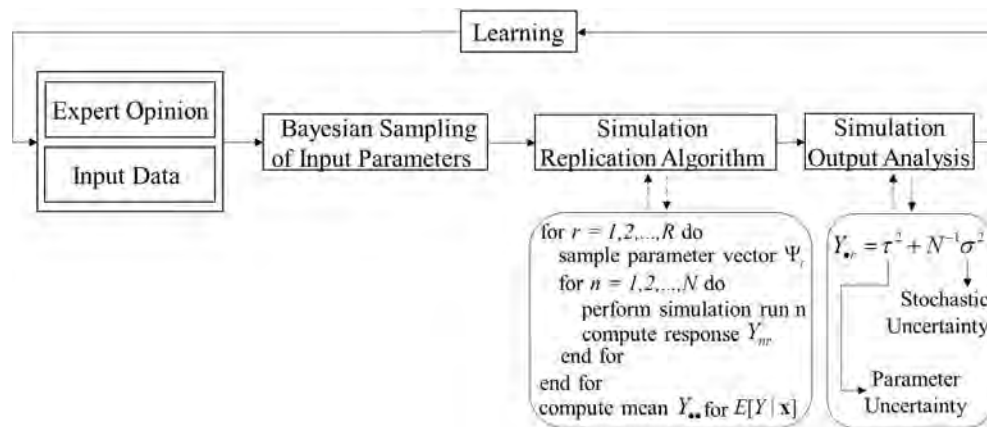


Figure 2: Representing input uncertainty in a stochastic simulation.

After the identification of a method of sampling input parameter values to reflect uncertainty around their assumptions, the simulation replication algorithm (Zouaoui and Wilson 2003) can be utilized for the propagation of the uncertainty through the simulation. The statistical summary of the simulation outputs can be presented together with the decomposition of the overall simulation output variance into stochastic uncertainty (due to the use of pseudo-random numbers in each replication of a simulation experiment) and assumption (input) uncertainty (due to the lack of knowledge of the parameters that are used to model the input random variables). We provide the illustration of our description of representing input parameter uncertainty in a stochastic simulation in Figure 2 together with the simulation replication algorithm. As

evident from this illustration, the simulation replication algorithm allows the representation of both stochastic uncertainty and input uncertainty by sampling input distributions and distribution parameters from their uncertainty models in the absence of data or from their Bayesian posterior density functions in the presence of collected data before each simulation replication. The reason behind the consideration of this particular simulation replication algorithm in our work is due to its ability to separately quantify the amounts of stochastic uncertainty and input uncertainty in the simulation output data.

Next, we present a numerical example demonstrating the use of the simulation replication algorithm with the purpose of gaining insight into the significance of the input parameter uncertainty in a 95% confidence interval relative to the stochastic uncertainty. Specifically, we consider a ten-step process flow where each step is dedicated to a different equipment except the second and fourth steps where only manual tasks are performed. There is a single operator that performs loading and unloading operations and is not needed during the machining process. Also, the pieces flow through the production line in a batch of size 25. We illustrate this short production line in Figure 3 with the details of each process step in Table 1. We note that there is a 10% variation around the operator loading and unloading times (given in minutes) while the machining portion of the process is automated and exhibits negligible variation. We use AnyLogic Professional Hybrid Simulator with Java source code to mimic the dynamics of this production line.

At this stage of the project, the values of the unknown parameters of the input distributions are assumed to be experts' opinions and/or be coming from the equipment vendors. The simulation for this production line is assumed to have been developed to support the solution of the equipment portfolio selection problem; i.e., the determination of the number of equipment of each type to purchase to reach the business production target with minimum shortfall and CAPEX, representative of a use case we have recently worked on at General Electric (Biller et al. 2017). The goal of this numerical study is not to show how we would solve the equipment portfolio selection; instead, we aim to understand the level of risk surrounding the recommendations made by a discrete-event stochastic simulation constructed with limited information. The outcome of this analysis will inform the decision-maker about the strategic delivery capability of the facility under design and lead to the learning of the sources of randomness in the system in terms of the contribution(s) they make to the overall system variability.

Because the objective of the simulation is to make a strategic delivery prediction, we set the duration of the simulation to 100 weeks, set the length of the warm-up period to 50 weeks, and perform steady-state output analysis in a way to recognize the additional layer of input uncertainty represented in the simulation output data. We first follow the method of Ankenman and Nelson (2012) to quantify the ratio between τ —

Table 1: Process-flow description for the numerical example using index of the process step, equipment name, capacity (i.e., number of parallel units), batch size, average batch load time (minutes), average batch process time (minutes), average batch unload time (minutes), MTBF mean time between failures (days), and MTTR mean time to repair (hours).

Step	Equipment	Capacity	Batch Size	Load Time	Process Time	Unload Time	Days MTBF	Hours MTTR
1	Machine 1	1	1	10	146	10	28	10
2	Manual	1	25	0	20	0	12.5	8
3	Machine 2	4	25	10	67	10	28	12
4	Manual	1	25	0	20	0	12.5	8
5	Machine 3	4	50	10	150	10	12.5	12
6	Machine 4	1	1	10	25	10	21	8
7	Machine 5	1	1	30	164	30	12.5	24
8	Machine 6	1	1	10	25	10	21	10
9	Machine 7	3	1	10	146	10	28	10
10	Machine 8	4	1	10	250	10	28	10

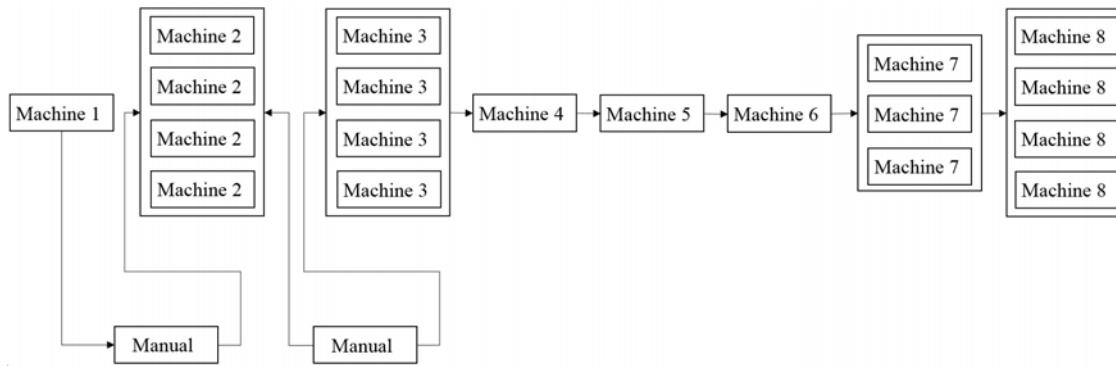


Figure 3: Illustration of the production line assumed by the numerical experiment.

the standard deviation of the simulation output due to input uncertainty — and σ/\sqrt{N} — standard error of the simulation output due to stochastic uncertainty — each of which is illustrated in Figure 2. We choose $R = 10$ and $N = 20$ to perform a total of 200 simulation replications. We note that Ankenman and Nelson (2012) assume the availability of limited historical data and thus, $R = 10$ corresponds to the number of bootstrap samples. Although we do not assume the availability of any historical data at the strategic level of decision making, we sample input distribution parameters from a region of uncertainty surrounding the process-step assumptions; therefore, the nature of the algorithm and the analysis of the simulation outputs are the same. We compute the ratio of input uncertainty to standard error of the simulation to have a point estimate of 5.90 and a 95% confidence interval of (3.85,11.35). Thus, with 95% confidence, the input uncertainty is 3 to 11 times as large as the simulation standard error for the mean throughput.

Next, we calculate the 95% confidence interval for the mean throughput. In the case of ignoring input uncertainty and accounting for only the stochastic uncertainty in the simulation output analysis, we identify the 95% confidence interval for the mean throughput as (3,669, 3,761) with a mean of 3,715 units. However, when we account for both stochastic uncertainty and input uncertainty, we identify the 95% confidence interval as (2,984, 3,723) with mean 3,353 units. Thus, accounting for the simulation input risk shows a reduction of 9.74% in the mean throughput. We illustrate these findings in the left-hand side of Figure 4 and show that we would overestimate the production capacity of the facility under design if we had not recognized the risk around the assumptions we made for the sources of uncertainty driving our discrete-event stochastic simulation.

The next step is to identify the sources of uncertainty that are making the largest amounts of contribution to the variation around the mean throughput prediction. While we do not discuss this step for the numerical experiment on hand, we refer the reader to Ankenman and Nelson (2012) for an example application of sequential bifurcation to search for important factors in a low-dimensional input environment and to Song and Nelson (2015) for the increasing number of inputs needed for larger scale stochastic simulations.

For the numerical experiment, we now assume that purchase decision is made for the recommended equipment portfolio, the full equipment has arrived the facility and as part of the equipment qualification process, data are being collected. Therefore, it is time to update our process-flow assumptions and improve the input models with the data available now. In this stage, often there would be changes to the process-flow itself in the form of a removal and/or addition to the ten steps illustrated in Figure 3. In the hypothetical experiment here, we assume that the structure of the process flow remains the same and that we only have access to input data which can be used for improving the accuracy of the input models. We now design the experiments under the assumption that the initial load-time, process-time and unload-time assumptions in Table 1 are 10% higher than the true (but unknown) load-time, process-time and unload-time parameters. An important simplifying assumption we make here is that beta distribution is the true distribution. Hence, the input uncertainty is, in fact, the parameter uncertainty, which is due to the estimation of the unknown input distribution parameters from limited historical input data.

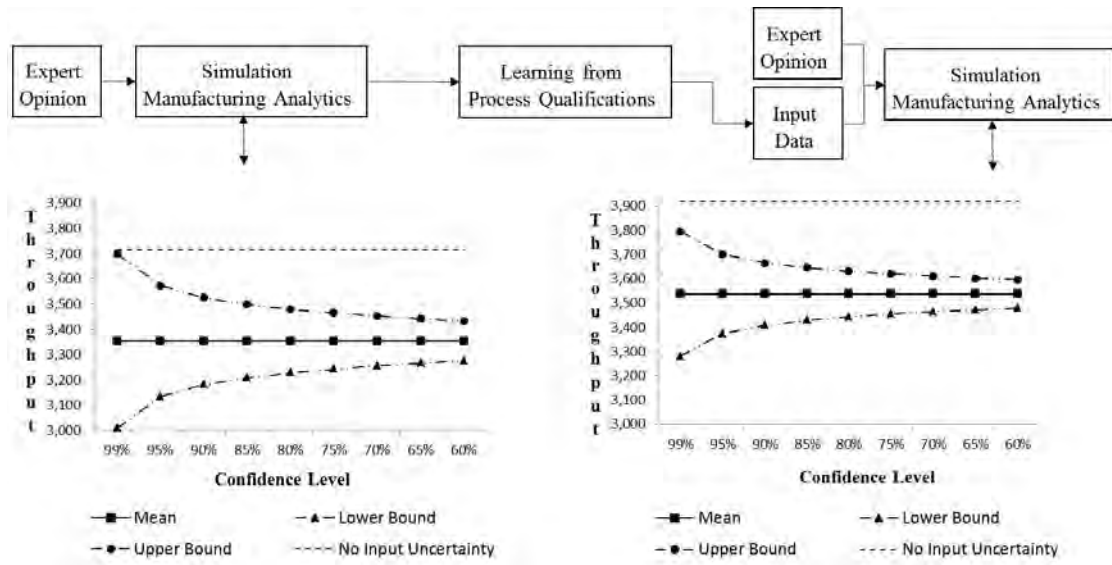


Figure 4: 95% confidence intervals for mean throughput before learning (left) and after learning (right).

We illustrate our findings in the right-hand side of Figure 4. Prior to the data collection, the 95% confidence interval constructed for mean throughput is found to have a total standard deviation of 191 units, decomposed into 32 units due to stochastic uncertainty and 188 units due to input uncertainty. After the collection of 100 data points for each process step, we now identify the total standard deviation of the 95% confidence interval for the mean throughput to decrease from 191 units to 143 units (i.e., a reduction of 25.13%). We also determine the decomposition of 143 units into 24 units due to stochastic uncertainty and 141 units due to input uncertainty. Despite the reduction of the variance in the prediction of the mean throughput, the effect of input uncertainty still dominates the effect of stochastic uncertainty, and accounting for input uncertainty continues to be important in the tactical planning phase of the project. A natural question to ask is “what is the impact of 25% variability reduction on the estimation of the mean throughput?” The answer is illustrated in Figure 5. Specifically, we identify the 95% confidence interval as




Figure 5: Illustration of the impact of learning on throughput and variability.

(3,335, 3,739) with 3,537 for the mean throughput. Thus, we now predict that we can obtain 5.49% more throughput than we initially predicted. Under the assumption of CONWIP (CONstant Work In Process) operating policy for joint wafer release and inventory management, we also find that the production capacity assumed at the strategic level can be achieved with 15% less inventory in the system. However, it could very well have been the case that we obtain less throughput from the system than we initially estimated. In that case, the project would switch attention to new scenarios of capacity planning (e.g., out-sourcing) as a way of hedging against the system uncertainty. In either case, it is critical to do a good job in maintaining the discrete-event simulation of the system and adding new modeling features to better capture operations of the shorter time horizons. The goal is to have the simulation to be ready for being driven by the real-time data when the project moves from tactical to operational level of decision making.

We conclude this section with the note that as we move from strategic planning to tactical planning and start collecting data for the simulation inputs, alternatives to distributions constructed with three-point estimates (i.e., minimum, mode (or mean), and maximum) may be preferable to be able to capture a wider variety of distributional characteristics, despite their ability to represent symmetric as well as positively and negatively skewed distributional shapes. In fact, there exists a well-established literature on simulation input modeling with flexible system of distributions. The representation of input parameter uncertainty can be readily performed by following the Bayesian approach. We again first select a joint prior density function for the unknown parameters of the flexible distribution family, next update the prior density function with the likelihood function of the collected input data to obtain the posterior density function, and finally sample input parameter values from the posterior density function to drive the simulation with the simulation replication algorithm. For an example application, we refer the reader to Biller et al. (2017) where we perform these steps for the Johnson translation system in a semiconductor manufacturing simulation.

4 OPERATING POLICY LEARNING VIA SIMULATION OPTIMIZATION

One of the challenges of using stochastic simulation for strategic planning is the lack of the opportunity to validate the model until the system implementation. Nevertheless, the step of verification presents the opportunity to understand why the simulation outputs may react unexpectedly to certain deviations from the modeling assumptions. As we move from strategic planning to tactical planning, we start working on the validation of the model with the data collected during the qualification step. Not only we can update the stochastic input assumptions but also modify process-flow assumptions. It is critical to assess the impact of the process-flow learnings that are the results of the logic changes in the simulation. In a recent project of designing a new production line at General Electric, we quantify the benefit of a similar learning process as an improvement of 60% in key performance index estimation accuracy through reinforcement learning.

Assuming that we have already reached the stage of operational planning, the designed system is now in operational mode and the model to support its decisions must be calibrated and validated. However, the calibration-validation task should not be a one-time activity. Ideally, a red flag is presented to show what we simulate is not what is actually occurring in reality; hence, the anomaly corresponding to this discrepancy is to be detected building on the state-of-the-art anomaly detection algorithms. Once an anomaly is detected, it is also important to determine the root cause of the event and quickly respond to mitigate the impact on the system performance. Observing the system and accumulating information of the types of events occurring will eventually enable us to establish the potential modes of operations (i.e., regimes) for the facility. Such clustering of operations of the facility would be instrumental in the development of robust operational policies and smooth regime switches within the facility. Hence, when simulation is used for operational level of decision making, artificial intelligence presents itself as an opportunity for discrete-event simulation to integrate itself into the facility as a near-real decision-support tool.

It is important to note that a big obstacle to the utilization of simulation for short-term decision support is often the lack of a simulation logic that can capture the behavior of the system's intelligent objects (to the extent the simulation will provide decision support). The failure to do so often explains why a simulation model that is used for strategic planning falls short of meeting the needs of the decision

support necessary in near real time. This brings out the challenge of supporting real-time simulation design and analysis by within simulation optimization, which is often not addressed in simulation methodology research (to the best of our knowledge). An example application of this situation occurs when the objective is to identify automated guided vehicle routing patterns along available paths to minimize the mean time spent by different kinds of parts in the system. Dulgeroglu (1994) proposes a solution to this problem by combining the learning automata theory with simulation analysis and demonstrates the importance of integrating improvement heuristics and artificial intelligence with simulation to provide prescriptive power.

When our project is in the stage of strategic or tactical planning, we often perform simulation design and analysis by following the principles of steady-state analysis. In the stage of operational planning, however, it becomes critical to account for the current system configuration, for example, the number of units of inventory kept at different locations in the system. This concept is known as hot-starting the simulation and it plays a critical role in the model validation step. It further pushes discrete-event stochastic simulation, originally developed for strategic decision support, towards being a data-driven simulation. As an example of a data-driven simulation, Akbay, Dulgeroglu, and Toledano (2011) describe the use of General Electric's generic modeling toolkit for hospital operations. The authors discuss how snapshots of the system are taken several times each day and they are fed into the model to partially hot-start the supporting simulation with the current state of the hospital, presenting additional opportunities to incorporate learning into simulation design and analysis. Today, we see similar practices of designing and analyzing simulations across a wide spectrum of industries. Tannock et al. (2007) describe concept and operation of a supply-chain model builder for aerospace manufacturing. Meng et al. (2013) describe a data-driven modeling and simulation framework developed for material handling system of coal mines. General Electric has been developing generic auto-simulation capability to provide data-driven simulation capability targeting manufacturing applications (Annunziata and Biller 2012). Cheng and Law (2017) discuss data-driven simulations in the present age of big data, providing a summary of work that has been presented primarily at the Winter Simulation Conference over the years.

Despite the increasing computing capabilities, we still encounter situations in which running a large-scale system simulation may not be possible in a short amount of time. In those situations, it becomes important to have the capability to characterize the simulation output response by running simulations with fewer design points. In the case of deterministic simulation (i.e., the stochastic uncertainty is negligible), it often suffices to rely on the statistical learning (also, known as machine learning) tools (James et al. 2015) for response surface modeling. Especially when local approximation is all that is needed, the generalized additive model that allows customization to individual factors often presents good solutions to the simulation response-surface modeling problem. However, these response-surface models fall short of providing good predictions at points that are not included in the design space of the training data. While there may be alternative remedies to this problem, a close look at the existing literature reveals kriging as the method of choice for successful global meta-modeling. We refer the interested reader to Jones, Schonlau, and Welch (1998) for an excellent introduction to the method of kriging where the concern of the authors is primarily geared towards simulation optimization for the automotive industry. The latest review of the advances made in this research stream is, on the other hand, provided by Kleijnen (2017), including the customization of the method of kriging to stochastic simulation experiments which require the modeling of the stochastic uncertainty in simulation outputs. In recent years, we also see the research emerging in this area to be utilized for input uncertainty quantification as alternative to the use of the random-effects model for representing the simulation output process. As the complexity of the simulation model increases, the amount of output data collected from simulation increases and it becomes even more critical to support simulation design and analysis with the tools of statistical learning under limited time and budgets.

5 CONCLUSION

This article discusses learnings that naturally arise in simulation projects to provide decision support in different stages of (strategic, tactical, and operational) planning. Specifically, we provide a brief review

of learning from experts at the strategic stage and from qualification data at the tactical stage of the project, accompanied by a numerical experiment. We also discuss how simulation optimization presents opportunities for learning when the focus is especially on calibration and validation of the simulation model. We conclude with the statistical learning tools that may be effective for near real-time decision support.

REFERENCES

- Akbay, K., O. Dulgeroglu and D. Toledano. 2011. "Designing, Improving and Efficiently Operating Hospitals with Help from General Electric's Generic Modeling Toolkit". INFORMS 2011 Healthcare Conference, Montreal, Canada, June 20 - 22, 2011.
- Ankenman, B., B. L. Nelson, and J. Staum. 2010. "Stochastic Kriging for Simulation Meta-Modeling". *Operations Research* 58:371–382.
- Annunziata, M. and S. Biller. 2012. "The Future of Work". General Electric White Paper.
- Barton, R. R. 2012. "Tutorial: Input Uncertainty in Output Analysis". In *Proceedings of the 2012 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 67–78. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Barton, R. R. and L. W. Schruben. 1993. "Uniform and Bootstrap Resampling of Input Distributions". In *Proceedings of the 1993 Winter Simulation Conference*, edited by G. W. Evans, M. Mollaghasemi, W. E. Biles, and E. C. Russell, 503–508. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Barton, R. R. and L. W. Schruben. 2001. "Resampling Methods for Input Modeling". In *Proceedings of the 2001 Winter Simulation Conference*, edited by B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, 372–378. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Barton, R. R., B. L. Nelson, and W. Xie. 2014. "Quantifying Input Uncertainty via Simulation Confidence Intervals". *Inform Journal on Computing* 26:74–87.
- Biller, B., O. Dulgeroglu, M. Hartig, R. J. Olson, P. Sandvik, G. Trant, and C. G. Corlu. 2017. "Semiconductor Manufacturing Simulation Design and Analysis with Limited Data". Forthcoming in *Proceedings of the 28th Annual SEMI Advanced Semiconductor Manufacturing Conference*, Saratoga Springs, NY, USA, May 15 - 18, 2017.
- Cheng, R. C. H. and W. Holland. 1997. "Sensitivity of Computer Simulation Experiments to Errors in Input Data". *Journal of Statistical Computation and Simulation* 57:327–335.
- Cheng, R. C. H. and W. Holland. 1998. "Two-point Methods for Assessing Variability in Simulation Output". *Journal of Statistical Computation and Simulation* 60:183–205.
- Cheng, R. C. H. and W. Holland. 2004. "Calculation of Confidence Intervals for Simulation Output". *ACM Transactions on Modeling and Computer Simulation* 14:344–362.
- Cheng, R. C. H. and A. Law. 2017. "History of Input Modeling". Forthcoming in *Proceedings of the 2017 Winter Simulation Conference*, edited by W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Chick, S. E. 2001. "Input Distribution Selection for Simulation Experiments: Accounting for Input Uncertainty". *Operations Research* 49:744–758.
- Dulgeroglu, O. 1994. "Intelligent Simulation Modeling of a Flexible Manufacturing System with Automated Guided Vehicles". Manuscript MU-SEAS-CSA-1994-001. Department of Computer Science and Systems Analysis, Miami, University Oxford, Ohio.
- Edwards, J., P. Fearnhead, and K. Glazebrook. 2017. "On the Identification and Mitigation of Weaknesses in the Knowledge Gradient Policy for Multi-Armed Bandits". *Probability in the Engineering and Informational Sciences* 31:239–263.
- Frazier, P., W. B. Powell, and H. P. Simao. 2009. "Simulation Model Calibration with Correlated Knowledge-Gradients". In *Proceedings of the 2009 Winter Simulation Conference*, edited by M. D. Rossetti, R.

- R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 339–351. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Gilks, W. R., S. Richardson, D. J. Spiegelhalter. 1996. *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall, Inc.
- Helton, J. C. 1997. “Uncertainty and Sensitivity Analysis in the Presence of Stochastic and Subjective Uncertainty”. *Journal of Statistical Computation and Simulation* 57:3–76.
- James, G., D. Witten, T. Hastie, and R. Tibshirani. 2015. *An Introduction to Statistical Learning with Applications in R*. 6th ed. New York: Springer, Inc.
- Johnson, N. L. 1949. “Systems of Frequency Curves Generated by Methods of Translation”. *Biometrika* 36:149–176.
- Jones, D. R., M. Schonlau, and W. J. Welch. 1998. “Efficient Global Optimization of Expensive Black-Box Functions”. *Journal of Global Optimization* 13:455–492.
- Kelton, W. D., R. Sadowski, and N. Zupick. 2014. *Simulation with Arena*. 6th ed. New York: McGraw-Hill, Inc.
- Meng, C., S. S. Nageshwaranier, A. Maghsoudi, Y. Son, and S. Dessureault. 2013. “Data-Driven Modeling and Simulation Framework for Material Handling Systems in Coal Mines”. *Computers and Industrial Engineering* 64:766–779.
- Raftery, A. E., D. Madigan, and C. T. Volinsky. 1996. “Accounting for Model Uncertainty in Survival Analysis Improves Predictive Performance (with discussion)”. In *Bayesian Statistics*, edited by J. M. Bernardo, J. O. Berger, A. P. David, and A. F. M. Smith, 323–349. Oxford: Oxford Press.
- Ryzhov, I. O., W. B. Powell, and P. I. Frazier. 2012. “The Knowledge Gradient Algorithm for a General Class of Online Learning Problems”. *Operations Research* 60:180–195.
- Song, E., B. L. Nelson. 2015. “Quickly Assessing Contributions to Input Uncertainty”. *IIE Transactions* 47:1–17.
- Tannock, J., B. Cao, R. Farr, and M. Byrne. 2007. “Data-Driven Simulation of the Supply-Chain – Insights from the Aerospace Sector”. *International Journal of Production Economics* 110:70–84.
- Zouaoui, F. and J. R. Wilson. 2003. “Accounting for Parameter Uncertainty in Simulation Input Modeling”. *IIE Transactions* 35:781–792.

AUTHOR BIOGRAPHIES

BAHAR BILLER is a Senior Research Scientist at General Electric Global Research Center. Her research lies in the area of quantitative risk management and advances the design and analysis of large-scale business system simulations to aid decision-making under uncertainty. Her e-mail address is bahar.biller@ge.com.

STEPHAN R. BILLER is a Chief Scientist of Enterprise Analytics at General Electric Global Research Center. His research lies in the area of enterprise risk management and digital manufacturing. His e-mail address is biller@ge.com.

CANAN G. CORLU is an Assistant Professor of Administrative Sciences Department of Metropolitan College at Boston University. Her research interests are in the areas of design and analysis of stochastic simulations with applications to operations management. Her email address is canan@bu.edu.

ONUR DULGEROGLU is a Senior Research Scientist at General Electric Global Research Center. His research lies in the area of business system risk modeling and discrete-event simulation of systems ranging from transportation to energy to financial systems. His e-mail address is dulgerog@research.ge.com.