

## MODEL-BASED MEASUREMENT OF SITUATION AWARENESS

W. Scott Neal Reilly  
Sean L. Guarino

Charles River Analytics  
625 Mount Auburn Street  
Cambridge, MA 02138, USA

Bret Kellihan

DCS Corporation  
1330 Braddock Place  
Alexandria, VA 22314, USA

### ABSTRACT

Decision making in complex environments in the face of uncertain and missing information is a daunting task. We describe a modeling and simulation based approach to providing planners, analysts, and decision makers with a better understanding of the effect of imperfect information on the reliability of decisions made in such situations. We use techniques adopted from Sensitivity Analysis to evaluate the sensitivity of particular decision-making procedures to the uncertainty associated with the information that is being used to make the decision. We use this analysis to support the development of more robust decision-making procedures and effective and efficient information-gathering plans. We demonstrate how these tools can be used in both on-line decision analysis and off-line decision evaluation and development, and we describe how these tools can be used to support complex simulation systems such as the U.S. Army's Modeling Architecture for Technology and Research EXperimentation (MATREX).

### 1 INTRODUCTION

Knowing what is going on is obviously an important part of making good decisions. This basic insight underlies the significant amount of work in Situation Awareness in military and aerospace domains, but the same idea is applicable to any dynamic domain. Unfortunately, reliable information is often an elusive target. For instance, it might be based on sources that are not fully reliable, such as human reports or sensors that are only accurate to a certain resolution, or it might be old data that has possibly changed since it was gathered. And, in some cases, the information might be missing altogether.

We would like to provide decision makers with the ability to answer a variety of important questions, including: In the current situation, given the available information and its reliability, can I make a correct decision? If not, how wrong could the decision be? If I were to gather more

information, where should I focus my efforts? And if I gather more information, how much would that information improve the reliability of my decision-making ability?

The answers to these questions will differ from decision to decision and from situation to situation for similar decisions. For instance, knowing the locations of nearby aircraft is critical for many aviation decisions relating to navigation, but plays no role in the decision about when to extend the aircraft's landing gear. Also, knowing the distance to other planes to a great level of accuracy is more important when the other planes are near than when they are far.

The field of Sensitivity Analysis provides a number of tools and techniques that can be used to evaluate and visualize the impact of the uncertainty of inputs on the outputs of complex systems, including decision making systems. We have adopted and adapted a number of these ideas and applied them to our current effort.

In this paper, we describe one of the tools that we are currently developing for the Situation Understanding and Reasoning Evaluation (SURE) toolkit. SURE is a suite of tools used by simulation analysts to evaluate the situation understanding and reasoning processes of complex decision-making processes. The tool we describe uses computer-based models of decision-making processes and uncertainty associated with the information used to make those decisions to evaluate the stability of the decision-making process in the face of the uncertain information.

This is a general-purpose tool that can be used in a wide range of domains. In this paper, we describe our current efforts to use the SURE tool to support the U.S. Army's MATREX (Modeling Architecture for Technology and Research EXperimentation) simulation environment (Kellihan and Washington 2006). The goal of this effort is to support the development of more robust decision-making processes, to support the development of effective and efficient information-gathering procedures, and to guide the design of more complex simulation-based tests of complex decision-making environments.

In the following section, we describe a formal decision-making model that will help to ground our discussion. In Section 3, we describe methods for measuring and analyzing the impact of uncertainty on decision making. In Section 4, we describe a number of techniques for visualizing the results of this analysis. In Section 5, we describe future work, including plans for applying these techniques in a simulation system being developed by the U.S. Army for analyzing future technologies and processes. In Section 6, we relate this effort to work in a number of related fields, including Situation Awareness, Decision Theory and Information Value Theory, and Sensitivity Analysis.

## 2 DECISION MAKING MODEL

In a simple but general sense, decisions can be viewed as functions that take inputs and produce outputs. A medical diagnosis takes symptoms and medical history and produces a diagnosis and treatment plan. An investment decision looks at business and market factors when determining what, how much, and when to buy. A military commander uses information about the current state of a conflict to decide on tactics. While this model may not easily fit all decision-making processes, we have found it quite general in its expressiveness. Figure 1 shows this model graphically.

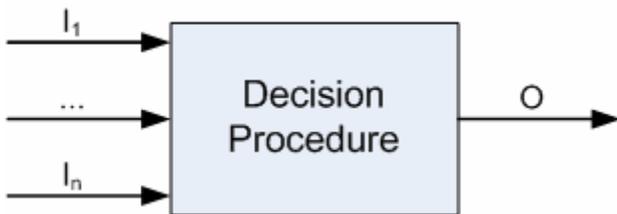


Figure 1: Decision procedure

To be compatible with the analysis we want to perform on decision procedures, we place two restrictions on this model: decision procedures must be deterministic and they must produce an output.

The Decision Procedure (DP) needs to be deterministic for our purposes as randomness makes it difficult to determine when variability in a decision is based on changes in the inputs instead of as the result of a stochastic process. A stochastic DP can still be modeled by pulling the randomness out as one or more separate input variables. For instance, if the decision process is partly determined by the result of a random number generator, this random number can be pulled out as a new input to the process, which allows the effects of the randomness to be evaluated independently of the other inputs.

The other constraint is that the DP produce an output for any given set of inputs, even if some of the inputs are missing. The output can be to do nothing, but the decision procedure must respond.

The inputs passed to the DP are not necessarily accurate representations of ground truth; they simply capture what information is currently available. For instance, the information might come from a source with limited accuracy (e.g., a thermometer that is known to be accurate to within 1 degree), or it might not be current (e.g., an enemy tank’s position might be reported, but a tank can move; the older the information, the further the tank can be from where it was last reported). So, given the input values and a level of uncertainty associated with each, we can define a space of inputs that indicates where the true values of the inputs might be. Figure 2 reflects the input and output spaces that are plausible. That is, the real input values are uncertain, but we have bounds on those values that indicate where the true values can be.

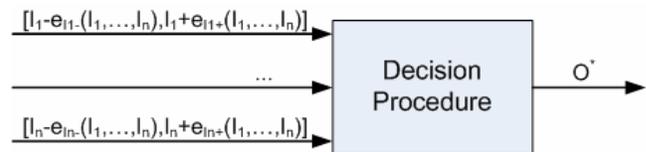


Figure 2: Input/output spaces for decision values

The bounds of uncertainty associated with each input are represented as a pair of input-specific functions ( $e_{I^-}$  and  $e_{I^+}$ ). The parameters to these functions are the current input values ( $I_1, \dots, I_n$ ). So, for instance, we can model that the error associated with a human report of an enemy target’s location is less accurate the further away the target is and is adversely affected by other inputs, such as the weather.

The output values associated with the range of possible input values form a set called  $O^*$ . If, for instance, we found that no matter where in the input space we are, the decision procedure (DP) mapped the full range of possible inputs to the same value (i.e.,  $O^*$  is a singleton), then the information we have is sufficient to make the decision and gathering additional information is not helpful. We often refer to this as being a *stable* point in the decision space.

We also assume that there is some function  $E_O(O_1, O_2)$  that computes the distance between any two output values, which will allow us to evaluate how similar two decisions are. The construction of such a function is simple in some cases and more challenging in others (such as when the outputs are members of an enumerated set). The error function is intended to provide the analyst with a general sense of how inaccurate a decision may be, as discussed below, and therefore does not have to be especially precise. So, in the case of actions from a set, the error function can simply return 1 if the outputs are different and 0 if they are the same. More subtlety can be added if some members are closer than others and this is represented by higher or lower values returned by this function. The relative values returned are of most interest to the analyst.

### 3 MEASURING THE EFFECTS OF UNCERTAINTY

Using this decision-making model, we can use a number of techniques from Sensitivity Analysis to aid in the evaluation of the stability of a particular decision procedure. Sensitivity Analysis provides general-purpose techniques and tools for measuring the influence of inputs to a function or model on the outputs. Sensitivity Analysis techniques come in two basic types: global analyses and local analyses (Saltelli et al. 2004; Saltelli, Chan, and Scott 2000).

Global sensitivity analyses look at the relative impact of each input on the output. The analysis is global in that it looks at the full range of possible values for each input to determine the full range of output effects.

Local sensitivity analysis is comparable to partial differentiation where the instantaneous slope of the function is measured against each input variable. This form of analysis is local in that it is done at a particular point in the input space. While comparable to partial differentiation, many instances of local sensitivity analysis are designed to work even with procedures that are not differentiable. This is important for techniques being applied to decision-making models where, for instance, non-differentiable threshold functions are common.

Neither of these approaches, however, is completely suitable for our task. Global approaches do not provide local information. This means that, for instance, they can inform the analyst that knowing the distance to other aircraft is important, but not that the accuracy of this information is more important when the aircraft are closer or when weather is bad. Pure local approaches are also not appropriate as they will give the slope at a point, but do not take into account the distance from that point that reality might actually lie. For instance, knowing that small errors in a distance measurement can lead to large decision errors is only important if there is some chance of error in that measurement. Similarly, knowing that the effect of a particular input only matters if there is a large uncertainty associated with that input does not help us if we don't know whether or not there is such a level of uncertainty. So, by ignoring the amount of uncertainty associated with each input, local analyses do not provide enough information for a thorough analysis.

The SURE tool uses a hybrid technique that provides local information, but does so with an awareness of the uncertainty level of each input. Our approach can be viewed as a global analysis within a very small, localized input space. Our approach and two applications of it are described in the next two sections.

#### 3.1 Situation-Specific Decision Analysis

It is often the case that we want to evaluate the stability of a decision procedure in a particular situation. That is, how

much does the current level of uncertainty affect the ability to make the correct decision in the current situation? This can be useful either as part of an off-line analysis process or as part of an on-line process that supports making better decisions about when to act, when not to act, and when to gather additional information.

If we have a specific situation to evaluate, which includes a DP and particular vector of  $i$  input values, we use a variant of Morris's local Sensitivity Analysis approach. Morris incrementally perturbs each input value by some small  $\delta$  to measure the influence of that input at a point in the input space (Morris 1991). We also use a perturbation-based analysis, but we evaluate the DP at the given point and at the  $2i$  points that result from independently modifying each input to the extreme values associated with that input's uncertainty function, instead of by a fixed value.

This approach assumes that the uncertainty is small enough that the effects on the output are monotonic. If this is not the case, we can easily change this approach to sample the possible inputs in between the sample point and the extremes as needed to capture these effects. We could also extend this approach to evaluate all of the vertices of the hypercube that represents the uncertainty surrounding the given input vector, though this is exponential in the number of inputs and does not scale to many practical problems, especially where we want to evaluate large numbers of situations, such as the full decision procedure analysis that we describe in the next section.

This analysis allows us to answer the questions we introduced in Section 1:

- *In the current situation, given the available information and its reliability, can I make a correct decision?* In this case, we can look at the decision that would be made given the current inputs and how different this decision could be if we were to have access to ground truth. We call this the *instability* of the decision procedure at this point. If the uncertainty associated with the inputs does not result in different decisions, then we can make the decision with confidence. We call this a *stable* decision procedure.
- *If not, how wrong could the decision be?* Again, since we have evaluated the decision based on the given information as well as in a representative set of possible situations that might represent ground truth, we can determine how far the decision is from what it might be if we had access to ground truth by using the  $E_0$  function described in Section 2. If the decision is close, we might decide to use the decision instead of gathering additional information.
- *If I were to gather more information, where should I focus my efforts?* In this case, if we have means of reducing the uncertainty associated with certain inputs, we want to know which inputs we should focus on to maximally improve the likelihood of making a correct decision. Since we have evaluated the potential deci-

sion error by sequentially perturbing each input value based on its uncertainty function, we know how much the uncertainty of each input effects the output, which, in turn, provides insights into which information is most productive to gather. We call the instability that results from perturbing an individual input the *local instability* or the *instability with respect to* that input.

- *And if I gather more information, how much would that information improve the reliability of my ability to make decisions?* As just described, since we know the potential error associated with the various input uncertainties, we know how much of an effect gathering additional information about any particular input will have on the reliability of the overall decision reliability. This can be compared to the costs of gathering the additional information to determine whether it is cost-effective to gather more information or to act under uncertainty.

### 3.2 Situation-Independent Decision Analysis

In addition to the analysis of specific situations just described, SURE also provides a means for performing an off-line, global analysis of where a decision procedure might break down in the face of uncertainty. For instance, when developing a decision procedure, it is useful to understand where the decision procedure could break down due to information uncertainty. This can enable the developer to modify the decision procedure to make it more robust. Such a global analysis can also speed up the local-situation analysis described in the previous section by, in effect, pre-computing the role of information uncertainty in a wide range of situations.

In theory, what we want to do for such an analysis is to walk through the possible values of the input vectors, evaluate the DP to produce an output value, and compare that output value to the other output values that are produced by input values that are within the space of anticipated uncertainty surrounding the initial input vector. We then report the maximum error between the output associated with each input vector and other uncertainty-based outputs to the analyst. This *instability* value represents the maximum error that is associated with making the decision with the current information and associated uncertainty.

In practice, however, this approach is too computationally expensive for most non-trivial problems. Instead, we use a Monte Carlo process to approximate the same basic approach. The current SURE tool breaks the  $i$  inputs into  $n$  subranges, where  $n$  is provided by the analyst. Larger values of  $n$  provide more accurate results, but require more computation time. We choose a set of  $n$  Latin Hypercube samples using these subranges, where a value from each subrange of each input variable is randomly assigned to one of the samples.

At each of the points in the sample, we perform the local decision analysis described in the previous section. That is, we evaluate the DP at the sample point and at the  $2i$  points that result from independently modifying each input to the extreme values associated with that input's uncertainty function.

This process is repeated  $m$  times, where  $m$  is specified by the analyst who can tradeoff accuracy for speed, with a different set of Latin Hypercube samples chosen each time. This provides a greater coverage of the whole input space and provides data for analyzing cross-input relationships.

The final result is a set of  $nm$  points distributed throughout the input space and the  $i$  maximum error values associated with the uncertainty around each input at each of those points. This data provides insights into the impact of uncertainty of each input on decision making in a wide range of situations.

One of the difficulties of this analysis is that it produces a great deal of data for the analyst. To support the analyst in sifting through this data for the relevant features, we provide a number of visualization tools that are described in the following section.

## 4 VISUALIZATION

The results of the analysis described in the previous section can provide significant insights into the role of uncertain information in the decision-making process. Unfortunately, the resulting dataset is also large and multi-dimensional, making it difficult for the human analyst to understand the results. To address this issue, it is useful to provide visualization tools that support the human analyst in this process. We have used two approaches to visualization: scatter-plots and parallel-coordinate plots.

Scatter plots provide two- or three-dimensional Cartesian graphs of relevant data. For instance, we have found it useful to create a series of plots where the x-axis is the value of each input and the y-axis is the local instability with respect to the given input that is associated with decisions made at that input value. That is, instead of plotting the maximum error value of perturbing all of the inputs, we plot the error associated with the perturbation of that particular input.

Typically, there will be a number ( $m$ ) of plots for each subrange of the input. For instance, imagine a decision procedure that simply returns the value of the function in equation (1). While simple, this function provides both a threshold ( $x_1 > 0.5$ ) and a continuous element ( $x_2$ ) that is only sometimes referenced, and therefore makes a good demonstration example. Thresholds in particular, are common in decision procedures and tend to be difficult for some sensitivity analysis procedures to handle. For this example, we posit a constant uncertainty value of 0.5 for all values of  $x_1$  and  $x_2$ .

$$F(x_1, x_2) = \begin{cases} 0.5 & \{x_1 > 0\} \\ \sin(x_2) & \{otherwise\} \end{cases} \quad (1)$$

Figure 3 demonstrates a scatterplot of  $x_1$  against the instability of the decision procedure with respect to  $x_1$ . In this case, it is visually obvious that the uncertainty associated with  $x_1$  becomes a problem around the threshold value of 0. As  $x_1$  moves away from the threshold in either direction, the DP is fully stable.

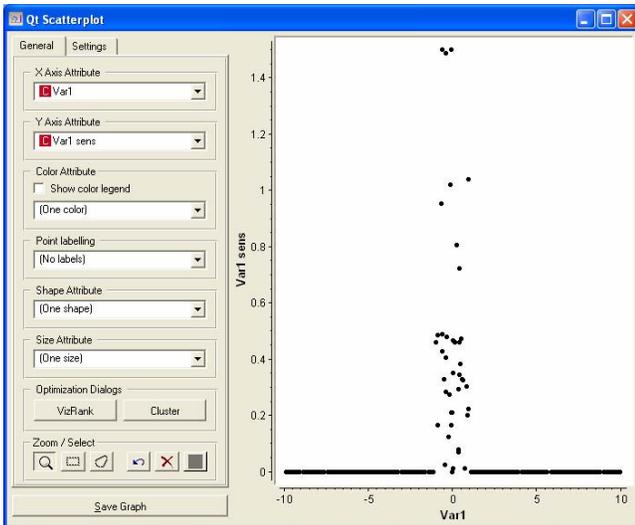


Figure 3: Scatterplot example. The function  $F(x_1, x_2)$  in eq. (1) is sensitive to  $x_1$  near  $x_1=0$ .

This sort of structure is seen often even in more realistic DPs. For instance, a DP that determines when to fire artillery at an enemy target is only unstable near the maximum range of the artillery. If the enemy target is far enough outside or inside the range of the artillery, the decision is stable. Near the boundary, uncertain information can produce incorrect decisions.

We often, however, want to view the interaction of more than two variables at a time on the scatterplot. Plots with three axes are one approach. We can also add additional variables through other drawing features of the data points, such as the color, shape, or size of each point. Figure 4 provides an example of this approach. In this case, we have plotted  $x_2$  on the x-axis and stability with respect to  $x_2$  on the y-axis. We have also included the  $x_1$  value as the saturation value of the point on the graph. We see in this figure, that light-colored dots appear at the bottom of the graph while dark-colored dots create arc shapes above. This indicates that the instability value is low whenever  $x_1$  is low, which is true as  $x_2$  is ignored and  $x_1$  is stable (since we are plotting local sensitivity with respect to  $x_2$ ). It also indicates that when  $x_1$  is sufficiently high, the shape of the instability curve is an arc with a period of  $\pi$ .

We have also evaluated parallel coordinate graphs (Inselberg 2007), which display each input variable and the

sensitivity values as separate parallel vertical lines. Each point in the global analysis represents a line that passes through the corresponding values on each of the parallel coordinates. This approach scales to significantly more dimensions and is a common tool for sensitivity analysis visualization. We have, so far, found that scatterplots are simpler to read and understand and with the approaches described above, scale to sufficient dimensions to provide analysts with a powerful visualization tool for many problems.

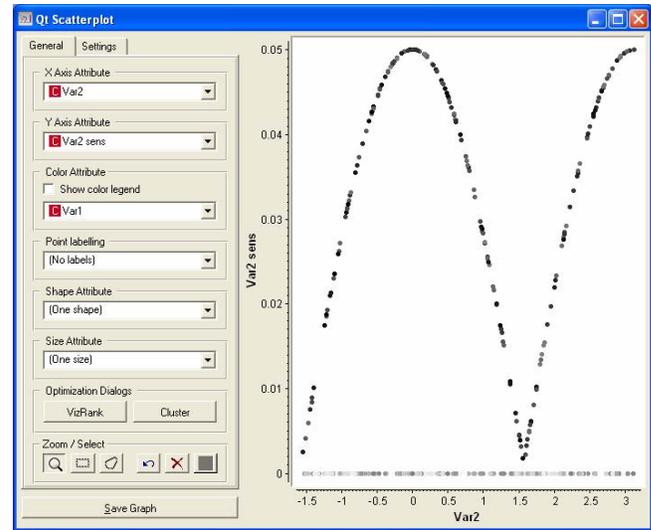


Figure 4: Scatterplot with three dimensions. The function  $F(x_1, x_2)$  in eq. (1) is periodically sensitive to  $x_2$  when  $x_1$  is sufficiently large. The value of  $x_1$  is indicated by saturation.

## 5 FUTURE WORK

We break our discussion of future work into two parts. First, we describe our planned efforts to apply this analysis approach to support a number of different tasks within the U.S. Army’s MATREX simulation-based analysis architecture. Second, we describe functionality that we plan to implement in future versions of the SURE tool to provide additional power, efficiency, and usability.

### 5.1 MATREX Support

The U.S. Army’s MATREX (Modeling Architecture for Technology and Research EXperimentation) is a simulation-based architecture for modeling complex, network-centric operations (Kellihan et al. 2006). The goal is to support the development, evaluation, and acquisition of new technologies and tactics, techniques, and procedures (TTPs) that are required to effectively use the newly developed technologies. MATREX uses the C3HPM human performance model to model the human decision makers within these simulations (Kellihan and Washington 2004).

The MATREX tools and models are still being refined, but we have designed a three-part integration plan for using SURE functionality to improve the analysis power of MATREX. In this section, we describe this planned work. Since SURE has only been tested using small-scale DPs, this discussion is meant to suggest ways that SURE can be used in practical, large-scale simulation systems. We are exploring three possible uses for the SURE tool within MATREX:

- **Doctrine Analysis.** The models that we will be evaluating with the SURE tool are based on current military doctrine as modeled by TRADOC using the C3HPM modeling architecture. Many of these doctrines are well-established. Others are newer doctrines and TTPs designed to work with the latest technologies being adopted by the Army. We plan to use SURE to provide an additional level of evaluation for these doctrines to ensure that they are robust enough to be effective even in situations with uncertain and missing information, such as the battlefield.

We are developing an off-line analysis routine that uses SURE to generate inputs and evaluate the resulting outputs of individual C3HPM decision models. Such a system is obviously not a complete solution as these decisions will eventually be made in a larger decision-making context that evolves over time and that this tool is not designed to evaluate. It does, however, provide a means of evaluating individual decisions to ensure that the larger, temporally extended decision-making processes are based on sound building blocks.

- **Model debugging.** The C3HPM models that represent the decision-making processes used on the battlefield are complex and are the result of a complex, multi-stage development process. It is expected that not all of the models will be perfect representations of the underlying doctrinal decisions that they are intended to capture. In particular, one issue that is a concern in this effort is that the models are complete; that is, do they respond to the full set of situations that they are expected to be able to handle. In this case, the situations are typically messages containing some information or request that needs to be responded to. The Monte Carlo testing approach that we described above provides a means of evaluating these models in a wide range of different situations, including an evaluation of each type of message and a wide variety of message combinations and message contents. This provides the developers with a quick, off-line means of testing and improving the completeness of the models
- **Scenario development.** Running the full MATREX simulation tool and analyzing the results requires more effort and time than is required by the SURE tool, which is able to process many possible input states and does not need to wait for each situation to arise during the course of a simulation run. To this end, we are

evaluating the use of the SURE tool as a means of suggesting simulation scenarios that are most likely to cause difficulties for the technologies and TTPs being evaluated. It is valuable to identify such problematic situations during a simulation run, before similar situations are found during live testing or even on the battlefield. By using SURE to identify potential areas of decision instability, we can help with this early testing process.

## 5.2 Additional Future Work

The effort described in this paper has focused on evaluating the effect of uncertain information. A similar approach can be used to evaluate the effect of missing information. In this case, instead of evaluating the error at the extremes of the error range, we can sample the possible values of the missing information and use those to determine the maximum error caused by the lack of information. This same approach is also applicable to uncertain information in cases where we expect the range of uncertainty to be great enough that our assumption of local linearity is no longer valid.

Also, in the SURE tool, the analyst has full control over which data plots they want to see. With a large number of input variables, however, the number of plots can be large, especially when the analyst wants to evaluate the interrelationships between input values. To assist in this problem, SURE performs a pairwise linear regression of each pair of input values to determine which input pairs are related. We are also working on a mechanism that will suggest appropriate visualizations for the regressions that suggest detected correlations between input variables. For instance, SURE might note in the example in Section 4 that  $x_1$  and  $x_2$  have some non-random relationship and suggest the presentation in Figure 4 as being useful to the analyst.

Another issue that we expect to have to address is efficiency. While the approach we described has already been significantly simplified from the theoretically optimal solution for this purpose, as we aim for more accurate results, efficiency will reemerge as an issue. There are a variety of methods available, including Response Surface Models (Lim et al. 1989), that provide a more efficient mechanism for performing the large number of Monte Carlo simulation and runtime evaluations of a situation to support runtime decision making.

We are working on the design of a system that will provide support for more complex decision-making processes than those described in Section 2. While the simple model we described is quite powerful, it is also useful to describe many decisions in a hierarchical manner, where some decisions are made in support of other, more complex decisions. While each of these decisions can be viewed as an instance of a DP model as described above, there is additional structure that we expect will be useful to

the analyst evaluating these decision structures. We are investigating analyses and interfaces that support the analysis of such decision structures.

Finally, we have adopted a system that assumes that the analyst is most interested in knowing how bad a decision can be in the face of existing uncertainty. In real problems, it is also sometimes useful to know how likely those bad cases are. The proposed approach can readily be extended to support such an analysis, though the difficulty of creating probability distributions for the uncertainty associated with each input raises issues of usability and practicality that will need to be addressed.

## 6 RELATED WORK

We have already mentioned a number of relevant areas of related work. Sensitivity Analysis, which measures the sensitivity of a complex model to the inputs, is obviously closely related and we have adopted a number of ideas from Sensitivity Analysis. None of the existing approaches to Sensitivity Analysis were directly applicable to our problem, so we adapted the borrowed ideas to provide a hybrid solution that effectively provides a global analysis of a local space or over a representative sample of local spaces.

Work in Situation Awareness is also closely related to this effort. Situation Awareness is concerned with understanding what human decision makers need in terms of information and providing that information in a way that it is easily understood and usable by the human (Endsley 1992). Situation Awareness is largely concerned with topics such as information and display design. The work described in this paper provides a means of evaluating what Pew refers to as *Ideal SA* (Pew 2000). This is the best level of information available to the decision maker. By using computational models of decision processes, we can provide a very clear view of exactly what level of information is needed to reliably make any particular decision. The issues involved in presenting information to the human in a clear manner are still important issues that our work does not attempt to address. We believe, however, that using computational decision-making models can provide additional guidance to those designing human-system interfaces. For instance, Endsley (Endsley 2000) describes an effort to improve the interface for a fighter aircraft system that increased the pilots' Situation Awareness, but it did not increase their mission performance because the tactics that the pilots used were not affected by the additional information. A model of the tactics built ahead of time and evaluated using the techniques described in this paper could have theoretically predicted such a result.

Other relevant work in Situation Awareness includes some efforts at computational models (e.g., (McCarley et al. 2002; Burdick and Shively 2000; White, Young, and Kelsch 1998; Shively, Brickner, and Silbiger 1997)) but

the results are often domain specific and do not provide a direct way of mapping Situation Awareness to the expected effects on particular decisions or in particular situations.

Finally, work in Decision Theory (Von Neumann and Morgenstern 1944) and Information Value Theory (Howard 1966) are closely related. Decision Theory evaluates the optimal decision in a given situation given a network of costs, benefits, and probable outcomes of actions. Information Value Theory builds on Decision Theory and provides a means of measuring the effect of additional information on the decision making process. In many ways such efforts are very similar in their objectives. The advantage of the approach described in this paper is that it works more effectively with continuous decision functions. Decision Theory assumes all variables are in any of a finite number of states with a probability and that the probabilities are either provided or their relationship to other variables are provided. The SURE system supports this with continuously valued variables and does not assume that the probabilities of variable values are known, just that the range of uncertainty is known; we have found this information is typically easier to get than a probability distribution representing the uncertainty.

## ACKNOWLEDGMENTS

The work described in this paper was funded under U.S. Army Contract W911QX-06-C-0087. We would like to thank Carolyn Dunmire for her support and guidance during this effort. We would also like to thank Karen Harper, Randy Washington, Dina Dalton, Peter Harding, and Zack Henson for their efforts in support of this work.

## REFERENCES

- Burdick, M. D. and R.J. Shively. 2000. *Evaluation of a computational model of situational awareness*. NASA Ames Research Center.
- Endsley, M. 2000. Direct Measurement of Situation Awareness: Validity and Use of SAGAT. In M. Endsley & D. J. Garland (Eds.), *Situation Awareness Analysis and Measurement*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Endsley, M. R. 1992. Situation Awareness: The development and application of a theoretical framework. In *Proceedings of Human Factors and Ergonomics Society 37th Annual Meeting*.
- Howard, R. 1966. Information Value Theory. *IEEE Transactions on Systems Science and Cybernetics*, SSC-222-26:22-26.
- Inselberg, A. (2007). *Parallel coordinates: visual multidimensional geometry and its applications*. Springer.
- Kellihan, B. L., Washington, R. (2006). Human-centric modeling and simulation of network-enabled battle

- command. In *Proceedings of Behavior Representation in Modeling and Simulation (BRIMS)*. Baltimore, MD.
- Kellihan, B. and R. Washington. 2004. C3HPM: An application of IMPRINT within the MATREX. In *Proceedings of 2004 Conference on Behavior Representation in Modeling and Simulation (BRIMS)*.
- Lim, J. T., H.J. Gold, C. G. Wilkerson, C. D. Raper, Jr. 1989. A Monte Carlo/response surface strategy for sensitivity analysis: application to a dynamic model of vegetative plant growth. *Applied Mathematical Modelling*, 13:479-484.
- McCarley, J. S., C. Wickens, J. Goh, W. J. Horrey. 2002. A computational model of attention/situation awareness. In *Proceedings of 46th Annual Meeting of the Human Factors and Ergonomics Society*. Santa Monica, CA.
- Morris, M. D. 1991. Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33:161-174.
- Pew, R. W. 2000. The state of situation awareness measurement: heading toward the next century. In *Situation Awareness Analysis and Measurement*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Saltelli, A., K. Chan, E. M. Scott. 2000. *Sensitivity analysis*. West Sussex, England: John Wiley & Sons, Ltd.
- Saltelli, A., S. Tarantola, F. Campolongo, and M. Ratto. 2004. *Sensitivity analysis in practice*. West Sussex, England: John Wiley & Sons, Ltd.
- Shively, R. J., M. Brickner, and J. Silbiger. 1997. *A computational model of situation awareness instantiated in MIDAS*. Moffett Field, CA: Army/NASA Rotorcraft Division.
- Von Neumann, J. and O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton: Princeton University Press.
- White, D. A., L. Young, G. R. Kelsch. 1998. Modeling and simulation of situational awareness in the tactical internet. In *Proceedings of Military Communications Conference (MILCOM 1998)*.

## AUTHOR BIOGRAPHIES

**W. SCOTT NEAL REILLY** is a Principal Scientist at Charles River Analytics. Dr. Neal Reilly's research interests include human behavior modeling architectures, models of human emotional processes, and scalable tools for developing human behavior models. Dr. Neal Reilly received his B.S. in Computer Science from Duke University and his M.S. and Ph.D. in Computer Science from Carnegie Mellon University.

**SEAN GUARINO** is a Senior Software Engineer at Charles River Analytics. His research interests include the design and development computational models of human cognition and behavior, distributed systems of interactive

intelligent agents, and advanced adaptive interface design. Mr. Guarino earned his B.S. at Cornell University majoring in Computer Science, with a concentration in Cognitive Studies, and is currently pursuing his Masters in Liberal Arts in Psychology at Harvard.

**BRET KELLIHAN** is a Modeling and Simulation Engineer at DCS Corporation. He leads the design of constructive human performance modeling and decision making tools for use in distributed simulations. Mr. Kellihan has been the lead developer on several IMPRINT based simulations for the Army Research Laboratory and the Tank Automotive Research, Development and Engineering Center. His interests include the fields of Knowledge Representation and Human Behavior Modeling. Mr. Kellihan received his B.S. in Electrical Engineering from the University of Virginia with concentrations in Digital Systems and Applied Electrophysics, and a minor in Computer Science.