

MODELLING AND IMPROVING HUMAN DECISION MAKING WITH SIMULATION

Stewart Robinson
Thanos Alifantis
Robert Hurrion

Operational Research and Systems Group
Warwick Business School
University of Warwick
Coventry, CV47AL, UK

John Ladbrook

Dunton Engineering Centre
(15/4A-F04-D), Laidon, Basildon
Ford Motor Company
Essex, SS15 6EE, UK

John Edwards

Management Science Group
Aston Business School
Aston University
Birmingham, B47ET, UK

Tony Waller

The Oaks, Clews Road
Lanner Group
Redditch, Worcs, B98 7ST, UK

ABSTRACT

Modelling human interaction and decision-making within a simulation presents a particular challenge. This paper describes a methodology that is being developed known as 'knowledge based improvement'. The purpose of this methodology is to elicit decision-making strategies via a simulation model and to represent them using artificial intelligence techniques. Further to this, having identified an individual's decision-making strategy, the methodology aims to look for improvements in decision-making. The methodology is being tested on unplanned maintenance operations at a Ford engine assembly plant.

1 INTRODUCTION

Most operations systems include significant elements of human decision-making and interaction. A manufacturing plant may involve any number of manual processes, and human decision-making may affect many aspects of the plant's operation such as scheduling of production and maintenance operations. The situation is normally more exaggerated in service operations where service personnel interact with customers.

Modelling these elements of an operations system presents a number of challenges. Most simulation tools do not provide a full set of functions that would be useful for modelling human behaviour. A more fundamental challenge is determining the behaviour of the human actors within the operations system. Further to this, the purpose

of the simulation exercise may be to establish if the human decision-making process could be improved. This requires a methodology for determining, modelling and looking for improvements in human interaction.

The purpose of this paper is to outline a methodology that is being developed to help model and improve human decision-making. The methodology, known as 'knowledge based improvement', is based on the use of visual interactive simulation (VIS) with artificial intelligence (AI) approaches, particularly neural networks and expert systems. The viability of the concept is being established by modelling an actual engine assembly facility at a Ford plant in Wales. Following a brief review of previous work on modelling human decision-makers, the background to the Ford case study is described. The stages in the knowledge based improvement methodology are then outlined, before progress to date on its implementation is described.

2 SIMULATING THE HUMAN DECISION-MAKER

The typical approach to representing human decision making in simulation models is to try to elicit the decision rules from the decision maker. In some cases this amounts to little more than a guess on behalf of the modeller. Following this, the rules are included in the model using the constructs of the simulation language or simulator. This normally requires the use of a series of 'if', 'then', 'else' statements, which can result in an excessive amount of code that is difficult to interpret.

One approach to overcoming these problems might be to use AI to represent the human decision maker, and link it with a simulation model. Indeed, some have already attempted to do this (Flitman and Hurrion, 1987; O’Keefe, 1989; Williams, 1996; Lyu and Gunasekaran, 1997, Robinson et al., 1998). This could be implemented in two ways:

- elicit the decision rules from the expert and represent them within an expert system
- use the simulation model to prompt the expert to make decisions, building up a set of examples from which an AI system could learn

These correspond to the two fundamental approaches to knowledge acquisition: elicitation by human knowledge engineer and machine learning from examples, respectively. The first approach would employ the constructs of an expert system and so make it easier to accurately represent the decision process than using simulation language constructs. It should also be easier to interpret and easier to change since expert systems are specifically designed to facilitate this. In this way the approach should aid model development and improvement to decision-making. What it does not provide, however, is a simple means for knowledge elicitation. This remains a well-known problem in AI generally (Waterman, 1986). It is in the second approach where the link to a simulation model could provide significant advantages. Most work on machine induction (e.g. Hart (1987)) treats the set of examples as somehow “given”, and devotes little or no discussion to the process of obtaining the examples. By getting the simulation model to present the human decision maker with realistic conditions and asking for a decision, a set of examples could be obtained at an accelerated speed (assuming the model runs faster than real-time!). In this way the approach acts as an aid to obtaining input data, that is, the process by which a human decision-maker works.

3 CASE STUDY: FORD BRIDGEND ENGINE ASSEMBLY PLANT

The Ford engine plant at Bridgend is one of the main production facilities for the ‘Zetec’ petrol engine. The plant consists of a number of transfer lines (Ladbrook, 1998) that feed the main engine assembly line. In engine assembly, blocks are placed on a ‘platen’ and pass through a series of automated and manual processes. For the purposes of this research, the maintenance operations on a self-contained section of the engine assembly line are considered.

Prior to this research one of the authors (Ladbrook) had already developed a simulation model of the complete engine assembly facility. The model, developed in the WITNESS simulation software (Lanner Group, 2000), was used to identify bottlenecks and to determine viable operat-

ing alternatives. The maintenance logic in the model assumed that when a machine fault occurred, the decision would be to make an immediate repair. Random sampling was used to determine the skill level of the engineer required to service the fault. These assumptions were considered to be adequate for the purposes of the study that was performed.

In the early stages of this research one of the authors (Alifantis) spent some time observing the production facility and in discussion with plant engineers. This showed that reality is somewhat different from what is represented within the simulation. Although the obvious action to take when a machine breaks down is to repair it immediately (RI) this may not always be the most appropriate action for a variety of reasons:

- *Inappropriate*: If there is a long queue of parts downstream from the machine requiring repair, then immediate repair may not be the most appropriate action, and the maintenance engineers may be better deployed elsewhere.
- *Insufficient*: Repairing a machine takes time. Meanwhile the rest of the production facility continues to process parts and to move them around. This means that during the repair of the machine queues may occur upstream, while downstream the process will be starved of parts. Simply repairing the machine may be insufficient to reach target throughput.
- *Impossible*: Sometimes it may not be possible to repair the machine immediately since all the maintenance engineers are busy. There is always the option to interrupt the repair of another machine and so to release one of the engineers, but this may not be the best course of action. Further to this, on occasions spare parts required for the repair of the machine may not be available.

From the above it is clear that apart from repairing immediately other policies should be considered when a machine breaks down.

Stand-by (SB) can be considered as an alternative to repair immediately. In this case an engineer processes the parts manually and pushes them to the next machine through the conveyor. In general it is not possible to repair the machine at the same time as stand-by is being operated because of space restrictions. The type of fault, the extent of queues and labour availability among other attributes are the key determinants of this decision.

Stop the line (SL) is another option, which might be considered as a complement or substitute to repair immediately. In this case the maintenance supervisor (‘group leader’) should decide whether it is useful to stop the whole line or part of it. This might be used, for instance, to

avoid a build up of work-in-progress in a section of the line.

Do nothing (DN) is an alternative decision that might be the desired course of action under certain circumstances, for instance, close to the end of a shift. Obviously this decision must be revised eventually and the machine repaired later.

Beyond the above list of options, the group leader may want to revise the decision for a particular machine at a later point. For example, the group leader might decide to stop the repair of a machine because he needs the labour to repair another machine that is down.

Although it is possible to identify the types of decisions the group leaders might make, it is another to determine how those decisions are taken. In determining what course of action to take the group leaders rely upon their knowledge and experience (tacit knowledge). Direct questioning of group leaders showed they are unable to directly express this knowledge. Therefore, the methodology that underpins this research, first aims to elicit information on the group leaders' decision-making strategies and then to look for ways to improve their decision-making. This methodology is now described.

4 METHODOLOGY: KNOWLEDGE BASED IMPROVEMENT

The methodology outlined in this paper, known as knowledge based improvement, consists of five key stages:

- Understanding the decision-making process
- Data collection
- Determining the experts' decision-making strategies
- Determining the consequences of the decision-making strategies
- Seeking improvements

Each of these stages is now described.

4.1 Stage 1: Understanding the Decision-Making Process

The first step in determining the experts' decision-making strategies is to identify the component parts of the decision-making process: decision variables, decision options, attribute variables and attribute levels. For instance, in a simple maintenance scheduling problem where there are 2 actions, that are not mutually exclusive, and 2 engineers who can be asked to act if they are available; as such, there are 4 *decision variables*. The first 2 variables correspond to the actions and the other 2 to the engineers. Each of them has 2 alternative *decision options*: the action can either be taken (denoted 1) or not taken (denoted 0); the engineer can either be asked to act (denoted 1) or not asked to act (denoted 0). Assume, for the simplicity of the example,

that the decisions are determined taking into account an estimate of the repair time and the type of fault. It is clear that there are two *attribute variables* in the decision-making process. The first attribute can take the value of any real number that represents the estimated repair time, albeit that in practice it is likely to be rounded to at least the nearest whole number. The second attribute (type of fault) can take values that represent the code of any particular fault. The range of estimated repair times and number of fault codes define the *attribute levels*.

Although interviews and discussion with the decision-makers can reveal some information about the decision-making process, usually the human expert cannot explicitly identify and list the decision-making components. To do so the modeller should observe the human experts as they take decisions. In addition, in order to build a complete model of the decision-making process the modeller may need to make assumptions by considering other rational decision that can be taken by the decision-maker.

A decision-making process can be represented as two row vectors. The first vector corresponds to the decision, with each element representing a decision variable (*d*). The second vector corresponds to the attributes of the decision, with each element representing an attribute variable (*a*). In the context of the simple maintenance scheduling example described above, the decision-making process can be represented as follows:

$$\mathbf{D}_{i,j} = f_j(\mathbf{A}_i) \quad (1)$$

Where:

$$\begin{aligned} \mathbf{D}_{i,j} &= [d_1 \quad d_2] \\ \mathbf{A}_i &= [a_1 \quad a_2] \end{aligned} \quad (2)$$

The subscript *i* indicates the time at which the decision was taken and the subscript *j* indicates the human expert that took the decision. The function *f* represents the decision-making strategy of the individual expert, taking into account the attributes of the vector *A_i*. The purpose of stages 1 to 3 of the methodology is to determine the function *f* by applying AI techniques to a set of collected example decisions.

4.2 Stage 2: Data Collection

Having identified the decision components, the next step in determining the decision-making strategies is to collect examples of decisions from each expert. Each example in the data set should include the value of each decision and attribute variable. The data set should have the form of two matrices: **D_j** and **A**. **D_j** represents the decisions made by decision-maker *j* under specific attribute values (identified in **A**). Each row of the matrix **D_j** corresponds to the row vector **D_{i,j}**, that is, the decisions taken at time *i*. Each col-

umn in the matrix **D** corresponds to a decision variable. Each row of the matrix **A** includes the attribute values at a particular decision point (i). Each column corresponds to an attribute variable.

For the simple decision-making process outlined above, the data set to be used in determining the decision-making strategy of expert j should have the following form:

$$D_j = \begin{bmatrix} d_{1,1} & d_{1,2} \\ \cdot & \cdot \\ d_{i,1} & d_{i,2} \\ \cdot & \cdot \\ d_{l,1} & d_{l,1} \end{bmatrix} \quad A = \begin{bmatrix} a_{1,1} & a_{1,2} \\ \cdot & \cdot \\ a_{i,1} & a_{i,2} \\ \cdot & \cdot \\ a_{l,1} & a_{l,1} \end{bmatrix} \quad (3)$$

One method of collecting these data would be through observation of the experts at work. This, however, would be extremely time consuming, particularly if the elapsed time between decision points is large. It would also be difficult to record the full set of many attribute values at a specific moment in time, and because the values are likely to change continuously, inaccuracies would occur if there were any delay. As a result, the methodology uses a VIS. The expert interacts with a visual simulation of the system in question. The simulation model stops at a decision point and reports the values of the attribute variables. The expert is then prompted to enter his/her decision to the model. The model records the value of each decision variable and attribute variable to a data file. As a result a set of values for the matrices D_j and **A** are collected.

The methodology suggests the use of VIS for a number of reasons. First, it is less time consuming than observation, because the simulation runs much faster than real time. Second, because the simulation stops at a decision point, it is possible to capture all attribute values at that moment in time. A third benefit is that a simulation run can be replicated exactly, enabling the system state to be interrogated further at a later date, should this be required. For instance, it may become apparent that the decision-maker takes into account attributes that have not previously been identified. This also provides the benefit that different decision-makers can be presented with the same series of attribute values.

Of course the use of VIS as a data collection method is not free of problems. Three specific difficulties arise. First, the model needs to contain and report all the key attributes in the decision-making process. This probably requires a very detailed model which in itself could be time consuming to develop. Accurate data, required to support such a detailed model, may not be available either. A second problem is the need to involve the human decision-maker in entering decisions to the model. A very large number of example decisions may be required to obtain a

full set of data, which in itself could be time consuming. A third problem is whether the human decision-makers are likely to take realistic decisions in a simulated environment. It is quite likely that they will take greater risks, as there are no real consequences to their decisions.

Anticipating the second problem, the methodology suggests the collection of a limited number of decision examples using the VIS. This data set is not expected to be large enough to be able to determine the decision-making strategy of an expert. It is believed, however, that it could be used to train a neural network model. This in turn could be used to increase the example set D_j and **A**. A feed forward neural network with 3 layers is believed to be computationally sufficient for this purpose. A separate neural network would, of course, need to be trained for each decision-maker.

4.3 Stage 3: Determining the Experts' Decision-Making Strategies

Having collected a series of examples using the VIS and the neural network model, the next step is to use the data in the matrices D_j and **A** to determine the decision-making strategies of the individual experts. A decision-making strategy can be represented by the use of a decision tree; a separate decision tree being constructed for each decision-maker. Experts systems software is capable of constructing a decision tree from a set of examples, such as those collected via the VIS and neural network. One such method for constructing a decision tree is Quinlan's ID3 algorithm; see, for example, Mingers (1986). The algorithm prioritises the attributes according to the degree to which they match the data set with the correct decisions.

Figure 1 shows an example of such a decision tree for the simple maintenance problem described above. This assumes that each attribute has only two levels L1 and L2.




Figure 1: Decision Tree for Simple Maintenance Scheduling Example

4.4 Stage 4: Determining the Consequences of the Decision-Making Strategies

Having determined the decision-making strategies, that is, a decision model f_j for each expert j , the next step in the knowledge based improvement methodology is to assess and compare the performance of each expert. The ultimate performance measure in most manufacturing facilities is the level of throughput. This means that each expert can be assessed on the basis of the throughput that is achieved in the simulation when the decision-making process is controlled using his/her decision-making strategy. To predict the throughput, conditional to each human expert, the VIS can be linked with the expert systems software (or, indeed, the neural network).

The expert system is used in place of a decision-maker to interact with the simulation. Each time the simulation reaches a decision point the simulation stops and the expert systems software is invoked. The value of each decision attribute is passed from the simulation to the expert system software. In turn, the expert system returns the values of the decision variables to the simulation before the simulation run is continued. For a description of how to link a simulation package with expert systems software see Robinson et al (2001).

When the simulation has reached the end of the run, the throughput of the production line provides an indicator of the performance of the expert whose decision tree was used during the run. Running the simulation under each expert's decision-making strategy for a number of replications to eliminate stochastic variability, enables the most efficient strategy to be found by comparing the output from each run.

Of course having identified the most efficient expert does not mean that the most efficient strategy has been found since there is no guarantee that the best current strategy is the optimal one. Although the best strategy may not be optimal it can still be used to train less efficient decision-makers, providing improvements in overall performance.

4.5 Stage 5: Seeking Improvements

The last stage in the methodology uses the decision-making strategies of the most efficient experts as starting point to search for an improved strategy. The search could be made informally by combining strategies and by making incremental changes. Alternatively, heuristic search methods could be implemented, in order to seek for improvements. In each case, the alternative strategies can be tested by running them with the simulation model in order to determine their effectiveness.

5 IMPLEMENTING THE METHODOLOGY ON THE FORD ENGINE ASSEMBLY PLANT

Having described the knowledge based improvement methodology, the progress in applying it to the Ford engine plant at Bridgend is discussed. Although there is a simulation model of the line, no study has been undertaken to evaluate the effectiveness of alternative maintenance strategies.

As already stated, the methodology is being implemented by considering a self-contained section of the production facility. In this section a team of 5 engineers (2 skilled, 2 semi-skilled and 1 unskilled engineer) perform unplanned maintenance, among other tasks. One of the skilled engineers acts as group leader for that part of the line. One of his duties is to decide what to do when a machine breaks down. When a failure occurs a message is received via a pager that is carried at all times. The message reports the name of the equipment and a short description of the fault. After an inspection of the machine the group leader decides what action is the most appropriate. The main options available are the following:

- SB: Stand-by
- RI: Repair immediately
- SMLB: Stop the main line before the machine
- SMLA: Stop the main line after the machine
- SSB: Stop the section before the machine
- SSA: Stop the section after the machine

Alternatively, the group leader may decide to do nothing, which is equivalent to a null response to all these options.

Having decided what course of action to take, the group leader should decide who must act. Available engineers who can be asked to act if they are available are the following:

- L1: Group leader
- L2: Second skilled engineer ('M/Elec')
- L3: Semiskilled engineer 1 ('IMS1')
- L4: Semiskilled engineer 2 ('IMS2')
- L5: Unskilled engineer.

Given the above discussion the row vector Di,j in this particular decision-making process should include the following elements:

- SB, RI, SMLB, SMLA, SSB, SSA, L1, L2, L3, L4, L5

Each element of the above matrix represents a decision variable that can take the value 0 or 1. Zero means that the decision-maker is not taking the particular action or that the particular resource will not be asked to act. On the contrary, 1 means that the decision-maker is taking the par-

not report extremely important attributes such as the estimated repair time for each fault. In the second session it was found that the breakdown scenarios reported were very similar (the estimated repair time for all of them was very short), so the decision by the human expert was the same at every decision point (repair immediately). As a result of these findings, improvements have been made to the model: the estimated repair time is now reported, and the trace data on breakdowns has been adjusted to provide a wide range of scenarios. The next step is to collect example decisions from a number of experts before moving to the AI stage in the methodology.

6 FINDINGS AND DISCUSSION

Having almost completed the stage in the methodology that requires the use of a VIS for collecting example decisions, a number of strengths and weakness in the methodology have been identified.

First of all it has been found that VIS is probably a unique way of obtaining data about decisions in a reasonable time interval, since the simulation runs faster than real time. In addition, VIS is a very efficient approach for data collection since it is an experimental environment where the modeller can control the values of the attributes that are generated and reported to the user of the model. This means that the modeller can guarantee that a wide range of attribute value combinations can be created during a data collection exercise. There is no such control in a real life data collection exercise.

Although the approach is quite promising, some problems have been identified in its application. First of all, it is difficult to isolate and understand the decision-making surrounding unplanned maintenance since the team performs many other tasks as well. Indeed, the team find it difficult to articulate how they make decisions concerning only this part of their work. Another challenging point is the fact that the simulation model may not be capable of reporting all the attributes that the decision-maker takes into account. In the above application, for example, the decision-maker takes into account the physical condition of the machine. This attribute is all but impossible to simulate in a model of this nature.

7 CONCLUSION

The paper has explained the development of a methodology that can be used to evaluate the effectiveness of decision-makers. In addition it proposes a method to improve the current decision-making practices based on the experts' extant knowledge. The methodology described in the paper includes an innovative data collection method based on the use of a VIS combined with a neural network algorithm.

The methodology is being tested with unplanned maintenance operations in a Ford engine assembly line. Initial data on example decisions has been collected via a simulation model. The next stage in the research is to collect a full set of data for different maintenance supervisors. Following this, the use of AI for learning and improving current decision-making strategies is to be explored.

ACKNOWLEDGEMENTS

This work is being jointly funded by the EPSRC (grant reference GR/M72876), Ford Motor Company and the Lanner Group.

REFERENCES

- Flitman A.M. and Hurrion, R.D. (1987). Linking Discrete-Event Simulation Models with Expert Systems. *J. Opl Res. Soc.*, **38** (8), pp. 723-734.
- Hart, A. (1987). Role of Induction in Knowledge Elicitation. In *Knowledge Acquisition for Expert Systems: a Practical Handbook*, ed. A. Kidd, 165-189. New York: Plenum.
- Ladbrook, J. (1998). Modeling Breakdowns: An Inquest. MPhil Thesis, University of Birmingham.
- Lanner Group (2000). WITNESS 2000 User Manual. The Oaks, Clews Road, Redditch, UK.
- Lyu, J. and Gunasekaran A. (1997). An Intelligent Simulation Model to Evaluate Scheduling Strategies in a Steel Company. *International Journal of Systems Science*, **28** (6), pp. 611-616.
- Mingers, J. (1986). Expert Systems: Experiments with Rule Induction. *J. Opl Res. Soc.*, **37** (11), pp. 1031-1037.
- O'Keefe, R.M. (1989). The Role of Artificial Intelligence in Discrete-Event Simulation. *Artificial Intelligence, Simulation and Modeling* (L. E. Widman, K.A. Loparo and N.R. Neilsen, eds.), pp. 359-379. Wiley, NY.
- Robinson, S., Edwards, J.S. and Yongfa, W. (1998). An Expert Systems Approach to Simulating the Human Decision Maker. *Winter Simulation Conference 1998* (D.J. Medeiros, E.F. Watson, M. Manivannan, J. Carson, eds.), The Society for Computer Simulation, San Diego, CA, pp. 1541-1545.
- Robinson, S., Edwards, J.S. and Yongfa, W. (2001). Linking WITNESS to an Expert System to Represent a Decision-Making Process. *Simulation* (in review).
- Waterman, D.A. (1986). *A Guide to Expert Systems*. Addison-Wesley, Reading, Mass.
- Williams, T. (1996). Simulating the Man-in-the-Loop. *OR Insight*, **9** (4), pp. 17-21.

AUTHOR BIOGRAPHIES

STEWART ROBINSON lectures in Operational Research and Systems at the Warwick Business School in the UK. He holds a BSc and PhD in Management Science from Lancaster University. Previously employed in simulation consultancy, he supported the use of simulation in companies throughout Europe and the rest of the world. His research interests are in finding ways to improve the use of operational research, and specifically simulation, within industry. He is author/co-author of two books on simulation. Current work involves an investigation into the use of expert systems to represent a human decision-maker in simulation models, the analysis of complex output behaviours from simulations, and developing an understanding of quality in relation to operational research and simulation studies. His email and web addresses are stewart.robinson@warwick.ac.uk and www.btinternet.com/~stewart.robinson1/SLRHOME.HTM

THANOS ALIFANTIS is a PhD researcher at Warwick Business School in the UK. He holds a BSc in Economics and Operational Research from Lancaster University. He is investigating the use of artificial intelligence for representing human decision-makers in simulation models. His research interests are in econometrics, optimization, heuristic techniques (simulated annealing, taboo search and genetic algorithms) for timetabling problems, simulation, data envelopment analysis, visual basic programming, game theory, direct rule extraction from artificial neural networks and analysis of the structure of feed forward neural networks. His email and web address are a.alifantis@warwick.ac.uk and <http://www.warwick.ac.uk/staff/A.Alifantis/>

ROBERT D. HURRION is a reader at Warwick Business School in the UK. He is the originator of the visual interactive simulation technique. He has acted as a consultant to organisations such as ICI and the Lanner Group to help them in their development of visual interactive simulation systems such as Vision, See Why and WITNESS. He has published extensively on the use of visual simulation for decision support systems and intelligent simulation methods that link rule based expert systems and simulation. More recently his research has investigated the potential of neural network technology for artificial intelligent simulation. His email address is r.d.hurrion@warwick.ac.uk

JOHN S. EDWARDS is Senior Lecturer in Operational Research and Systems at Aston Business School, Birmingham, UK. He holds MA and PhD degrees (in mathematics and operational research respectively) from Cambridge University. His principal research interests are knowledge-based systems and decision support systems, especially

methods for their development. He has written more than 30 research papers on these topics, and two books, *Building Knowledge-based Systems* and *Decision Making with Computers*, both published by Pitman. Current work includes a study of the use of artificial intelligence techniques in management accounting. His email address is j.s.edwards@aston.ac.uk

JOHN LADBROOK has been involved in computer simulation with Ford since the early 1980s. In that time he has developed, or managed the development of, numerous simulation models. He has also been involved in a number of collaborative research projects with Cranfield, Birmingham, Bath and Cambridge universities. John Ladbrook recently completed an MPhil at the University of Birmingham, School of Manufacturing and Mechanical Engineering, on the modelling of breakdowns. This work identified a series of difficulties with modelling breakdowns in simulation software, as well as proposing potential solutions. His email address is jladbroo@ford.com

TONY WALLER has been involved in the field of Computer Simulation for more than 17 years. His work has included international consultancy for several well known companies in a variety of market sectors. His current role is as product manager for the WITNESS simulation software package where he establishes the requirements for such software and manages the delivery to market. His email address is twaller@lanner.co.uk