MODELLING HUMAN DECISION-MAKING

STEWART ROBINSON

Operational Research and Systems Group
Warwick Business School
University of Warwick
Coventry
CV4 7AL
stewart.robinson@warwick.ac.uk

Abstract

A series of projects have been, and are being, performed, that look at modelling human decision-making in simulations. The focus is on using a simulation to elicit knowledge about human decision-making. Artificial intelligence methods are then used to learn the humans' decision-making strategies. By linking the trained artificial intelligence system with the simulation, it is possible to assess the performance of the decision-maker. Results are presented from the most recent project. The motivation for modelling human decision-making is also discussed. In this work the prime motivation is to understand and improve decision-making, rather than to develop more accurate simulation models.

1. Introduction

Since the mid-1990s the author has been investigating the use of artificial intelligence methods as a means for representing human decision-making in simulations. This paper describes the history of this work and future work that is being undertaken. Starting from an idea generated when attempting to model rail marshalling yards, an artificial example of simulation and expert systems working in collaboration was generated. The ideas were then applied to a real case of maintenance operations at an engine assembly plant. Future work is looking into simulation as a means of knowledge elicitation. The paper briefly describes each of these phases of work and concludes by discussing why it is important to model human decision-making.

2. Forming Ideas

In the mid-1990s the author undertook an ESPRIT funded project looking into the simulation of industrial rail marshalling yards. This work, carried out in collaboration with a Belgian consultancy, aimed to identify the requirements for a rail yard simulator. Previous experience had shown inadequacies in the commercial software available. It was particularly difficult to represent the movement and shunting of individual wagons, backwards and forwards in a yard.

In discussing the nature of rail yard operations another important issue arose. A supervisor is employed to receive incoming trains and direct the splitting up of wagons within the yard. The supervisor is also tasked with selecting wagons from different locations in order to form outgoing trains. This involves complex decision-making, especially if wagons are to be directed and removed from locations in order to minimise the movement and disturbance to the yard. Since the supervisor's knowledge is largely tacit, it is difficult for him/her to express the strategies that are employed. As a result, there is no direct means for representing the decision-making strategy in a simulation model. Indeed, it became apparent that modelling the human decision-making was more problematic than modelling the physical movement of wagons.

3. Proof of Concept

A possible to solution to this issue was the use of expert system, or potentially other artificial intelligence methods. Researchers had previously attempted this with some success (Flitman and Hurrion, 1987; O'Keefe, 1989; Williams, 1996; Lyu and Gunasekaran, 1997). Some of this work had been carried out a number of years earlier and none of it seemed to entail the use of commercial software, which was the focus of the rail yard study.

A small study was set-up to answer two questions:

- Could commercial expert systems be linked to commercial simulation software as a means of modelling human decision-making?
- Could the simulation model be used as a means of knowledge elicitation, by getting a decision-maker to interact with the model?

The aim of the second question was to see if the problem of tacit knowledge could be overcome by creating decision scenarios in a simulation and getting an expert to respond to those scenarios.

A simple simulation, based on a real case in a steel factory, was developed in Witness (figure 1). Lorries arrive at a lorry park requiring loads of between 5 and 20 items. On arrival the lorries are allocated to a loading bay by the bay supervisor, should a suitable one be available. In making this decision the supervisor must take account of the restrictions on the bay capacities. Lorries requiring more than 10 items must be allocated to bay 2 or 3, since bays 1 and 4 only have capacity for up to 10 items. Should a bay not be available then the lorry waits in the park until a suitable bay becomes available. Once a lorry is allocated, it moves to the bay where it is loaded before departing from the system.

Figure 1 *Lorry Loading Bay Example*

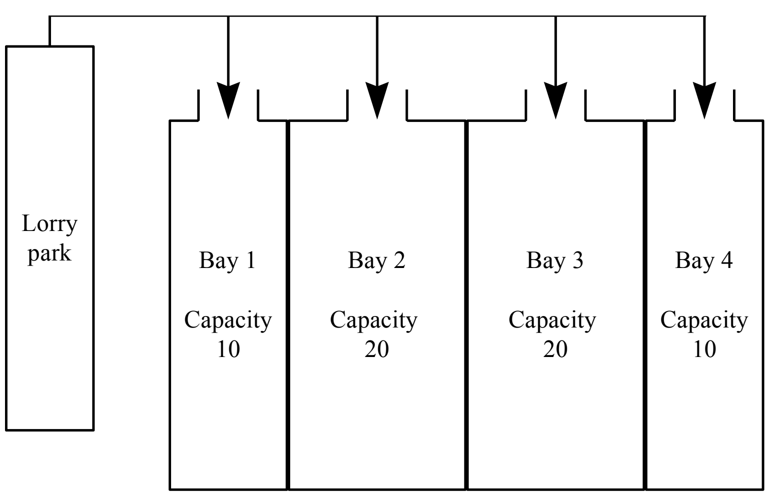

XpertRule was used to develop the expert system that represents the supervisor's allocation decisions. This package was selected for two reasons. First, it adopts a rule induction approach. Second, XpertRule is one of the few expert systems packages available that has a true Windows implementation and is OLE compliant.

Since Witness can only work as an OLE slave, it was necessary to develop a model controller (MC) in Visual Basic (figure 2). The MC initiates the run of the simulation model. At a point where an allocation decision is required, the simulation model automatically stops and waits until the MC returns a decision and continues the run. Once the MC has detected that the model is not running, it extracts data from the model which it passes to the expert system for a decision. The decision is returned to the simulation model via the MC. Some effort was required to ensure that this sequence of events was adhered to. A particular difficulty was encountered in detecting whether the Witness model had stopped running before seeking a decision from XpertRule. If Witness could act as an OLE client it could call XpertRule directly, removing the need for the MC. This would have simplified the linking of the packages significantly.

Figure 2 *Linking Witness to XpertRule*

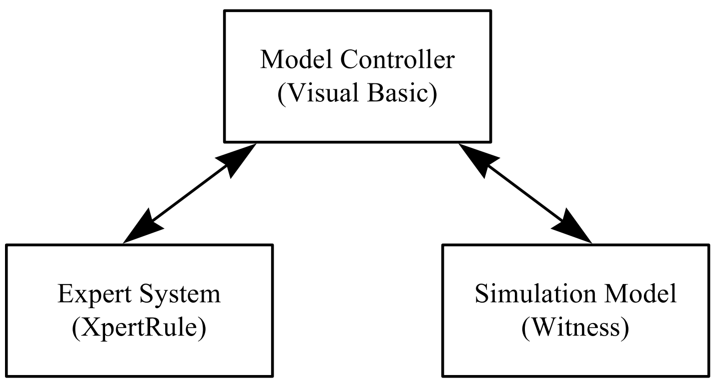

In order to develop the expert system decision tree, the simulation was first used as a knowledge elicitation engine. The simulation was run and at a decision point the user (the author) was prompted for an allocation decision. These decisions were logged in a data file along with variables describing the state of the system. These were then used to train the expert system. The decision tree shown in figure 3 was developed using this approach.

Figure 3 *Decision Tree Induced from Examples*

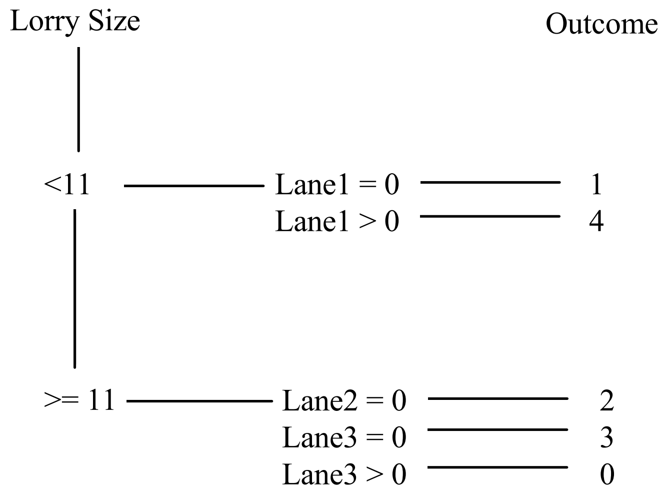

Once the decision tree had been defined, the simulation could be run with the expert system (rather than the decision-maker), in order to determine the effect of the decision-making strategy on the operation of the loading bay. Full details of this work can be found in (Robinson et al, 1998)

4. Modelling Maintenance Decisions in a Manufacturing Plant

The example above showed that commercial software could be linked for the purposes of representing human decision-making and that simulation could be used to good effect as a knowledge elicitation approach. As with previous work, however, this was an artificial example. The question, therefore, arose: could this approach be used in a real and complex case? In 1999 a three-year EPSRC funded collaborative project began, looking at this very issue. The project was a collaboration between Warwick Business School, Aston University, Ford Motor Company and the Lanner Group. The case considered was an engine assembly plant and the decisions taken by supervisors when a machine fails. The work is described briefly below and more fully in (Robinson et al., 2001)

In the engine assembly plant, blocks are placed on a 'platten' 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 were considered. Prior to the research a simulation model of the complete engine assembly facility had already been developed. The model, developed in the WITNESS simulation software, was used to identify bottlenecks and to determine viable operating 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 practice, however, a maintenance supervisor has a number of options beyond repairing the machine immediately:

- Stand-by: an engineer manually processes parts until the end of the shift, when the machine is repaired.
- Stop the line
- Do nothing

The question was, could the simulation that already existed be used to elicit knowledge from the maintenance supervisors on how they made these decisions, and could this information be used to develop an artificial intelligence representation of the decision-maker? The aim was not so much to be able to develop a better simulation model, but to devise a means for identifying and then improving decision-making.

To address this issue, the knowledge based improvement (KBI) methodology was devised which consisted of five stages:

- *Stage 1*: Understanding the decision-making process
- *Stage 2*: Data collection
- *Stage 3*: Determining the experts' decision-making strategies
- *Stage 4*: Determining the consequences of the decision-making strategies
- *Stage 5*: Seeking improvements

These stages are described in detail in (Robinson et al., 2001).

Following a process of knowledge elicitation, up to 63 example decisions were collected from each of the three maintenance supervisors (one for each shift). Knowledge elicitation sessions lasted about one hour. This seemed to be a limit on the time the supervisors had available and on their ability to concentrate on making decisions in the model.

A number of artificial intelligence methods were trained using the examples obtained with varying degrees of success. Table 1 shows the proportion of example decisions that were incorrectly classified by the different methods. A zero score implies a perfect classification. The poor performance of the neural network in unsurprising, since they are known not to perform well with small training sets.

Table 1 *Misclassification Comparison*

	Decision-maker		
	1	2	3
ID3 (XpertRule)	0/63	0/63	0/53
CART (SPSS)	5/63	4/63	4/53
Neural Network (Matlab)	19/63	10/63	24/53
Logistic Regression (SPSS)	0/63	0/63	0/53

The simulation was then run with the ID3 decision tree. The results in figure 4 show the day-to-day throughput resulting from employing the three different decision-making strategies, as well as the results obtained from the decision logic in the original model developed by Ford. This shows some differences in the plant throughput as a result of the different decision-making strategies.

Figure 4 *Throughput under Alternative Decision-Making Strategies*

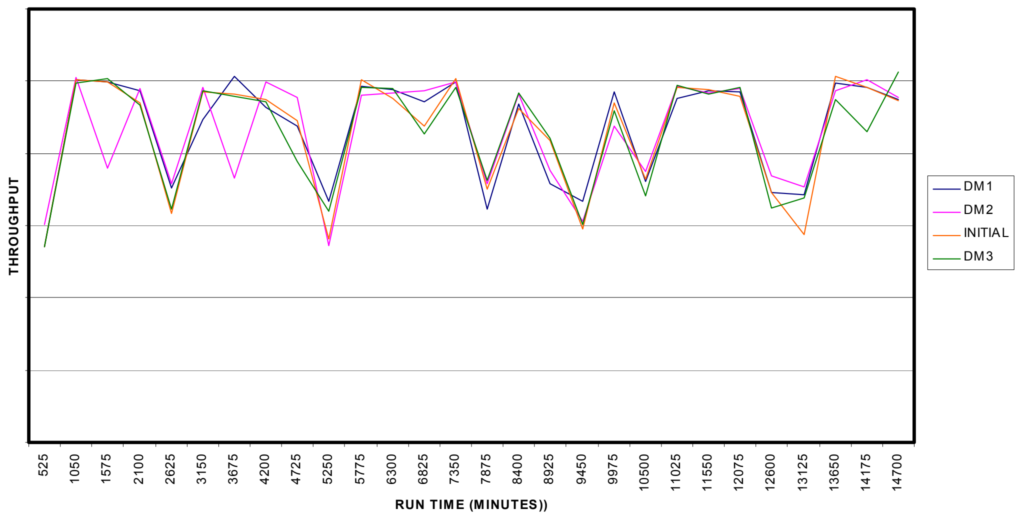

5. Knowledge Elicitation through Simulation

The work on the engine assembly case demonstrated the possibility of using the KBI methodology in a real situation, as well as some of the difficulties in its use. One particular difficulty was in obtaining realistic decisions from the supervisors and in obtaining sufficient example decisions to enable valid artificial intelligence representations to be trained. A further three year project started in October 2002 which will address these specific issue of knowledge elicitation. This work is also funded by the EPSRC with collaboration from Ford, Lanner Group and Aston University.

The specific objectives of the project are:

- To determine alternative mechanisms for eliciting knowledge from decision-makers using a visual interactive simulation
- To compare the alternative methods in terms of their efficiency (speed of data collection)
- To compare the alternative methods in terms of their effectiveness (accuracy of data collection)
- To compare the data collection methods in terms of the ability to train various artificial intelligence methods from the data sets collected

This will involve considering the following issues:

- *Level of visual display*: paper based, none, 2D, 2½D, 3D
- *Interactive interface*: number of decision-making attributes (key data upon which decisions are taken) that are reported to the decision-maker
- *Scenario generation*: use of historic scenarios, adapted historic scenarios to give more extreme examples, random sampling of scenarios, adapted random sampling of scenarios to give more extreme examples
- *Self learning*: learning responses to specific scenarios as the data collection progresses and automatically responding to future iterations of the same scenario

6. Conclusion: Why Model Human-Decision Making?

In conclusion, it is worth discussing the motivation for modelling human-decision making. Is it to enable the development of better models, or to help better understand and possibly improve human decision-making?

Checkland (1981) describes four types of system. Two are of interest here:

- *Designed physical systems*: systems designed by humans with no human interaction in day-to-day operations e.g. an automated warehouse.
- *Human activity systems*: systems of human activity, existing purely as human interaction e.g. political and social systems.

These represent systems at two extremes in terms of human interaction. In simulation modelling of the type considered in this paper (operations system modelling), we are rarely dealing with systems at either extreme, but somewhere between the two. Operations systems typically consist of a designed physical system in which humans interact e.g. a manufacturing line or a bank. Human decision-making is a key aspect of that human interaction. Since human interaction and decision-making are central to operations systems, there is a clear motivation for modelling that interaction.

But are we looking to develop better models by including elements of human decision-making? There is a potential problem with this motivation. Robinson (1994) presents the diagram shown in figure 5. This shows that there are diminishing returns, in terms of accuracy, from increasing the level of complexity in a model. Indeed, it is argued that there comes a point at which added complexity reduces the accuracy of a model because there is insufficient knowledge to support the detail being modelled. Modellers would argue that the optimum point, or best model, is around point x. This is the point at which the model is sufficiently accurate and beyond which there is little gain from additional complexity. The exact location of point x depends upon the purpose of the model, which in turn determines the required level of accuracy.

Figure 5 *Simulation Model Complexity and Accuracy (Robinson, 1994)*

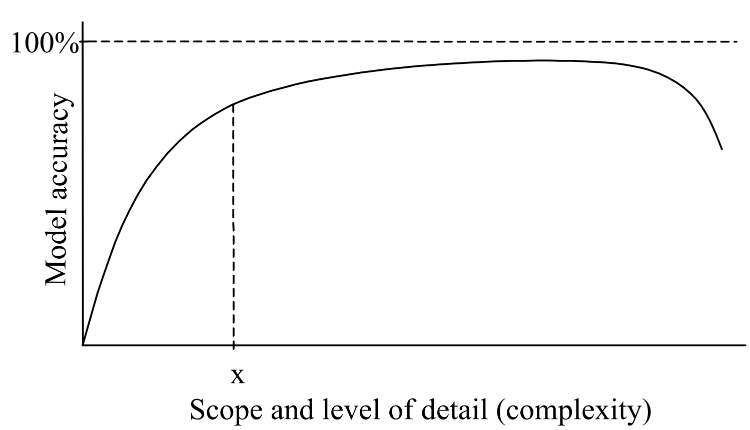

One motivation for modelling human decision-making is to add extra complexity to a model in order to improve its accuracy. The danger of this approach is that it could be trying to climb along the flat part of the curve in figure 5 and so gives little gain. Indeed, it could be argued that although a slightly more accurate model is generated by modelling human decision-making this does not represent a better model, since a large amount of effort is required to obtain only an incremental improvement in accuracy. This argument depends very much on the modelling context and the required level of accuracy. There are cases where incremental gains in accuracy are needed, since a high level of fidelity is required.

Another motivation is to model human decision-making so it is better understood and it can be improved. This should help to improve the performance of the systems in which the humans are interacting. The concentration is no longer on making models more accurate, but on using the models to assess the effects of human interaction and to look for ways of changing the human interaction in order to improve system performance. In this case model accuracy plays a secondary role to generating insight and understanding. This is the motivation behind the knowledge based improvement methodology.

Acknowledgements

The author wishes to acknowledge the support and collaboration of the EPSRC, Ford Motor Company (John Ladbrook), Lanner Group (Tony Waller) and Aston University (Professor John S. Edwards).

References

Checkland, P.B. (1981). *Systems Thinking, Systems Practice*. Wiley, Chichester, UK.

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.

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.

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. (1994). Simulation Projects: Building the Right Conceptual Model. *Industrial Engineering*, 26 (9), pp. 34-36.

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., Alifantis, A., Edwards, J.S., Hurrion, R.D., Ladbrook, J. and Waller, T. (2001). Modelling and Improving Human Decision Making with Simulation. *Proceeding of the 2001 Winter Simulation Conference* ed. B.A. Peters, J.S. Smith, D.J. Medeiros, and M.W. Rohrer. The Society for Computer Simulation, San Diego, CA, pp. 913-920.

Williams, T. (1996). Simulating the Man-in-the-Loop. *OR Insight*, **9** (4), pp. 17-21.