

Using a Simulation Model for Knowledge Elicitation and Knowledge Management

J.S. Edwards^{*a}, T. Alifantis^b, R.D. Hurrion^b, J. Ladbrook^c,
S. Robinson^b, A. Waller^d

^a *Aston Business School, Aston University, Birmingham, B4 7ET, UK*

^b *Warwick Business School, University of Warwick, Coventry, CV4 7AL, UK*

^c *Ford Motor Company, Dunton Engineering Centre, Laindon, Basildon, SS15 6EE, UK*

^d *Lanner Group, The Oaks, Clews Road, Redditch, B98 7ST, UK*

Abstract

The work reported in this paper is part of a project simulating maintenance operations in an automotive engine production facility. The decisions made by the people in charge of these operations form a crucial element of this simulation. Eliciting this knowledge is problematic. One approach is to use the simulation model as part of the knowledge elicitation process.

This paper reports on the experience so far with using of a simulation model to support knowledge management in this way. Issues are discussed regarding the data available, the use of the model, and the elicitation process itself.

Keywords: Knowledge Elicitation, Knowledge Management, Manufacturing, Simulation, WITNESS

1. Introduction

1.1. Maintenance

Maintenance is one of the largest elements of indirect operating costs for manufacturing organisations. A key element of the maintenance process is the scheduling of maintenance tasks for both unplanned maintenance (equipment failure) and planned maintenance (preventative maintenance and tool-change). This is of critical importance to maintaining manufacturing throughput and minimising work-in-progress (WIP). For instance, in a typical Ford automotive engine plant in the order of 2 to 4 engines are lost, or a WIP buffer of 2 to 4 engines is required, for every minute that a machine is not working.

* Corresponding author: e-mail j.s.edwards@aston.ac.uk Phone: +44 (0) 121-359 3611 x5029

Various suggestions have been put forward for methods of automating maintenance scheduling. These include algorithmic [24] and heuristic approaches [8]. However, there has been only limited application of these techniques in practice. In general, maintenance scheduling decisions remain the task of the maintenance supervisor.

Reliance on human decision-makers brings the benefits of experience and flexibility. The maintenance supervisor often has considerable experience with the manufacturing facility and is able to cope with many different circumstances. Inevitably, however, some supervisors will be new to the job and/or the particular facility. In addition, when a new manufacturing facility is set up, by definition the benefits of experience of that facility are not present and may take many years to obtain.

There are more general issues. Human decision-makers are unable to manage large quantities of information in short periods of time and are open to making poor and inconsistent decisions. Indeed, research has shown that there is great variation in the effectiveness of individual 'expert' decision-makers [2]. This has been borne out by the success of expert systems projects intended to bring all decision-makers up to the level of the best available decision-maker, such as Bonissone and Johnson [4]. Perry and Moffat [19] have studied the extent to which even experienced senior decision makers (naval commanders in a simulation of a maritime combat scenario) take different decisions when provided with the same information.

It is important that maintenance supervisors obtain a knowledge of the process that they are managing and of the effectiveness of their decisions as quickly as possible, as well as identifying alternative strategies for making decisions. In this way the costs of maintenance and WIP, and the loss of throughput from machine downtime can be minimised. Knowledge management in general is high on the agenda for many organisations at present [3, 5, 15], but there has been virtually no work specifically directed to managing knowledge about scheduling decisions. Indeed, much of the discussion of knowledge management in the literature takes place in an abstract "world of knowledge", inspired by the work of Nonaka and Takeuchi [17], whose concepts are sometimes difficult to operationalise, especially in manufacturing operations.

1.2. The Project

The work described in this paper has been carried out as part of a project investigating the simulation of unplanned maintenance operations in an automotive engine production facility. From the client's point of view, their main interests for the overall project are twofold:

- improve the maintenance strategies of the facility
- improve the learning processes relating to unplanned maintenance operations

From the modellers' viewpoint, unplanned maintenance operations in an automotive engine production facility present a complex system incorporating a key human element, the maintenance supervisor or group leader.

There are two potential uses of a simulation model for knowledge management in this context: as a tool for helping to elicit knowledge and as a vehicle for supporting organisational learning. Both are part of this ongoing project. This paper concentrates on the former of the two, as it is a necessary pre-condition for the type of organisational learning support envisaged. This will be discussed in more detail later. First we describe the specific situation in the Ford manufacturing facility in more detail.

When unplanned maintenance is required, i.e. a breakdown occurs somewhere on the production line, the maintenance supervisor has to decide what should be done, and (in

most cases) who should do it. A full description of the decisions faced may be found in [20]. The principal options available are:

Repair immediately (RI). Clearly the most obvious course of action to take when a machine breaks down is to repair it, but for various reasons this may be an inappropriate response, an insufficient one, or indeed impossible.

Stand-by (SB) can be considered as an alternative to repair immediately. It does not mean waiting, as might be supposed, but literally standing by the faulty machine; 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 either a complement to, or a substitute for, repair immediately. In this case the maintenance supervisor 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, or complete depletion, of work-in-progress in a section of the line. The line can be stopped either before (in front of) the machine that has developed a fault, or after it, or both.

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

If the decision includes RI or SB, then the maintenance supervisor also needs to decide whom to allocate to the job. This will depend on who is free, as well as the type of fault.

A Visual Interactive Simulation (VIS) model of the manufacturing facility had already been developed using the WITNESS software [12]. The advantages of VIS in manufacturing based problems, such as the development of alternative maintenance strategies, are well documented; see for example [10, 13]. A particular benefit anticipated in this project was that the VIS model could be used as a vehicle for experimentation, both with and without the human decision-maker. However, the existing model in effect assumed that the decision would always be RI, random sampling being used to determine the skill level of employee needed to effect the repair and to generate the repair time. In order to help in the development of alternative maintenance strategies, a better model of the human element in this decision needed to be constructed.

Two ways forward presented themselves:

- use an algorithmic approach as part of the simulation
- link some form of knowledge-based system (KBS) to the simulation model

The choice made by the project team was the second option, on the basis that using a KBS for this part of the overall model would give more flexibility in knowledge representation and also improve the ease of construction. It was recognised that technical issues might arise from having to link two packages together, but the literature does contain examples of the linking of a KBS to a simulation model for training purposes. Examples can be found in [7, 9, 14, 18, 25].

1.3. "The Difference"

A further point of research interest to the modelling team was that the "human element" in this case was not the same as in most situations where a KBS is proposed. Most KBSs are developed with the intention of making the best possible decisions. What was needed here was to represent the human decision-maker as closely as possible,

which is a slightly different requirement. Although finding ways to improve the human performance was one objective of the project as a whole, it was not the immediate objective of the KBS.

1.4. Possible Technologies

Various technologies could be used to construct the KBS. These include:

- Neural networks
- Rule-based systems, either produced directly by machine learning (“rule induction”) or devised by humans
- Data mining

All of these require some kind of knowledge elicitation process. Neural networks, rule induction and data mining need, at minimum, a set of examples; values of appropriate variables, plus the decisions made. Typically this set is then sub-divided into a training set, whose examples are modelled by the software package, and a testing set which is used to validate the model produced. Modelling the relevant knowledge by a human team requires a more flexible process, tailored to the domain and to the human decision-makers concerned. This may include interviews, observation of the decision-maker in the real situation, protocol analysis (working through simulated cases) and/or more formal methods such as the use of repertory grids. Kidd [11] gives a good introduction to the various methods. A point that should not be overlooked is that even when an approach based on examples is used, it is still necessary to have the human decision-maker involved in order to determine which variables might be relevant to the decision concerned.

The first phase of work concentrated on the rule-based approach, with the collection of examples of real decisions to be used as an input both to a machine learning/rule induction process, and for human knowledge engineering. The rule-based approach was preferred to neural networks or data mining, as such an approach makes it easier to construct explanations for how the KBS is reaching its decisions. The downside is that the rule-base may become very large and unwieldy. Recognising this problem, Moffat [16] uses an approach based on Bayesian inference for representing command and control in military campaigns. In the present context, however, it was not expected that the rule-base would be too large.

This work uses the XpertRule package [1], which allows either a rule induction or a human knowledge engineering approach to be taken. XpertRule also had the advantage that the technical problems of linking it to the WITNESS simulation had been overcome in previous work by Robinson et al [22]. The next section of the paper reports on the knowledge elicitation process, and the role of the simulation model in it.. The third section discusses the issues that have emerged so far, under three headings: data, the knowledge elicitation process, and knowledge management. The project outcomes are then discussed as well as the connections between simulation and knowledge management. The paper concludes with some views on the future directions of the research.

Although the example presented in this paper is specific to maintenance operations in the manufacturing sector, the concepts could be applied in many contexts. Many processes that are the subject of simulation involve human interaction and decision-making. For such processes, simulation could play a role in eliciting and managing knowledge. This role is outlined in section 5.

2. Knowledge Elicitation

2.1. Starting point

As mentioned in the previous section, one approach to knowledge elicitation is that of protocol analysis, using either real or simulated cases. By using the simulation model that was already in existence, the knowledge elicitation process could combine the validation benefits of using real cases and the speed and control benefits of using simulated ones. The starting point for the knowledge elicitation process was to use the “trace” – an extensive log of past actual situations regarding unplanned maintenance on the line. The production facility has an automated system for collecting data on failures that occur (times and fault types), and so a realistic trace could easily be generated.

In the real facility, the maintenance supervisor is alerted to the need for unplanned maintenance (and so to make a decision) by a message received on a pager. An interface to the simulation model was developed so that the simulation model could be stopped at the point where the maintenance supervisor’s decision was required, and the human decision-maker presented with the information that would be in the pager message. The simulation also provided other appropriate information that would be available from the production control systems and the supervisor’s presence on the line. Fuller details of its development may be found in [21]. This interface and model would then be used for protocol analysis sessions with the actual decision-makers (the maintenance supervisors). This is shown in Fig. 1.

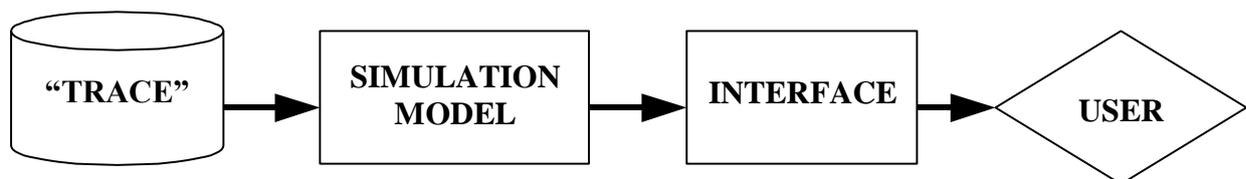


Fig. 1. Use of the simulation model in knowledge elicitation (early approach)

2.2. “One Rule”

In early discussions with the staff of the facility, it was claimed that all the knowledge possessed by the supervisors could be represented in one rule, as follows:

If the estimated repair time is less than n minutes, repair immediately
If not, stand by

There was, however, some debate about the value of n . Values of 10 or 20 minutes were mentioned. Indeed, later it emerged that senior management had recommended the maintenance supervisors to use such a decision rule, although as will be discussed, it is not clear that they actually do use it.

3. Emerging Issues

The various issues that emerged during the process of knowledge elicitation may be grouped under three headings:

- Data
- The knowledge elicitation process itself
- Knowledge management

3.1. Data issues

Two issues arose concerning the data presented to the decision-maker. The first was obtaining a value for one of the key variables in the decision. As mentioned above, the need for a decision in the real facility is signalled by a message on a pager indicating the machine that has broken down and the nature of the fault. Observation of the subsequent decision-making process revealed that the first step in the real decision is to look at the condition of the machine that has broken down in order to estimate the repair time. Although it might be theoretically possible to simulate this given unlimited resources, for example by a system that displayed a video sequence of the machine fault in question, this is impracticable with the resources available for the project. It was therefore necessary to generate and report an estimated repair time for the fault, to be used as a surrogate for the maintenance supervisor's own judgement in this matter.

The second issue concerned the variability in the data – or rather the lack of it. Although built up over a reasonably long period of time, the trace contained only a limited selection of the possible faults that might occur, especially as it was only possible to present a subset of the data during the protocol sessions (see below). As a result, in the first knowledge elicitation session, the actual decision made in every case was RI (repair immediately).

The solution to this problem was combined with the solution for the next issue, and will be discussed there.

3.2. Knowledge elicitation process issue

Two issues arose under this heading. The main issue centred on the effort involved in interacting with the model. One of the advantages of using a simulation model is its ability to run faster than real time – much faster, in this example. In this case, a typical protocol analysis session took around 2 hours to collect 100 decisions from the maintenance supervisor. It was commonly the case that the supervisor seemed tired and/or fed up by the end of the session. This was not surprising when one considers that a normal 8-hour shift might yield 5 or 6 decisions of this type, on average. The interaction with the simulation model was thus far more intense than the normal conditions under which the decisions are made, with consequently a possible (adverse) effect on decision quality.

In order to carry out effective protocol analysis sessions using the model, addressing both this problem and the limited variability in the data, two changes were therefore made. The first was to use a synthesised trace rather than part of the real one. Effectively a biased quota sampling approach was taken to the various possible faults, so that the less common (and generally more drastic) faults were very much over-represented in the synthesised trace, and thus the full range of different possible decision situations could be presented to the decision-maker. This does not affect the later stages of building the

knowledge-based system, since the most common form of rule-based systems, as used in this project, are essentially deterministic in operation. Only one example of each type of situation is therefore needed.

The second change was to run the simulation model in advance of the knowledge elicitation session, the interface becoming a separate front-end program with the decision situations stored and ready inside it, as shown in Fig. 2. Although this made the situation slightly less realistic, by removing the uncertainty of when the next decision would be needed, the greater speed with which decisions were presented eased the strain on the decision-makers during the knowledge elicitation session. This obviously meant that the decision-makers could not reference the simulation's visual display. It was apparent, however, that they had not done so in the earlier protocol analysis session. The number of decisions was also restricted to around 60, again to ease the pressure on the decision-maker.

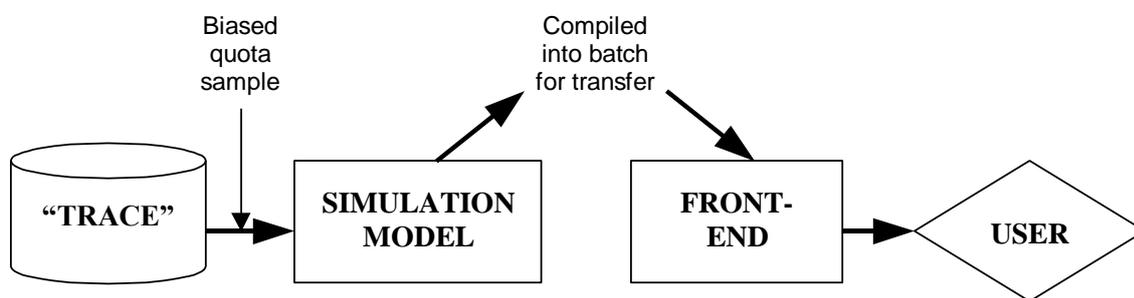


Fig. 2. Use of the simulation model in knowledge elicitation (later approach)

The second issue relating to the knowledge elicitation process concerned the apparent correctness of the model, as perceived by the maintenance supervisors. Although it is generally accepted that, within certain bounds, the more detailed a simulation model is, the more accurate it should be [20], this does not necessarily result in its being perceived as “more correct”. In this case, some of the maintenance supervisors, presented with certain examples based on the trace, pointed out that the machine concerned no longer existed on the line. They were therefore reluctant to take seriously any decision about repairing it. *A priori*, one would expect that as detailed a simulation as could be achieved would be needed to ensure that the simulated decision-making environment seemed as “real” as possible. However, here we found that too much detail in the simulation may negatively influence the perceived realism of the situation, particularly where that detail is incorrect. Interestingly, while some of the human decision-makers were bothered by this aspect of “non-existent machines”, others did not notice it at all. A similar phenomenon has been noted by Silver [23] regarding how users perceive the capabilities of a decision support system. Build in too broad a range of capabilities, and some users perceive the system as actually being very restrictive, because of the number of choices that have to be made in order to use it for anything at all. It is also possible that decision-maker fatigue, as mentioned above, had some influence on the perception of the model’s “correctness”.

3.3. Knowledge management issues

Three issues arose under this heading, although they are not completely independent of each other, or of the issues discussed previously.

First, and most importantly, it was apparent that different personnel approached the unplanned maintenance decision in different ways, assuming that the knowledge elicitation sessions that we carried out were valid. This was something that we had anticipated would occur, and is one of the reasons behind the “accurate representation” goal referred to in section 1.3. However, it is relatively difficult to observe in the real system because by definition there is only one person doing the maintenance supervisor’s job at any given moment. We found that some decision-makers did indeed appear to look first at the repair time, but that a more common approach was to consider the machine type first, then base the decision either on repair time or fault type, depending on the type of machine. The detail of the specific decisions made varied substantially between different decision-makers.

For example, the same set of 63, mainly “difficult” or “interesting” decisions was presented to four decision-makers (‘one’ of which was actually two people working together). Considering just the basic decision of RI (repair immediately) or not, only 19 (30%) of the 63 yielded the same decision in all four knowledge elicitation sessions. In addition, there were a further 9 breakdowns in which one of the decision-makers identified a machine that was no longer on the line and so declined to make a decision (see section 3.2), but the other three made a decision (the same decision in each case). In total there were therefore 28 (44.4%) cases where there was agreement, and 35 (55.6%) where there was substantive disagreement. We would like to stress the point that these percentages would be very different from those made in a sequence of real decisions, since more than 90% of decisions in the trace are very minor faults where the best decision is unquestionably to repair immediately. Table 1 shows an overall analysis of the decisions made in the 63 examples.

Table 1
Comparison of decisions whether or not to Repair Immediately (RI) made by four decision-makers

	All four made decision	Only three made decision; one declined	Total
Total number of examples	53	10	63
Unanimous agreement (4-0 or 3-0)	19	9	28 (44%)
Majority verdict (3-1 or 2-1)	19	1	20 (32%)
Split decision (2-2)	15		15 (24%)

The second issue, which surprised almost everyone involved with the project, is that the framing of the situation around a single decision-maker turned out to be not quite the case in reality. There are circumstances where the maintenance supervisor will consult a colleague or a superior before making the decision. This is especially likely where the supervisor is relatively new to the position. This became apparent during the data collection sessions, and was the reason why one of the sessions elicited knowledge from two people working as a pair.

A third issue is the possibility that the maintenance supervisors make inconsistent decisions. As mentioned in section 1, this is one of the identified weaknesses of human decision-makers. In the knowledge elicitation situation used here, when a given decision-maker is faced with exactly the same set of attribute values more than once, we might

expect to see that they would on occasions make different decisions. In the limited number of knowledge elicitation sessions held so far, we have not been able to demonstrate this unequivocally. Two sessions involved a set including some repeated examples, and these did lead to a few inconsistent decisions (four in total). However, these were all for examples where either the machine no longer existed on the line, or the fault type was (technically) not really a fault at all, e.g. machine door left open. The action in the latter case is of course to close the door, not to repair anything. It is also possible that this may again have been related to the fatigue factor, although these were by no means the last decisions made in that particular knowledge elicitation session.

That being said, the rules which best fit the decisions observed in the knowledge elicitation sessions nevertheless do not always make sense. One instance of this is that a sub-set of rules is so convoluted as to imply that, at the very least, there is inconsistency between the decisions made by one decision-maker in similar (but not identical) examples. The fact that this occurs is not necessarily a criticism of the knowledge elicitation approach used. It may indeed be that the decisions made in real life are inconsistent. This would tie in with our objective to represent the human decision-maker accurately, rather than to make the best possible decision. However, at this point our overall model of the process is not sufficiently robust to be able to test whether this is a genuine feature of the decision-making or not. It does raise problems for constructing decision making rules, especially using rule induction. With deterministic rules, a contradictory set of examples means that it is not possible to induce a rule set that completely describes those examples, although version 4 of XpertRule does now allow a set of rules to be created by including the contradiction as an outcome in itself.

An alternative instance of potential inconsistency is that even when contradictory examples were removed from the training sets, and a relatively simple set of rules clearly fitted the data, they sometimes still made little sense in practice. The most common problem would be a rule such as “if estimated repair time is less than m seconds, then stand by; if not, then repair immediately”. None of those involved in the project can think of a real world justification for this; stand by should become more likely the *longer* that a machine is expected to be out of action, and most rules reflect this.

3.4. The Fundamental Snag!

Taking all these issues together, we believe that we have now developed a useful knowledge elicitation process. As the description above shows, the process of reaching this stage has not been without its difficulties. At the heart of these difficulties is an interesting contradiction. The line being studied is a fairly new line in a modern facility, and so few major problems have occurred yet. From the management point of view, and indeed from a theoretical point of view, it makes sense to study the system and share/transfer knowledge *before* any major problems occur. However, because it is a new line, no-one necessarily *has* the knowledge that is required in all cases.

4. Project Outcomes

In spite of these difficulties, there have been several positive outcomes thus far. We have demonstrated that the “one rule” stated earlier is certainly not the whole story as far as modelling the human decision-making is concerned. Indeed, even the idea of the

single decision-maker is not entirely true. Machine type and fault type are at least as important as repair time in influencing the decisions made.

We have also demonstrated that different decision-makers not only appear to make decisions in fundamentally different ways, but as a result come to different conclusions. This confirms the findings of Perry and Moffat [19] in the far less clear-cut decision-making situation of maritime combat. Although the differences we have found only occur in a small minority of decisions, these are typically the decisions that potentially cause most disruption to production on the line.

An additional benefit is that a new maintenance supervisor has used the model developed for this project as a tool to help him learn his new job in a relatively safe environment. Thus in this particular way, the project has achieved part of the second objective, to improve the learning processes relating to unplanned maintenance operations.

5. Simulation and Knowledge Management

In the learning process described in the previous section, the simulation model played a central role, as it has done in the knowledge elicitation process. We have therefore seen examples of how a simulation model may be used to support various knowledge management activities. To look at this more systematically, we use Fig. 3, which shows a knowledge process model we have proposed elsewhere [6].

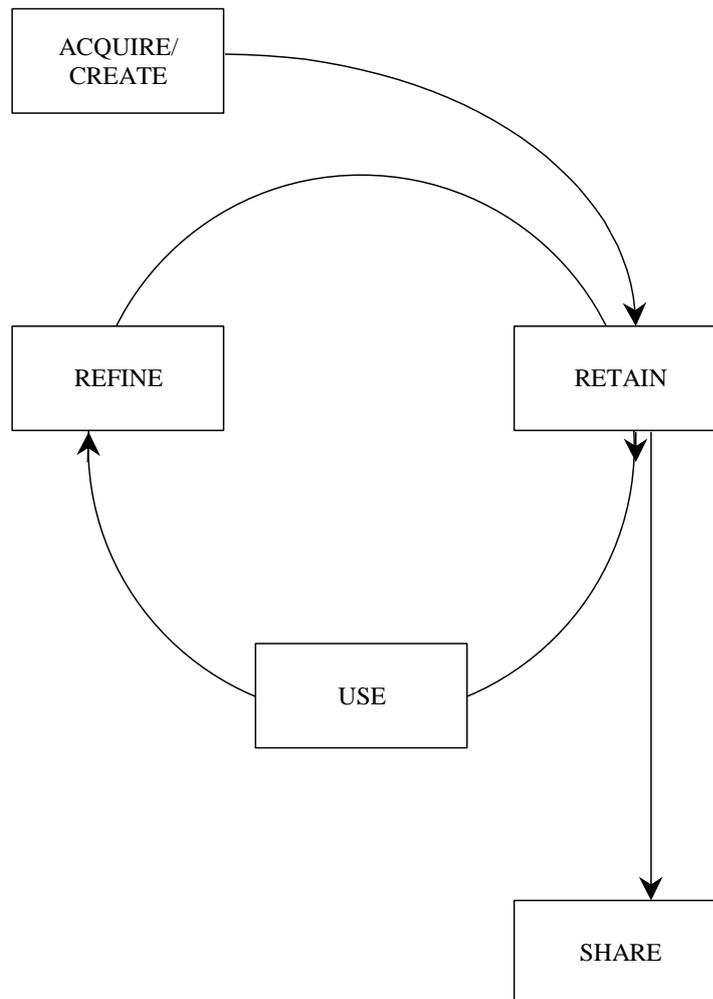


Fig. 3. A model of the knowledge management process focused on the knowledge itself. Source: [6]

From this we can identify three actual or potential uses of the simulation model and/or the knowledge-based system for knowledge management. We have already seen that a new supervisor used the simulation model as part of his training; this was Sharing the knowledge of others, but could also include Acquiring knowledge, if the approach that he develops turns out to be different from that of anyone else.

Future uses may include analysing the rule sets and decision trees for different supervisors (which we have done) and showing them to each other (which we have not yet done). This will again enable knowledge Sharing, and potentially knowledge Refinement. In addition, allowing the supervisors (or indeed their managers) to experiment on the system will support both the Acquisition of new knowledge and the Use of existing knowledge. For example, it has been suggested that Stand By operation (SB) should not be used. What effect might this have on the production line's output?

Overall, there is potential for the simulation model, perhaps with the addition of the knowledge-based system, to help with acquiring and surfacing tacit knowledge, and with sharing, refining and using explicit knowledge.

6. The Future

The work carried out to date has explored the use of simulation as a means for eliciting data from decision-makers and the representation of that data using rule based expert systems. The benefits and difficulties encountered in using this approach have been described. The role of the simulated knowledge elicitation approach in knowledge management has been briefly explored.

There are various directions that could be considered in future work related to this project and more generally to simulation and knowledge elicitation/management. Three ideas are discussed below.

Other approaches for modelling the examples that we already have could be adopted, for instance data mining techniques or logistic regression. However, the large proportion of decisions that are “repair immediately” does present us with some technical problems, whatever approach is followed.

A second area to pursue is the “learning” aspect of the work, looking more closely at the different ways in which different people make decision. This could consider not only differences between decision-makers, but changes to a decision-maker’s strategies over time.

A third area is to investigate alternative ways of presenting simulated examples to decision-makers in order to achieve the most realistic example set. This should consider the representational aspects (detail of information given and visual display) as well as the nature of the examples that are presented to a decision-maker.

Overall, this appears to be a very fertile area for continued research.

Acknowledgements

This research has been supported by the Engineering and Physical Sciences Research Council’s Innovative Manufacturing Initiative under grant reference GR/M72876, by Ford Motor Company, and by Lanner Group.

We are also grateful for the comments of the anonymous referees.

References

- [1] Attar Software, XpertRule User Guide (Attar Software, Newlands Road, Leigh, UK, 1997).
- [2] N.R. Augustine, Distribution of Expertise, *Defense Systems Management* (1979) 50.
- [3] I. Becerra-Fernandez and R. Sabherwal, Organizational knowledge management: A contingency perspective, *Journal of Management Information Systems* 18 (2001) 23-55.
- [4] P.P. Bonissone and H.E. Johnson, Jr., Expert Systems for Diesel Electric Locomotive Repair, *Human Systems Management* 4 (1984) 255-262.
- [5] E. Coakes, D. Willis, and S. Clarke, eds. *Knowledge Management in the SocioTechnical World: The Graffiti Continues*. Computer Supported Cooperative Work, ed. D. Diaper and C. Sanger. 2002, Springer Verlag: London.
- [6] J.S. Edwards, Artificial Intelligence and Knowledge Management: How Much Difference Can It Really Make?, in: *Proceedings of KMAC2000*. Knowledge

- Management Beyond The Hype: Looking Towards The New Millennium (Operational Research Society, Aston University, Birmingham, UK, 2000) 136-147.
- [7] A.M. Flitman and R.D. Hurrion, Linking Discrete-Event Simulation Models with Expert Systems, *J. Opl Res. Soc.* 38 (1987) 723-734.
 - [8] M. Gopalakrishnan, S.L. Ahire, and D.M. Miller, Maximizing the Effectiveness of a Preventive Maintenance System: An Adaptive Modeling Approach, *Mgmt Sci.* 43 (1997) 827-840.
 - [9] T.J. Grant, Report on the American-Paper-Institute Recovery Boiler Program, *Pulp and Paper-Canada* 91 (1990) 66-67.
 - [10] R.D. Hurrion, The Design, Use and Required Facilities of an Interactive Visual Computer Simulation Language to Explore Production Planning Problems, University of London 1976.
 - [11] A.L. Kidd, ed. Knowledge Acquisition for Expert Systems: A Practical Handbook. 1987, Plenum: New York.
 - [12] Lanner Group, WITNESS 2000 User Manual (Lanner Group, The Oaks, Clews Road, Redditch, UK, 2000).
 - [13] A.M. Law and W.D. Kelton, Simulation Modeling and Analysis, 3rd ed (McGraw-Hill, New York, 2000).
 - [14] J. Lyu and A. Gunasekaran, An Intelligent Simulation Model to Evaluate Scheduling Strategies in a Steel Company, *International Journal of Systems Science* 28 (1997) 611-616.
 - [15] K. Mertins, P. Heisig, and J. Vorbeck, eds. Knowledge Management: Best Practices in Europe. 2001, Springer-Verlag: Berlin.
 - [16] J. Moffat, Representing Command and Control in Simulation Models of Conflict', *Journal of the Operational Research Society* 51 (2000) 431-439.
 - [17] I. Nonaka and H. Takeuchi, The knowledge-creating company, how Japanese companies create the dynamics of innovation (Oxford University Press, New York and Oxford, 1995).
 - [18] R.M. O'Keefe, The Role of Artificial Intelligence in Discrete-Event Simulation, in: L.E. Widman, K.A. Loparo, and N.R. Neilsen, eds, *Artificial Intelligence, Simulation and Modeling*, (Wiley, New York, 1989) 359-379.
 - [19] W.L. Perry and J. Moffat, Developing Models of Decision Making, *Journal of the Operational Research Society* 48 (1997) 457-470.
 - [20] S. Robinson, Successful Simulation: a Practical Approach to Simulation Projects (McGraw-Hill, Maidenhead, England, 1994).
 - [21] S. Robinson, A. Alifantis, J.S. Edwards, R.D. Hurrion, J. Ladbrook, and T. Waller, Modelling and Improving Human Decision Making with Simulation, in: *Proceedings of The 2001 Winter Simulation Conference* (Institute of Electrical and Electronic Engineers, Piscataway, New Jersey, 2001) 913-920.
 - [22] S. Robinson, J.S. Edwards, and Y. Wu, An Expert Systems Approach to Simulating the Human Decision Maker, in: *Proceedings of Winter Simulation Conference 98* (The Society for Computer Simulation, San Diego, CA, 1998) 1541-1545.
 - [23] M.S. Silver, Systems that Support Decision Makers: Description and Analysis (John Wiley, Chichester, UK, 1991).
 - [24] R.W. Taylor, A Linear Programming Model to Manage the Maintenance Backlog, *Omega* 24 (1996) 217-227.
 - [25] T. Williams, Simulating the Man-in-the-Loop, *OR Insight* 9 (1996) 17-21.

