What Should the Agent Know? The Challenge of Capturing Human Knowledge

(Short Paper)

Emma Norling
Centre for Policy Modelling
Manchester Metropolitan University
Aytoun Building, Aytoun Street
Manchester, M1 3GH, United Kingdom
norling@acm.org

ABSTRACT

Reports of applications that include agent-based models of human behaviour tend to focus on the applications themselves and the success of the modelling exercise. They give little (if any) information on the process used to construct the models. Those who attempt to construct models of human behaviour quickly realise that this is a non-trivial task. This paper presents a methodological approach to eliciting knowledge for BDI-based models of human behaviour.

Categories and Subject Descriptors

D.2.2 [**Software Engineering**]: Design Tools and Techniques

General Terms

Design

Keywords

BDI, human modelling, knowledge elicitation

1. INTRODUCTION

Agent-based models of human behaviour have been used in a wide range of applications, from interactive games and story-telling to education, training and operations analysis. This paper is concerned with the construction of models that attempt to generate behaviour based on plausible models of human reasoning (as opposed, for example, to models which simply aim to create believable characters).

Capturing the underlying reasoning process is important in applications which are designed to explore different scenarios, and particularly unexpected scenarios. While it might be possible to know what behaviour would be expected in a typical situation, what is of more interest is what behaviour might arise in an unexpected situation. (In fact, it can even be the case that this type of model might illustrate *unexpected* behaviour in what is thought to be a typical situation, due to unexpected cognitive demands.) A major challenge in building these types of model is to capture the reasoning process. This does not require a full model of human cognition (indeed, a "full" model does not exist), but does require a model that does considerably more than simply react to stimuli.

Numerous agent-based models of this type have been constructed, particularly in the military domain, going back to TacAir-Soar [4] and SWARMM [5] in the mid nineties. However the reports of such applications focus on the applications themselves and the successes of the modelling exercises, giving little (if any) information on how the models have been constructed. Those which use BDI-based agents often discuss the natural mapping between the language used by experts and the core concepts of the paradigm (for example [7]), but it takes more than a casual conversation to gather the required information from the expert, and this process is not discussed.

The methodology presented here draws upon techniques from cognitive task analysis (CTA) to elicit knowledge for the development of BDI-based models of human behaviour. It could be adapted for use with other agent modelling paradigms, but is particularly suited to BDI due to that paradigm's folk psychological roots. The illustrative examples are taken from a project in which models of Quake 2 players were developed, but the methodology is intended to be applicable to the development of any BDI-based models of human behaviour.

2. KNOWLEDGE ELICITATION

There are a wide range of knowledge elicitation methodologies that have been developed for other purposes than agent-based modelling – for the development of training programs, to aid task and/or interface (re)design, and also for building models using other technologies. Cognitive task analysis is a family of methodologies that is particularly suited to knowledge elicitation for BDI-based modelling of human behaviour as the techniques focus on the cognitive aspects of the task – that is, *why* the subject performs certain actions – rather than trying to define the task in terms of procedures to be followed.

The methodology that is presented in this paper draws largely upon two forms of cognitive task analysis: applied cognitive task analysis (ACTA) [6] and the critical decision method (CDM) [3, 2].

Cite as: What Should the Agent Know? The Challenge of Capturing Human Knowledge (Short Paper), Emma Norling, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008,Estoril,Portugal,pp.1225-1228.

3. EXAMPLE MODELLING EXERCISE

The application of these techniques is illustrated here with examples taken from a project which involved the development of BDI-based models of expert players of the Quake 2 deathmatch game [8]. This particular application was chosen not because of a specific interest in computer gaming, but because the game environment provided a rich simulated world in which models could be evaluated, and subject matter experts were easily accessible. The techniques used are more broadly applicable, but the examples are limited to this one exercise for coherence.

The world of Quake 2 is specified by a three-dimensional map, through which the human players navigate with a graphical first person view, using mouse and keyboard inputs. The map can vary enormously from one game to another, from narrow hallways to wide open spaces, and can include hazards, secret hiding spots, bright lights, dark corners, etc. Scattered throughout the map are various items – weapons, armour, health, and others – each of which is periodically regenerated in the same place in the map, and each of which has its own particular set of characteristics. For example, a shotgun has very different range, firing speed, ammunition, etc to a rocket launcher. Players respawn at a number of fixed positions in this map, reappearing at a random one of these each time they die. Each time they are respawned, they have only the most basic weapon, full health, and no other items.

Within this world, the players must equip themselves, seek out other players, and kill them, all while avoiding being killed. (Being killed does not directly affect your score, but puts you at a distinct disadvantage due to lack of equipment.) There are a wide range of strategies used by different players, with none appearing to be significantly better than any of the others. That is, players with completely different styles can be very closely matched in terms of score. In this project, three players with considerably different playing styles were interviewed, and models constructed of two of them.

4. A STRUCTURED APPROACH

The methodology presented here uses elements from both ACTA and CDM. Specifically, it involves:

1. Preparation, preferably through direct observation as described below.
2. Development of a task diagram, as in ACTA.
3. Finally, an iterative process of:
   (a) Expansion of each of the elements in the task diagram through a combination of
      i. probes as per the knowledge audit process of ACTA,
      ii. presentation of hypothetical situations, and
      iii. CDM sweeps
   (b) Analysing the data and designing the models.

The number of iterations needed in the final step will depend upon a number of factors, including the care taken in planning the interview session(s), how forthcoming the subject is, and the complexity of the task being performed. *Knowing* when to stop can be tricky, as is discussed in Section 4.6.

4.1 Direct Observation as a Preliminary Tool

Direct observation can be a useful first step in the knowledge elicitation regardless of the level of familiarity of the interviewer with the task of interest. As well as familiarising the interviewer with the task (or, if the interviewer is already familiar with the task, alerting him/her to alternative means of performing the task), direct observation can serve two further purposes. Firstly it is a means of assessing the suitability of the subject(s): in the case of the project here, direct observation revealed that one of the three volunteers did not have the expertise that was required in later stages of the project. Secondly, direct observation can provide reference points for discussion in later interviews.

In the case of the example discussed here, it was possible for the three volunteers to play a tournament against each other. In other domains it would be necessary to observe each (potential) subject separately. In some cases (for example in some military scenarios) the logistics and/or cost of arranging for direct observation might be prohibitive. In such situations it is possible to proceed directly with the knowledge elicitation described below, but the interviewer must ensure that he/she has prepared adequately using other sources such as instruction manuals or similar literature.

4.2 Development of a Task Diagram

The first interview in each case was designed to develop the task diagram for each player, and consisted of questions such as those in Table 1. The data was used to construct a task diagram such as that shown in Figure 1. This was presented to the subject for approval and then each stage was considered in detail in later interviews. For the two subjects, the task diagrams that were generated differed only slightly in the timings.

Table 1: Sample first interview probes

When you play the game, do you perceive any distinct phases?
What are you main goals in each of these phases?
What are the relative priorities of these main goals?
Say you enter a new game, where you don't know the world map, and you may know some, but not all, of the players. What are the first things that you do?
Do you make an effort to get to know the style of the other players? How do you use that knowledge?

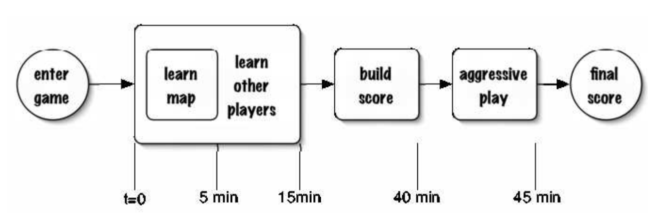

Figure 1: The task diagram for one of the subjects

4.3 Expanding the Task Diagram

The second stage of interviews probed more deeply into the phases that were identified in the task diagrams, using the three strategies identified previously: 1) probes about expertise, as in a knowledge audit, 2) probes using hypothetical situations to elicit details about strategies, and 3) asking the subjects to give examples of situations that were unusual or where they felt they performed particularly well, as in CDM. Due to the fast-moving nature of the game, the incidents presented by the subjects in this third case were of limited duration. Nevertheless the exploration of a number of these small incidents revealed much about the subjects' styles and strategies.

Whereas the first stage of interviews focused on the subjects' goals when playing the game, the second stage turned to the strategies that they used, and the way in which they characterised the world. Some examples of the types of probes are presented in Table 2.

Table 2: Sample probes from later interviews

You say the first stage of the game is when you don't know the map. When do you consider that you do know the map? Do you explore every nook and cranny?
What makes a good sniping spot?
If you'd just respawned and you could hear but not see a fight nearby, what would you do?
How important are the sounds in the game to you? What sorts of things do you listen for?
What sort of things most clearly differentiate novice players from expert players?
Say you'd identified a particular opponent as being a better player than you. Would you make an attempt to actively avoid him/her?
When you were talking about exploring the map, you said you explore cautiously. What do you mean by cautiously?

The interviews in this stage were open-ended. Before each interview, a list of questions were prepared, based upon the analysis of data from the previous interview. But it was also important to be flexible during the interviews, picking up on key points with the subjects' responses and following up on these in order to fill in the details. It was not critical if some of these details were missed during an interview, as the post-interview data analysis should identify them, but this would then require a further iteration in this stage of the knowledge elicitation. It was essential when framing the questions to remember that the aim was to uncover the *reasoning* used, rather than understand the task in a procedural sense. The task was complex enough that it would be impossible to cover every hypothetical scenario; instead, the aim was to understand the way they reasoned about their world, and hopefully by capturing this reasoning be able to generate realistic behaviours in unexplored situations.

4.4 Data Analysis and Model Design

As the aim was to develop BDI-based models human behaviour, it was necessary to identify the goals, plans and beliefs (or more accurately, the things about which the subject held beliefs) in the interview data. In the models were being implemented using a different modelling paradigm, it might be necessary to search for different elements in the text. However it should be stressed that the methodology described here works *because* of BDI's folk psychological roots: people naturally discuss their reasoning in these terms, and probes regarding goals, plans and beliefs are easily addressed. It would not be advisable to use this type of technique to try to obtain knowledge for a modelling language that required *sub*conscious elements of reasoning, as subject matter experts generally do *not* have ready access to this type of knowledge. As such, they are not reliable sources for this type of knowledge, and it should be obtained from other sources.

Consider the fragment of interview below:

> **Subject:** ...hearing doors open is a good one, or hearing lifts activate. Like for instance if you're up the top on that level you can hear the lift coming before it gets there, so...
> **Interviewer:** Right, so you're ready to shoot at them?
> **Subject:** As soon as you hear the lift activate, you can start lobbing grenades into it.
> ...
> **Subject:** One of the things that it's common to do on this level... and... a lot of other levels, is to run across the lift, activate it so it comes up... umm... people will start... lobbing grenades in. Then if you're quick enough you can actually run around the corner and umm... beat the lift to the top. And... shoot them as they're lobbing grenades into the lift.

This fragment gives information for the basis of two plans: one to attack players in lifts, and another to trigger the lift and sneak up behind the player who thinks there is someone *in* the lift. It also highlights the way in which the subject reasons about various objects, as shown in Figure 2. The agent must have a relationship between a particular type of sound and the location of lifts (elevators), and the lifts have some attributes.

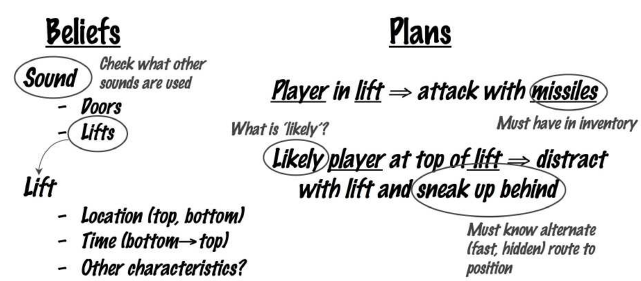

Figure 2: Working notes associated with previous interview fragment.

4.5 Further Iterations

The analysis of interview data as described above served both to construct the agents and to highlight the gaps that needed to be filled. For example, follow-up questions to the above included things such as

- "You said you used sounds like doors opening and lifts activating. Which other sounds do you pay attention to?"
- "With your trick of activating a lift and then trying to run around and surprise a person lobbing grenades in, do you just assume that there will be someone there, or do you use other information to decide when/if to do this?"

- "If there was more than one route that would get you to the top of the lift in time, how would you decide which one to use?"

4.6 Knowing When to Stop

These two steps of detailed interviews followed by analysis which in turn generates further questions must be repeated until sufficient data has been gathered to construct a model. How then should the interviewer recognise that sufficient data has been obtained? In the example here, the models connect to the Quake 2 engine in the same way that a player's client software would. This meant that the models needed to generate actions at the level of keystrokes and mouse clicks and movements. However subjects do not naturally express their behaviour and reasoning in such low-level terms, nor is it appropriate to expect them to do so. Thus it was necessary to determine the lowest level in which the subjects naturally referred to actions and beliefs and build a bridge in the model between this point and the low level actions that could be performed.

The lowest level in which the subjects naturally expressed actions was that used in expressions such as "fire into the lift shaft" or "move to the doorway." Below this level, the actions taken were part of an action-feedback loop, where constant adjustments to movement were made with mouse or keyboard based on the visual feedback from the game, but the subjects did not consciously think about it. This gap between the level at which the players described their behaviour and the level at which the engine accepted input was addressed using an extension similar to that presented in [9].

Finally it must be noted that It is important to follow through all avenues of questioning with *each* subject, even when they appear to be saying the same thing. In the example presented here, the subjects were sometimes using the same expressions to mean quite different things. Take for example the idea of "sufficient" health to keep fighting:

> **First subject:** ...basically I try to keep my health above sixty percent, roughly sixty percent.
>
> **Second subject:** ...I think it'd be probably around twenty-five percent health [that I would run away from a fight]

Obviously the subjects had very different ideas about when to quit a fight. If it had been assumed that "low health" (in terms of the time to run away from a fight) was the same for both subjects, one of the models would have been quite inaccurate.

5. DISCUSSION

The methodology described above is suitable for eliciting knowledge from subject matter experts in order to build BDI-based models of human behaviour. It has been demonstrated with the development of two models of Quake 2 players, and has been applied to other modelling problems. Gathering the knowledge required for complex models of human behaviour will never be a trivial task, but this methodology at least provides a systematic approach.

The approach assumes that subject matter experts will be the primary source of knowledge for the models, with additional sources as needed to bridge the gap between the knowledge that the SMEs are capable of supplying and what is required by the modelling environment. Bridging this gap *can* be a difficult task in itself, but it is important to realise what the limits of the SMEs' knowledge are, and access other sources (such as [1]) for the knowledge required to fill this gap.

The iterative nature of the interview-analyse-design process may be regarded as an expensive approach, but careful preparation can minimise the number of iterations required. In particular, it is important for the interviewer to be able to 'speak the same language' as the experts yet not be biased by his/her own biases.

6. ACKNOWLEDGEMENTS

Support for this work was provided by a Strategic Partners with Industry Research and Technology (SPIRT) award from the Australian Research Council and Agent Oriented Software Pty Ltd. The author would like to give thanks to the anonymous reviewers of this paper who wished to see more detail. Unfortunately the constraints of a short paper prohibit this, but further details may be found in [8].

7. REFERENCES

[1] K. R. Boff and J. E. Lincoln, editors. *Engineering Data Compendium: Human Perception and Performance.* Harry G. Armstrong Aerospace Medical Research Laboratory, Wright-Patterson Airforce Base, Ohio, 1988.

[2] B. Crandall and K. Getchell-Reiter. Critical decision method: A technique for eliciting concrete assessment indicators from the "intuition" of NICU nurses. *Advances in Nursing Sciences*, 16(1):42–51, 1993.

[3] G. A. Klein, R. Calderwood, and D. MacGregor. Critical decision method for eliciting knowledge. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(3):462–472, 1989.

[4] J. E. Laird, R. M. Jones, and P. E. Nielsen. Coordinated behavior of computer generated forces in TacAir-Soar. In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation*, Orlando, FL, 1994.

[5] D. McIlroy and C. Heinze. Air combat tactics implementation in the Smart Whole AiR Mission Model (SWARMM). In *Proceedings of the First International SimTecT Conference*, Melbourne, Australia, 1996.

[6] L. G. Militello and R. J. B. Hutton. Applied Cognitive Task Analysis (ACTA): A practitioner's toolkit for understanding cognitive task demands. *Ergonomics*, 41:1618–1641, 1998.

[7] G. Murray, S. Steuart, D. Appla, D. McIlroy, C. Heinze, M. Cross, A. Chandran, R. Raszka, G. Tidhar, A. Rao, A. Pegler, D. Morley, and P. Busetta. The challenge of whole air mission modelling. In *Proceedings of the Australian Joint Conference on Artificial Intelligence*, Melbourne, Australia, 1995.

[8] E. Norling. *Modelling Human Behaviour with BDI Agents.* PhD thesis, The University of Melbourne, forthcoming.

[9] E. Norling and F. E. Ritter. Embodying the JACK agent architecture. In M. Stumptner, D. Corbett, and M. Brooks, editors, *AI 2001: Advances in Artificial Intelligence*, volume 2256 of *Springer Lecture Notes in Artificial Intelligence*, pages 368–377. Springer, 2001.