An Approach to Capturing Cognitive Processes
Extended Abstract

Emma Norling

Centre for Policy Modelling, Business School,
Manchester Metropolitan University
norling@acm.org

Abstract. Capturing the cognitive processes — the underlying reasoning behind action — of individuals is not always a goal in social simulation. Indeed many social simulations seek to pare back models of human behaviour to the bare essentials, simplifying both the construction and running of these models. For some types of simulation however, more detailed models of human behaviour are required, when it is important to understand not just *what* individuals are doing, but *why* they are doing it. In these cases, one must capture the cognitive processes of the subjects being modelled, but going about this is no simple process. People are not always easily able to explain the reasoning behind their actions, and it can take some skill for an interviewer to ask the right questions to untangle a person's description of what they do from how they decide what they are going to do. This extended abstract outlines a structured approach to this problem, which focuses on the cogntive processes of the subjects to be modelled.

1 Introduction

Why is gathering knowledge to model human behaviour so difficult? Essentially it is because we don't perfectly understand what drives people to do the things they do. If we are modelling a physical object, such as a car, it is possible to understand exactly how it works, in terms of the physical interactions of its component parts. From this understanding it is possible to produce a computational model that exactly mimics, at a mathematical level, the interactions of its various components (pistons, fly wheels, gaskets, etc) to produce its overall behaviour. That of course is often far more detail than is required in a simulation, so instead an object is often modelled in terms of its affordances (for example, the car can go forward, backward, turn, travels at a certain speed); an abstraction of the detail in the first model. For people, in general we don't know the details of *how* people work. Nor is it sensible really to think in terms of affordances; that type of model is practical where it is known that a particular action will cause a particular reaction (for example, turning the steering wheel to the left will cause the car to turn left). When it comes to people, there aren't in general such straightforward relationships. If there were, social simulation would be a whole lot easier, but life might be a bit boring.

Instead, we are left trying to understand the behaviour of actors in the scenario of interest, and trying to translate this into formal rules. The exact form of these rules will depend on the particular implementation language being used, but even the most "high-level" language is necessarily far more formal and precise than a narrative description of human behaviour. Somehow the model builder must bridge that gap, providing the precision required for implementation where perhaps in reality there is none. How should this be done? In many cases, for simplicity or from plain necessity, the model builder will "kludge" this, implementing something that is practical and seems sensible. It is desirable, however, to minimise such kludges (because they might unintentionally significantly influence the simulation outcome), instead trying to understand and implement the actual process.

The approach presented here is designed for models which attempt to model the *cognitive process* of (at least the key) human actors in a simulation, that is, not just *what* they do, but why they choose to take particular courses of action.

2 Drawing upon Cognitive Task Analysis

The methodology presented here draws largely upon two forms of cognitive task analysis: applied cognitive task analysis (ACTA) [1], and the critical decision method (CDM) [2]. Limitations of space prevent a full summary of these methodologies here; the approach described below explains the parts of each which are used. ACTA is in a sense a top-down approach, starting with the big picture, and working down to the details. CDM is complementary because it focuses on filling in the details, assuming the big picture is already there. Neither approach was designed with modelling human behaviour in mind. Understanding human behaviour, yes, but not necessarily providing the data required to implement a model of human behaviour. In order to obtain the data required to construct a model, a combination of techniques, primarily drawing upon the above two methodologies has been developed, which is outlined below.

The goal of this methodology is to capture cognitive processes, in order to generate behaviour in agent models of human behaviour that is based on (some abstraction of) the reasoning of the people being modeled. The level of abstraction required will depend on a number of factors. Firstly, how "smart" do the models need to be? Do you want them to be able to respond "sensibly" in unanticipated situations? If so, you probably need quite detailed reasoning models. But on the other hand, if your agents are performing routine tasks, and your focus is instead on the unexpected outcomes of the interactions of these tasks, a less detailed reasoning model may well be adequate. Secondly, from a pragmatic position, both the language used for implementation and the simulation environment will be significant factors in determining the level of detail required. And thirdly, it is important to recognise that this methodology looks at the subjects' *concious* reasoning; beyond a certain level of detail subjects will not be able to explain why or how they perform certain actions, they "just happen."

This methodology was used to develop models of two expert Quake 2 players. Quake 2 is a first-person shooter game, where players compete against each other to get the most kills. The two experts being modelled had quite different styles of playing the game: one was an aggressive player who would seek out combat, while the other was a “camper,” who would find a hiding place and wait for opportunies. When playing together, neither player was dominant. The game itself is fast-moving, and decisions must be made quickly in order for players to be successful. The aim of developing these models was to capture the two different playing styles, and ultimately, to develop models that more closely resembled human players than existing computer-generated characters did. For the purpose of this paper, the aim of discussing these models is to demonstrate how the above methodology can be used in developing fairly detailed models of human behaviour. For most social simulations, the level of detail will be far less than here. The person gathering knowledge should ensure he/she knows the level of detail that will be required before starting, so as to avoid wasted effort in gathering unnecessary detail, or the need for additional interviews to acquire missing information.

2.1 Direct Observation

Direct observation can serve three purposes, as elaborated below.

Filtering Depending on the purpose of the models, there may be a number of potential candidates to be modelled, but only a subset to be modelled. Observing the subjects in action can be useful for this purpose. It might be that a diversity of styles needs to be captured, for example, or there is interest in a particular approach. However filtering of subjects is not always required; it will depend on the application.

Preparation Direct observation can also be useful in helping to prepare for the interview stage. There is a delicate balance that must be met by the interviewer in terms of understanding the task. He/she should understand the task well enough to have meaningful discussion with the subjects, but also should avoid the pitfall of being focused on a particular way of performing the task (which might preclude him/her from recognising and/or understanding the alternative method(s) used by subjects). Direct observation can be a useful tool for avoiding either extreme.

Reference Points A third use for the direct observation is to provide reference points — particular incidents that can be revisited and discussed — for later interviews. While not an essential tool, in some cases it may be easier for either the subject or interviewer to refer to a particular instance that occurred during the demonstration. If the observation was recorded, it can be used to replay specific excerpts and discuss the reasoning behind the actions.

2.2 Development of a Task Diagram

The first interview in each case was designed to develop the task diagram for each subject. This is the "big picture" view of the task, breaking it down into the major stages (usually between three and seven). The stages are labelled and any timings are noted. The diagram that emerges serves as the basis for subsequent interviews. A set of questions is prepared in advance for this interview, but the key is to develop the task diagram, and the interviewer should drop or add questions as necessary to achieve this aim.

2.3 Expanding the Task Diagram

Expanding the task diagram was an iterative process of interviews followed by data analysis and model design. The number of iterations was variable, and depended in the interviewer's ability to identify the gaps in knowledge that were needed for developing the model during the interview. When the analysis and design phase identified gaps, the cycle must be repeated.

The second stage of interviews probed more deeply into the phases that were identified in the task diagrams, using three different strategies: 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. Whereas the first stage of interviews focused on the subjects' goals when performing their task, the second stage turned to the strategies that they used, and the way that they characterised the world. The interviews in this stage were open-ended. Before each interview, a list of questions was prepared, based upon the analysis of data from the previous interview. However reponses to questions at this stage often generated further questions.

The models that were constructed in this exercise used the JACK agent language [3], a BDI-based agent programming language that is implemented in Java. As such, the analysis of the data and design of the models involved identifying 1) the goals of the subject, 2) the plans used by the subject, and 3) the beliefs of the subject, or more accurately, the things about which the subject had beliefs. If the models were being implemented in a different agent language and/or paradigm, the interviewer should have a clear knowledge of form in which knowledge would be represented. This in turn would influence the form of the questions he/she would prepare for the interviews, but the same three strategies identified above would still be used.

3 Limitations and Future Work

The example given above obviously captures behaviour at a far more detailed level than is typically used in social simulation. Indeed, it was undertaken as part of a project looking at modelling individual human behaviour (though of course, this is almost always in a social context). In social simulation, it is rare that one would be required to model reasoning and behaviour at this level of detail, however the approach is appropriate for more abstract models as well, so long as the focus is on capturing cognitive process. The difference would be managed by a different level of probing in the second stage interviews.

In some circumstances, it may be appropriate to model "key players" in a simulation in detail, using the approach here, while the remainder of the population is implemented using a simpler approach. What remains open for exploration is to consider how to combine detailed models of behaviour in small numbers with a simpler model of the larger population. One approach would be to consider a small subset of the larger population, model them in detail, then somehow generalise these results to a simpler model of the larger population. Quite how this generalisation would happen is very much an open question. An alternate approach would be to identify "key players" in a simulation, modelling them in detail, but with simple models of the wider population. In this case, firstly, what consitutes a "key player" and secondly, how would the detailed and simple models of behaviour interact?

Another avenue of enquiry is to explore a potential synergy between this approach and *grounded theory* [4, 5]. Grounded theory takes a similar approach to knowledge acquisition, but the aim is the discovery of theories, rather than the development of models. However it can be argued that a model is in fact an implementation of a theory, so would a grounded theory in fact provide the information required to build a model, or indeed could the process of building a model from a grounded theory help to identify gaps in the theory?

References

1. Militello, L.G., Hutton, R.J.B.: Applied Cognitive Task Analysis (ACTA): A practitioner's toolkit for understanding cognitive task demands. Ergonomics **41** (1998) 1618–1641
2. Crandall, B., Getchell-Reiter, K.: Critical decision method: A technique for eliciting concrete assessment indicators from the "intuition" of NICU nurses. Advances in Nursing Sciences **16**(1) (1993) 42–51
3. Agent Oriented Software Pty. Ltd.: JACK Intelligent Agents. http://aosgrp.com/products/jack/index.html (Current at March 2012)
4. Glaser, B., Strauss, A.: Discovery of Grounded Theory. Strategies for Qualitative Research. Sociology Press (1967)
5. Strauss, A.: Qualitative analysis for social scientists. Cambridge University Press, Cambridge, England (1987)