PMI: Knowledge Elicitation and De Bono's Thinking Tools

M-M. Portmann & S. M. Easterbrook

School of Cognitive and Computing Sciences, University of Sussex
Falmer, Brighton, BN1 9QH. <mariap@cogs.susx.ac.uk>

Abstract. Much attention in knowledge acquisition has been directed at the question "What is Knowledge?". In this paper, we discuss a related question, which we consider to be of equal importance, namely "What is Thinking?". We present a definition of thinking that emphasizes the importance of arriving at new arrangements of knowledge, and discuss how having knowledge about something can be used to avoid thinking. Given this view, it is clear that stimulating an expert to *think* about the domain can provide more detailed knowledge about both the domain and about the expert himself. We have taken one of De Bono's thinking tools, the PMI (plus - minus - interesting) and built a knowledge elicitation tool for use in a domain where the expert's responses are likely to be based on unquestioned judgements. The tool requires the expert to think about the domain in ways that he is perhaps not used to, and the information elicited from this exercise gives an insight into the judgement policies of the expert. This, in turn, is of use when applying the knowledge, particularly where the resolution of conflicts becomes necessary.

1 Introduction

The problems of engineering knowledge-based systems inevitably beg the question "What is knowledge?". Many approaches to knowledge acquisition have arisen based on various perspectives on the nature of knowledge (see Boose 1989 for a survey), and experience with these has lead to a growing realisation that knowledge is not some objective essence to be mined and refined. Rather, knowledge seems to change its shape depending both on the task to which it is applied and the social setting. An extreme version of this view is expounded by Winograd & Flores (1986), who suggest that knowledge is socially constructed, and that it is in the act of communication that knowledge gets formulated.

The task of knowledge engineering then is to construct plausible domain models with the help of the expert. This involves stimulating the expert's thinking in order to discover how the knowledge is used. In this paper we examine the hypothesis that the use of de Bono's thinking tools in knowledge elicitation should prove advantageous in two main ways: they should allow a finer granularity of knowledge to be elicited; and they should help to characterise an expert, enabling better comparison between experts.

2 Background Work

Knowledge elicitation and knowledge representation have long been regarded as the bottle-neck in building intelligent systems. The human expert cannot be treated as part of a

machine, which means that psychological and sociological factors must be taken into consideration. This makes knowledge elicitation difficult and time-consuming. Three interconnected problems are involved:

1. The nature of the knowledge elicited: it has long been recognised that expert behaviour is not necessarily open to introspection, and that the knowledge used by the expert is compiled, and cannot be broken down into consecutive, neatly fitting chunks (du Boulay & Ross 1991).
2. The method of eliciting the knowledge: the type and quality of the knowledge acquired varies according to the elicitation technique, so that the availability of a range of techniques is desirable.
3. Conflict resolution: one cannot overlook the fact that the information an expert gives is coloured by his special interests in a domain, hence that knowledge needs to be elicited from more than one expert in order to cover the domain evenly. This in turn generates its own problems.

Similar problems also apply to users of a knowledge-based system: introspecting their requirements might be as difficult for the users as introspection of knowledge is for the experts, for very much the same reasons. Users of the system will also have special interests in the domain which might be different from the ones embodied in the system. In fact we are concerned with domains where there is no real distinction between experts and users. If the domain involves questions of subjective judgement, then the user might have just as much claim to "knowledge" as the expert.

2.1 Knowledge and the Interview

The most general elicitation technique is the interview. However, interviews are by no means trivial to conduct, nor is the data gathered from an interview always easy to analyse. The interview is made up of many different – possibly interfering – relationships:

- between expert and question: the expert might find the question threatening because it concerns an area of the domain in which he is not so hot;
- between interviewer and question: the interviewer might be personally interested in this;
- between interviewer and expert: differences of race, sex, class and age all might matter here;
- between interviewer and answer: does the interviewer understand the answer? Is it a matter of the interviewer being satisfied with the answer or not?;
- between expert and answer: where ‘answer’ needs to be read as ‘knowledge’.

LaFrance (1988) addresses the issue of knowledge being multi-modal, and makes use of the fact that different types of question elicit different *forms* of knowledge. Her *knowledge acquisition grid* organises knowledge and questions as "separate but interacting dimensions" (La France 1988, p.85). The "Types of Question"-dimension ranges from the open-ended to increasingly more specific and directed questions; the "Forms of Knowledge"-dimension ranges from more declarative ("what") to more procedural ("how") types of knowledge.

Gammack & Young (1985) justify the need for the "deeper" knowledge of the expert (beyond empirical rules) and hence of a more complex knowledge structure for expert systems, and also focus on a variety of methods for the elicitation of the different *kinds* of expert knowledge, but they admit that identifying some broad categories of knowledge can be done only for restricted domains.

There is merit in having a number of different question types, but it is hard to see how, in the process of an interview, the engineer can decide on the question type, as "...questions cannot only extract content but can follow it as well" (La France 1988, p. 89). A certain type of question does not necessarily guarantee a certain form of knowledge as a response, hence the information elicited by a question would need to be analysed before deciding on the type of the next question. The Grid may have its passive use, for decoding transcripts of interviews.

Whilst traversal of the Grid might be interesting for the expert, it is extremely demanding on the engineer. Further, it has been found that the relationship between question to be asked and question asked is not straightforward. Greenwell (1990) refers to Brenner (1982) as reporting that interviewers altered one third of the questions during the interview. This throws doubt on the usefulness of linear traversal of the grid.

LaFrance does admit that how the Grid is used (linear fashion or cyclical fashion) depends on the needs of different expert system teams, and that the use of the grid is a skill which engineers need to be trained for. However, DeLamater (1982) has found that "Using well trained interviewers ... [is] ...not related to responses they obtain" (DeLamater 1982, p. 37). He found that characteristics of respondents generate a greater difference of responses than question types.

2.2 Knowledge and Conflict

Shaw (1980; 1981) is concerned with conflicts not just between two or more experts, but within a single expert. Her work is based on Kelly's Personal Construct Theory (Kelly 1955) which holds that a person is made up of several sub-personalities, and that at any one moment, we might be operating with one of those sub-personalities. Furthermore, any communication act might be addressed not to another person, but to some sub-personality of this other person.

As a consequence of this Shaw identifies three aspects of conversation:

1. between sub-personalities of one person
2. between two sub-personalities, each of a different person
3. in a group of persons, which constitutes one or more sub-personalities.

Hence "adequate communion is dependent on the recognition and acceptance of difference both within and between people" (Shaw 1981, p. 132).

In Personal Construct Theory, a role, or sub-personality "implies a particular way of construing" (Shaw 1981, p. 129). This amounts to a particular conceptual system, and knowledge is treated as a widely shared and particularly significant conceptual system, seemingly "having an existence virtually independent of their carriers". This makes repertory grid methodologies based on personal construct psychology ideal for the elicitation and analysis of information from experts. The repertory grid is used as a conversational tool, or feedback device, relating to all aspects of conversation identified above. The tool takes over the tedious parts of the task of elicitation – hence the programs embody "content-free conversational algorithms" (Shaw 1980, p. 148). Shaw's tools do not force an analysis, and feedback facilities enable the user to change or eliminate constructs at any stage in the elicitation process.

In addition to eliciting the constructs and building the repertory grids, Shaw & Gaines (1988) compare and match grids, extract equivalences, measure similarities and highlight differences between grids to promote understanding between people and within individuals themselves. The tool SOCIO extends personal construct theory to knowledge acquisition. It uses the repertory grid not only for deriving conceptual systems but also for deriving relations between conceptual systems, of which four are identified:

consensus - between experts' views as a basis for communication.
correspondence - different terminology for shared concepts; a basis for mutual understanding.
conflict - same terminology for different concepts.
contrast - different terminology for different concepts; highlights difficulties for understanding.

SOCIO analyses and classifies the differences between conceptual systems. The derivation is entirely algorithmic, and there is no pressure on experts to reach consensus.

Critics of personal construct theory claim that it is very time-consuming, that it might intimidate the expert, and that Shaw & Gaines have not made it clear how the data is going to be used as the basis of a knowledge base. Greenwell (1990) concludes that "repertory grids have their place in knowledge engineering when the need arises in domains in which the conceptual structure is complex, limited in scope and without a clear nomenclature" (p. 55).

Certainly, personal construct theory is more appropriate and yields more useful information in some specific situations than others, i.e. in clinical psychology where the client voluntarily takes part in construct derivation as s/he has a vested interest to get to know their own various perspectives on the world. However, in knowledge acquisition, it is the knowledge engineer first and foremost who is interested in the expert's constructs, not the expert, and it is not certain whether experts are likely to be convinced of the usefulness of opening up the cabaret of their sub-personalities to the engineer. But without the co-operation of the expert, nothing would be gained, as Shaw & Gaines restrict themselves to the personal construct theory.

2.3 Conflict Resolution

Whilst Shaw & Gaines emphasise identification of conflict, Easterbrook (1991) focuses on encouraging conflict and the provision of productive resolution methods. By emphasizing collaborative resolution methods, such as education and negotiation, conflict is harnessed to explore the issues and assumptions underlying the experts' contributions. In this approach, a conflict is simply a "difference that matters". All different viewpoints are elicited, modelled and compared enabling differences to be captured and correspondences to be established. The comparison is intended to be exploratory, although it might offer some conflict resolution which might or might not be used. The support tool, Synoptic, provides mainly guidance through exploration of conflict, identification of conflict issues, generation of conflict resolution and evaluation of resolutions. Thus the final decision on the choice of a resolution happens in the final phase of the process.

The model described in Easterbrook (1991) relates mainly to requirements specification in systems analysis, where inevitably, information has to be gained from more than one person, and where each person's views might offer a different perspective of the function of the system to be built. It is an area where clashes occur between the needs of different groups of people within the organisation and where it is essential that none of the most important needs of any group should be compromised. While this requirement may be more acute in requirements specification, it is in essence the same requirement which has to be met in knowledge elicitation.

2.4 Summary

LaFrance's method treats the expert as a reactive system which can be prompted with questions and will come out with answers, and in this spirit she assumes that feeding it different question types will make it come out with different types of knowledge. This puts the emphasis on the external process, on what goes into the system (the questions), and what comes out of it; what goes on between the expert and the knowledge engineer, but not what goes on within the expert himself. Whilst Shaw and Gaines do take into account the psychological make-up of the expert, they employ a method which is primarily designed to get information about the expert himself, not about a domain.

The notion of encouraging conflict and then providing co-operative conflict resolution methods is a positive step forward. If knowledge is a social construct then the social context needs to be explored. Instead of maintaining a consensus, the different perspectives of various experts need to be captured and explored. The process of comparing conflicting perspectives then provides a productive focus for the exploration of the finer details of the expert's thinking. The comparisons reveal important information about the perspectives each expert takes, including hidden assumptions, and the relative importance attached to various issues.

Although various limited conflict resolution techniques have been proposed, these have yet to be integrated into knowledge acquisition methodologies. The model described in Easterbrook (1991) is limited in the sense that it leaves many questions still unanswered. In particular, it shares a fundamental flaw with all the techniques we have described: although it is assumed that the expert is a thinking being, there has been little consideration of what exactly this means - that is, the models take thinking for granted. We suggest that a fourth problem needs to be added to the list given at the beginning of this section, namely, "thinking".

3 What is "Thinking"?

If we are to define knowledge elicitation as getting information from a thinking subject, then the two crucial terms in this domain are knowledge and thinking. We shall attempt to define these terms:

Firstly, we will adopt a Wittgensteinian definition of knowledge, which amounts to knowledge being a matter of making a judgement about the truth of something. Any such judgement needs to be backed up by evidence, where such evidence is objective and publicly examinable. However, the judgement itself is subjective and does not say anything about the fact judged - it is merely "a characteristic of the manner in which I make judgements" (Wittgenstein 1984, para 149, page 47 - our translation). This means that the judgement says something about the person who is making it - a fact which this project exploits.

Definitions of thinking are harder to come by, as it is thinking behaviour, rather than the thinking process itself which has received the attention. De Bono regards thinking as a skill which can be learned, and hence can be done proficiently or badly, and that some ways of thinking are more profitable than others. We regard thinking as a kind of mental orienteering, which is best done with the help of tools. De Bono claims that traditional tools have led our thinking into rigid patterns and recommends attention directing tools to break those habitual and restricting patterns.

Our assumption is that the application of de Bono's thinking tools have a place here, and our enquiry starts with two hypotheses:

1. The application of different thinking patterns by the expert should in some way affect the information given by him. What exactly is to be expected has not been defined a priori, but we anticipate that the information should be of a finer grain than it would otherwise be, as the application of different thinking patterns should encourage the expert to consider different aspects of the subject.
2. The information gathered would not only say something about the domain, but about the expert, too - a fact which can be exploited in knowledge acquisition.

3.1 De Bono on Thinking and Conflict

In his Letters to Thinkers de Bono writes: "Thinking is our way of moving from one arrangement of knowledge to a better one." (De Bono 1988, p26). Taking this definition literally means that "if we had complete information in a situation then we would not have to think" (*ibid.* p26). A corollary of the definition is the paradox that "we might make better use of (the) information by not knowing about it" (*ibid.* p28). Knowing about something means having a specific arrangement of knowledge about something. Thinking about something is arranging the knowledge about that something in a specific way. Looking at all the information is looking at the domain via that specific arrangement.

This view of knowledge and thinking has some interesting implications. For example, if the knowledge claims authority in the domain, then we will not try to move to a better arrangement of knowledge, or to our own - we will not have to *think*. This corresponds to Heidegger's notion of *blindness* (Winograd & Flores 1986): the patterns of knowledge and the abstractions that we use to get by in everyday life blind us to other possibilities. De Bono suggests that it is better to have just enough information to be able to develop our own ideas about something, and then to look at the information in our own way. This gives thinking a chance to find new ways before being forced or tempted into the old concepts.

Conflicts between people often arise because people insist on their own arrangement of knowledge about something. This is the thinking style de Bono calls "small circle rightness" (De Bono 1988, p22) which only leads to further entrenchment of the conflicting parties within their position. de Bono has long been the advocate for a different style of thinking - lateral thinking - which is creative, as opposed to analytic thinking which is rigid and uses logic as its only resource, resulting in "small circle rightness". The tools of lateral thinking are of a quite different nature: provocation - mostly random, observation, exploration and intuition. The major difference between the two styles is this: Lateral thinking simply takes an *interest* in the topic, whilst analytic thinking is interested in the correctness of one's own argument or the fallacy of someone else's argument.

Analytic thinking has produced conflict resolution methods more appropriate to conflict perpetuation than conflict resolution - it merely changes the *shape* of the conflict: Instead of open warfare, there is a cold war - but still a war. In contrast, de Bono suggests a design approach to conflict. Traditionally, a conflict is a state of affairs in which the conflicting parties not so much try to win over one another, but in which both parties are trying to prolong their "state of victory": both sides are winning, because neither side is giving in to the other. Whilst both parties may wish for the end of the conflict, victory, in effect, consists of the prolongation of the dispute.

For de Bono, a conflict "is a situation with different perception, principles, needs and emotions" (*ibid.* p234). The conflict resolution task is "to design an outcome". This approach concentrates on a possible state of affairs *without* conflict, instead of "trying to reduce the concept to its basic confrontation simplicity" (*ibid.* p234-5) - in other words, concentrating on the clash of fundamental principles. Paradoxically, this simplification

results the conflict becoming even more irresolvable. The design stance, by doing the opposite, achieves the opposite: by complicating the concept or conflict situation, it "enrich[es] the situation so that a design can be made" (*ibid.* p236). This process of complication is a creative process. It relies on observation, not evaluation: new perspectives are added, new comparisons made, different scenarios thought of. This involves both parties in an exercise of exploration, not of confrontation.

This may even result in a re-definition of the conflict which might lead to the perception of different approaches to solving the conflict. Thus the design stance sets out to achieve a desired situation (instead of trying to eliminate an undesired problem). This turns the eyes of both parties in the same direction - i.e. of a common future.

3.2 PMI - the Thinking Tool

The idea of enrichment through exploration, as opposed to restriction through judgement, is inherent in PMI, one of de Bono's most powerful thinking tools "that is so simple that it is almost unlearnable" (De Bono 1985, p19). P stands for Plus or good points, M stands for Minus or bad points, I stands for Interesting or interesting points. The tool "should most especially be used when we have no doubt about the situation" (*ibid.* p23), in other words, when we are tempted to refuse to explore a situation further, as our judgement has been firmly made.

The PMI is intended to direct the attention of the thinker to those aspects which might otherwise be ignored. It would be wrong to think that doing a PMI consists of simply listing all the points we can think of concerning a situation, and then ordering them into plus, minus and interesting points. This would be a classification exercise, not an exploratory one. Doing a PMI consists of looking into the direction of Plus, Minus and Interesting, and giving roughly equal time to the exploration of each of these directions. The PMI "is always from the point of view of the thinker doing the PMI" (De Bono 1988, p85): the thinker is not required to list everything that is generally known about a situation, but to list the points he finds when looking at the situation through the channels of Plus, Minus and Interesting - but always from where he happens to be. Again, it is not a matter of finding out what is right, or correct: "No point of itself is Plus or Minus. Those are just directions in which the thinker looks" (*ibid.* p85).

At the same time, the PMI serves to by-pass our - in de Bono's view - naturally reactive and emotional judgements, by allowing a situation to be explored, whether it is liked or not. In the judgemental stance, any points listed would only ever be the points which back up the judgement already made. What is gained by the PMI - especially by looking through the channel Interesting - is the exploration of what is beyond acceptance (making a positive judgement) or rejection (making a negative judgement). To list points found in the direction of Interesting - though not necessarily leading to a reversal of one's opinion - might well lead to just that. Further, in situations of disagreement, I-points might allow a re-definition of the situation and to the perception of alternatives not previously observed. "'Interesting' is a signpost to nowhere and yet a signpost to everywhere" (*ibid.* p170). Again, Interesting points are not found to be interesting *after* they have been observed, but are *found* via the stance of the observer: "the sense of 'Interest' which leads the creative thinker to dwell on a point or an observation and then to look around that point to see what can be found." (*ibid.* p171).

4 Using PMI as a Knowledge Elicitation Tool

Using de Bono's tool, we have developed a program, also called PMI, which gives assistance to a knowledge engineer in the first stage of knowledge acquisition for an advice giving system. For this type of system the task of the knowledge engineer is not to get the *correct* information from an expert, but to get as finely grained information as possible from a variety of experts. Facts and goals are less well defined in advice-giving than in fault-finding or diagnosis - it is therefore harder to know what constitutes success. The one certain criterion for success is in fact the client's satisfaction with the advice given by the system. Such a system will need to satisfy a variety of clients from different economic and cultural backgrounds.

We take as an example of this kind of problem the domain of choosing a restaurant. We treat the problem as one of matching a client to a restaurant, and hypothesize that it should be possible to achieve this by matching a client with an expert. The problem of knowledge elicitation for such a system, then, is one of extracting the criteria which are responsible for the match between an expert and the restaurant s/he is giving information about. What emerges is that the task for an advice giving system in this domain - though not typically a diagnostic one - could be seen as just that: a client's requests concerning eating out are matched against the requests of a 'known case' - one of the experts - which then leads to the 'prescription' of a specific restaurant intended to satisfy the client's requests.

The system must be able to satisfy not only individual clients, but should be able to resolve the conflict between individuals within a group of clients with possibly differing motivations for eating out, and it is in this sense that the system goes beyond case-based reasoning. Traditional conflict resolution methods such as persuasion or negotiation leading to compromise are inadequate, and de Bono's thinking tools provide a framework for respecting the differences between people with clashing interests. Paradoxically, this is achieved by not focussing on them: Instead of the focus being on the conflict or the reasons for the conflict - that is the present, or past - it is directed towards finding a jointly acceptable way ahead, a future.

The program PMI elicits information from a variety of experts on a variety of restaurants, building up different perspectives of a restaurant and stereotypes of experts - to be matched with prospective clients for the restaurant. This is done by finding out not only what motivated the expert's choice, but which of his needs the choice is supposed to meet. The latter obviously determines the former, and it is a matter of working backwards from the information received to the original need to be satisfied, the purpose served by eating out. Though the basic requirement in the domain *is* eating out and the entity needed to meet this requirement is a restaurant, eating out is not necessarily motivated by the need to be fed. The purpose to be served by eating out *could* of course be to be fed, but it could also be to avoid having to cook, to enjoy the setting, to get away from home, to eat a speciality, and so on.

What might happen is that an individual expert, or client, could be at odds, not only with other experts, or fellow clients, but also with himself. For example, a person's judgement on a particular restaurant may change depending on the context in which that restaurant is being discussed. The rationale behind using de Bono's thinking tools for knowledge elicitation is that the system it belongs to should be able to offer more imaginative alternatives to the solution of conflicts than negotiation or compromise. When selecting a restaurant, compromise is usually a disappointment, whereas the alternative could be an adventure.

4.1 Characteristics of Test Domain

The utility of the PMI needs to be tested in a situation especially suited to it. Because the PMI is supposed to counteract reactive and emotional judgements, we have used a domain in which experts are required to express personal opinions rather than textbook knowledge. As the purpose of doing a PMI is not to get the correct information but to extract viewpoints, the domain needs to be one in which there are no canonical criteria for judging something right or wrong, yet one in which the expert has no doubt in the situation. In the restaurant domain, just about everyone has an opinion on at least one restaurant and feels himself expert enough to hold that opinion as any client to a restaurant is of course expert in judging his satisfaction with the restaurant. But what might be good for one expert (a steak eater) might be bad for another (a vegetarian). 'Noisy' is a Minus point only for someone who likes restaurants to be intimate and quiet. The same applies to the clients which an advice giving system would have to satisfy: Every client wants the restaurant to be 'good' - but what constitutes 'good' for one client is quite different from what constitutes 'good' for another.

There is bound to be disagreement between experts or clients, and the exploration of the 'Interesting' direction might help to find a jointly acceptable alternative for a group of disagreeing clients. The idea is not to get them to agree on points previously disagreed, but to find a way out of the deadlock by changing their motivation: instead of going to eat to satisfy the hunger, they might be persuaded to go to a restaurant to look at its decor, to enjoy the music, or to observe the weird clientele. This basically means that the clients are able to see the occasion of eating out in a different light.

4.2 Elicitation Method

Care has to be taken that knowledge elicitation does not become an exercise in classification or judgement for the expert. At the same time, it is obviously desirable that the information elicited covers as many aspects of the domain as possible. To present the expert with a questionnaire containing three columns - Plus, Minus and Interesting - forces him to concentrate on all three directions at once, which is not in the spirit of PMI.

As a pilot study for the project, two questionnaires were circulated to a test group. Both asked for information on two restaurants most of the subjects had been to. The first contained only the three columns in which to list Plus, Minus and Interesting points, together with a like / dislike box for each restaurant. The second had a list of numbered aspects pertaining to restaurants and eaters attached to it, with the intention that this would be used to number the points in the P, M and I columns *after* filling them in. We found that on the second questionnaire, people simply went down the list of aspects, writing their opinion next to the aspect, and then transferring it to either the P, M or I column. This is counteractive to the exercise of focussed exploration.

Clearly, listing all the aspects that ought to be covered on the questionnaire does, on the one hand, ascertain that no aspect gets overlooked for consideration. On the other hand, attention gets focussed on each aspect in turn, requiring a judgement, rather than an exploration. This defeats that aim of observing which aspects present themselves when focussing on one particular direction - P, M or I - in turn.

There is a trade off between closed questions and open questions: it has been found that certain aspects of, say, a job are mentioned 13 per cent less often in the open questions than with the closed questions, in which that aspect was one of the answer choices. Items

that are self evident are more often forgotten in open questions (Molenaar, 1982). Closed questions, however, go against the spirit of the PMI.

Consequently, passive knowledge elicitation methods are less well suited to the use of the PMI. Directed questions (requiring yes/no answers) too are to be avoided as much as possible. For these reasons we rejected the use of questionnaires. In order to explore the use of de Bono's tool, an open, but guided, interview with the expert needs be conducted, hence the development of the PMI program.

4.3 Ontology

The experiment with the questionnaires, and the considerations described in the previous section, led us to develop an ontology of the domain (Regoczei & Plantinga 1988) which would be accessible to the engineer only. The two main actors in the scenario of eating out are obviously the restaurant and the eater. What relates these two entities is the world they are both in and the food to be eaten. Studying the answers received from the test subjects helped us to arrive at the provisional ontology shown in figure 1.

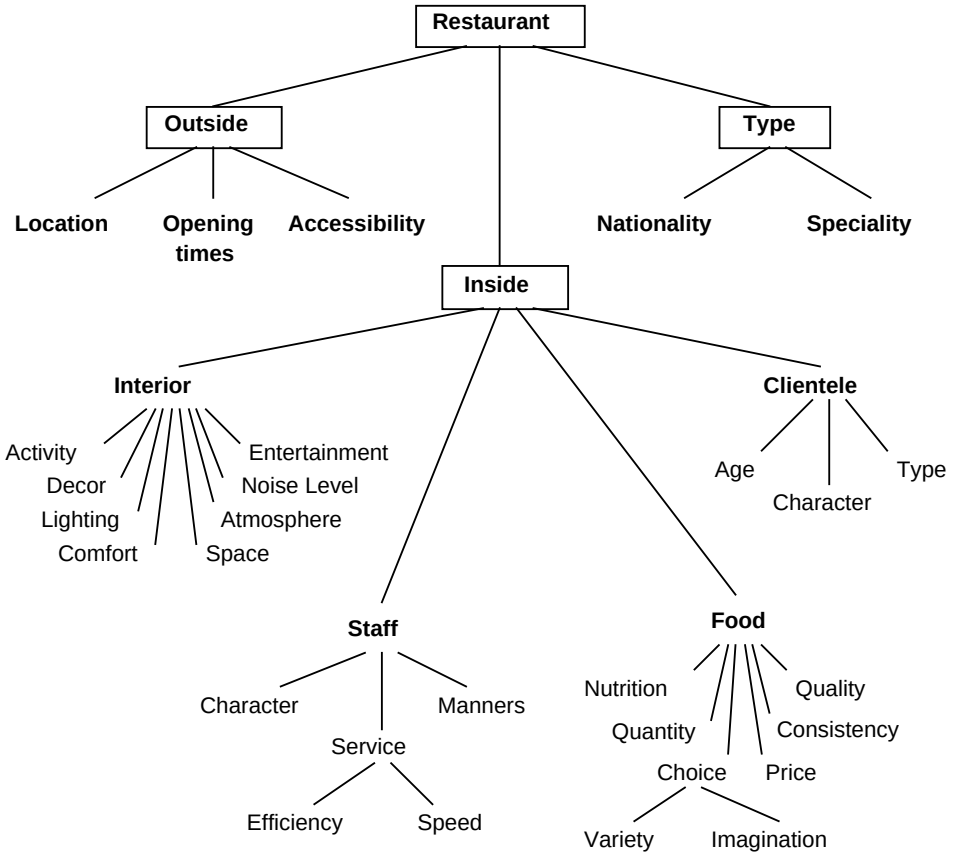

Fig. 1. A possible (hidden) ontology for the restaurant domain

At no time should the knowledge engineer force the expert's answers to fit the ontology. On the contrary, the ontology is flexible and will continue to evolve during the process of doing the PMI. Aspects commented on which are not in the ontology should be added to it. Certain groupings of aspects might have to be changed. Originally, the ontology included the aspect 'Menu'. After use of PMI with a number of experts, this was changed to 'Choice' and made a dependent of 'Food': people did not refer to the menu of a restaurant, but to the variety of dishes and the interestingness of choice.

The ontology should at most be used to provide consistency in the naming of aspects. When an expert refers to 'Price' he actually means a feature of the aspect 'Food', as this is what the restaurant sells. Hence 'price' would be entered as 'Food_price'. Again, this is not a matter of course, as the restaurant could be within a club which charges an entrance fee, in which case 'Price' could refer to this fee, or to the food. If it refers to the fee, then a new aspect 'Entrance_price' would have to be added to the ontology.

The ontology exists only on paper so far. It could be implemented as a look-up table, or as a hypertext. For the restaurant domain, the engineer can just about cope with it being on paper. In a domain with a large number of aspects to be discussed and where relationships between entities are more intricate, implementation would be a must. The structure would have to be flexible and updatable.

5 Implementation

The PMI knowledge elicitation tool consists of two modules, the first to gather the information, the second to analyse it. The first module can be said to be a prototype elicitation tool, the second module is still in its very early stages. The advantage of the use of PMI the thinking tool can be seen in the organisation of the elicited knowledge in the first module.

5.1 Module I - Elicitation

Doing the PMI. The data is collected via a straightforward interview between the knowledge engineer and the expert, with the engineer recording responses by interacting with the program, so that the expert is not distracted. The engineer explains the use of the PMI briefly to the expert, and - as far as is possible - unobtrusively monitors the amount of time spent on each of P, M and I. He should give the expert roughly the same amount of time for each of the three, even if the expert thinks he has no more comments to make.

At the start of the interview, the name of both the expert and the restaurant are recorded. The data collected will consist of lists of Plus, Minus and Interesting points. The expert is then asked whether he liked the restaurant or not. We could not avoid asking this yes/no question, as the answer to it will affect the way the information gathered will be used in the later analysis. However, it was important that the question should be asked as late as possible during the interview so as not to distract from the PMI.

In seeking alternatives to a closed question of liking or disliking a restaurant, we considered the possibility of asking whether the expert would *object* to eating in the restaurant under discussion, as this might disguise the like/dislike question. It turned out that the two questions elicit different answers: an expert might dislike a restaurant but not object to eating there. Although this fact could have been exploited in the analysis, for the purpose of this project we just used the like/dislike question. In fact, the pervasiveness of this question establishes the need for the PMI: as de Bono points out, the PMI is most useful when a person is in no doubt about something.

Building the Entities. The elicited knowledge sheds light both on the restaurant and on the expert or ‘eater’. The information gathered in the PMI constitutes the eater’s opinion on a restaurant (see figure 2). This is then used to characterise both the eater and the restaurant. Eventually, we want a fine-grained description of the restaurant which takes into account both people who like it and who dislike it. We also want a fine-grained description of a stereotypical eater who likes the restaurant, and a stereotypical eater who dislikes it. Stereotypes could of course be built up according to different criteria; for instance the stereotype "Student eater", or the stereotype "Businessman eater". Both the Eater stereotypes and the Restaurant entities are made up of more than one person’s information.

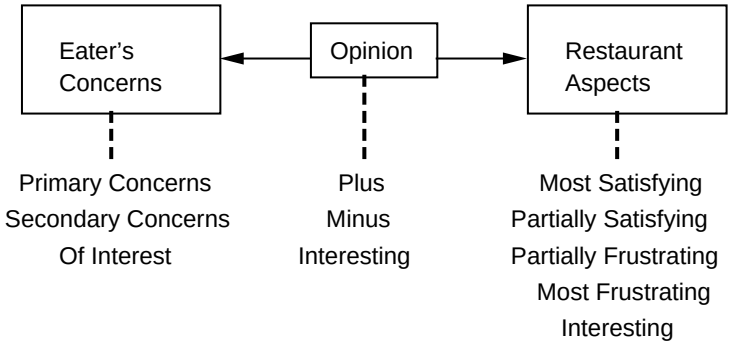

Fig. 2. The PMI gathers information about an eater’s opinion on a restaurant, which is then used to build eater stereotypes and restaurant descriptions.

Figure 3 shows the headings used for the eater stereotypes (Expert) and for the restaurant descriptions (Restaurant). Plus, Minus and Interesting are the lists of aspects (represented in the ontology), with the values the expert attaches to them. The three categories that make up the eater’s concerns are lists of aspects without corresponding values, as they abstract away from the particular restaurant. The categories provide information about which of the eater’s concerns are more likely to influence a decision about whether a restaurant is liked or disliked. The categories of the restaurant description are made up of aspects and the values attached to them, and provide information about which how different aspects affected eater’s opinions of the restaurant.

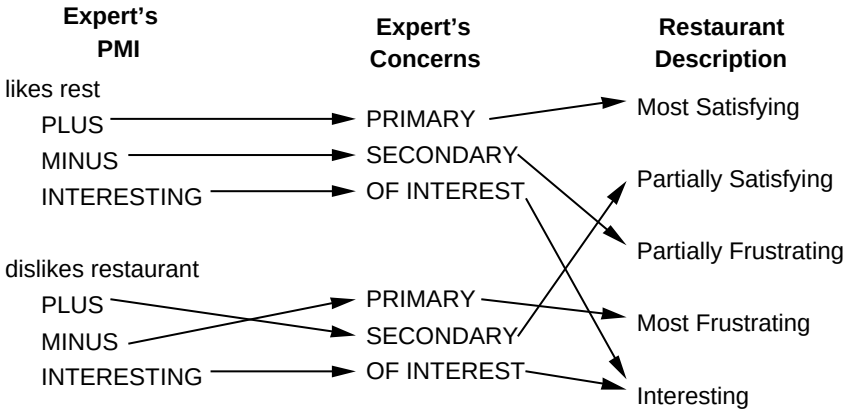

Fig. 3. Relationships between the P, M and I, and the Eater's concerns and the restaurant description.

The headings specify the relationship between expert and restaurant, but they might equally generalise as the relationship between restaurant and client. In generating these headings, we have made some assumptions about the expert's responses. In particular, we assume that there is a causal relationship between the plus and minus points elicited and whether the expert likes or dislikes the restaurant. For example, plus points for a restaurant that is liked overall are likely to be of greater importance to the expert than the minus points for that restaurant, while the reverse applies to minus points.

The following will clarify the considerations that led us to this choice of headings:

Expert

Primary Concerns are those points which need to be to the expert's satisfaction for him to like the restaurant. We assume that the plus points for a liked restaurant and the minus point for a disliked restaurant are significant enough to affect the expert's choice: they are primary concerns.

Secondary Concerns means the expert is not too bothered with these, whether they are to his general satisfaction or not. He will either like this restaurant, even if he is not satisfied by these points, or he will refuse to go to the restaurant, even if these points are to his satisfaction.

Of Interest are those points the expert finds interesting, whether he likes the restaurant or not.

Restaurant

Most Satisfying needs to be understood thus: If all these points are to the expert's or client's satisfaction, he will be satisfied by the restaurant and will like it. These are the plus points that caused people to like a restaurant.

Partially Satisfying reads: even if all these points are to the expert's or client's satisfaction, he may not be satisfied by this restaurant and may dislike it. While these may affect a person's opinion about a restaurant, they are not decisive enough to guarantee liking a restaurant.

Partially Frustrating reads: Even if all these points are frustrating the expert or client, he will still like the restaurant and be satisfied by it. These are the minus points that were not decisive enough for a person to dislike a restaurant.

Most Frustrating reads: if all these points are considered as minus points by the expert or client, he will find the restaurant frustrating and will dislike it. These are the minus points that caused a person to dislike a restaurant.

Amendments Options and General Concerns. PMI provides the expert with two amendment options: the Restaurant entity is built up after the first stage of the amendments option, the Expert entity gets built up after the second stage of the amendments option.

First Stage: Points can be added or deleted under any of the three P, M or I headings. Points can be moved from one heading to another. The wording used for listing the points can be altered. After each alteration, the new arrangement of the information is displayed.

The decision to give the expert this chance at all was considered carefully. One the one hand, this gives an expert who is not used to being interviewed the chance to avoid omitting the obvious. On the other hand, it could turn the whole of the PMI into what it tries not to be. Hence the expert should not be encouraged to add many aspects, nor should he be encouraged to re-classify them. The option to make changes should mainly be used to rectify mistakes the expert was aware of making during the PMI.

Up to this stage, the relationship between the entries under the 'Expert' headings and the entries under the 'Restaurant' headings correspond to the one described in Figure 3.

Second Stage: Any alterations made in this stage will change the entries for the expert's concerns, but NOT the entries under the 'Restaurant' headings. The reason for this way of doing things is that eventually we want to build up a stereotype of a *person* liking or disliking a restaurant, not a stereotype of a *relationship* between that person and the restaurant. Although the concerns have been extracted from the expert's discussion of the restaurant, they now no longer relate to that particular restaurant, but to eating out in restaurants generally.

The expert is given the chance to state what he is most concerned with when going to eat out. He is presented his primary concerns and his secondary concerns in relation to the restaurant discussed, and asked for changes to be made. The wording of the questions had to be carefully thought out and tested with several interview subjects in order to make them not too direct, yet to get a clear answer.

When an expert likes a restaurant, we know that he is obviously not too bothered by the Minus points he lists for it. But this does not mean that he is *generally* not bothered about these points. He may be bothered about them in principle but willing to ignore them in relation to a specific restaurant. We want to know which aspects might nevertheless be crucial for him generally. In other words, he is allowed to be human, which is to be inconsistent and ambiguous.

Similarly, if an expert dislikes a restaurant, we can be sure that he is not too bothered about the Plus points listed in relation to that restaurant. But we cannot know whether he is *generally* not bothered by these concerns. Some of them might be crucial for him in principle. We want to know which ones.

On the other hand, we know that all Plus points listed are to the satisfaction of the expert who likes the restaurant. But we cannot be sure that all of these points are equally important for the expert generally. Some of them might not need to be to his satisfaction for him to like the restaurant. We want to know which ones.

Similarly, if an expert dislikes a restaurant, we can be sure that in relation to that restaurant, he would be most concerned about the Minus points listed. (After all, he has rejected the restaurant despite its Plus points). But again, this does not mean that he is

necessarily bothered about them *in principle*. Some of them might not generally bother him: we want to find out which ones.

This approach offers a useful compromise between open questions and closed questions: the expert could not possibly be presented with an selection menu of all the points he might be concerned with when eating out. At the same time, we don't want the expert to have to think of points he might be concerned with *generally* in the middle of an interview about a *specific* restaurant. Thus *some* general concerns are extracted from the concerns emerging during the discussion of points relating to a specific restaurant.

Additional Information and Storage. The last section of Module I extracts some additional information about the expert such as age, gender, occupation, and the occasion of eating in the restaurant under discussion. This information may be used later in the analysis model for clarifying the stereotypes. For example, commonalities in the experts that contribute to a particular stereotype might be used to give a meaningful label to that stereotype. The tool then gives the engineer a chance to store the information extracted in the current interview cycle. The storing facility provides for data to be retrieved for use in a later interview session, and enables the engineer to gather and store a mass of information before analysing it in a separate session.

5.2 Module II - Analysis of Information

This module is under development. It is to provide retrieval facilities and information about data stored. It is also to prepare data for analysis and conflict resolution. Only some rudimentary facilities have so far been implemented, such as options to count and display items referred to several times, be they aspects of a restaurant, concerns of an expert, or values attached to aspects and concerns.

The main feature of this phase of the analysis is the combination of several entities into one. This is a must, as the analysis will be concerned with *stereotypes* of experts and *collective descriptions* of restaurants.

Stereotypes of experts. These are built up from many experts according to varying criteria: the stereotype liking or disliking a restaurant, the stereotype of a student, the stereotype of a twenty year old eating out, the stereotype female liking a particular restaurant, and so on. Many attributes of the eaters might be taken into account here: one might want the collective description of a restaurant given by all males between twenty and thirty, by all single people, by all couples with children. The choices will clearly depend on which characteristics were recorded for the experts in the interview: for the present study we selected a few of the more obvious ones. This provision is especially important in domains where there are reasons to believe that significant differences in description should emerge between groupings of experts according to different criteria.

Figure 4 shows an example of an stereotype eater from the restaurant domain. This stereotype is based on the criterion of liking a particular restaurant, and was built from five different experts' knowledge. The experts used for this example were all students, of roughly similar age. It is interesting to note that other commonalities occur across these experts, for instance, that in all cases the occasion was an evening meal. This tells us something about the relationship between this stereotype and the restaurant, ie. this particular restaurant is not frequented by this stereotype for lunch or tea.

General Information	
restaurant:	bystander
restaurant_type:	vegetarian (4); caribbean (1);
experts:	5
likes:	bystander (5);
gender:	female (4); male (1);
age:	25 (1); 23 (1); 27 (1); 28 (1); 26 (1);
occupation:	student (5);
occasion:	evening_meal (5);
Concerns of the Stereotype	
primary_concerns:	staff_manners (1); location (1); setting (1); lighting (1); atmosphere (3); food_choice (2); food_quality (4); food_price (5); service (2); food_choice_diet (1);
secondary_concerns:	space (2); accessibility (2); lighting (1); privacy (1); service_speed (4); food_choice_diet (1); windows (1);
of_interest:	music (3); flavouring (1); food_variety (2); food_choice (1);

Fig. 4. Stereotype of an Eater liking a particular Restaurant. The figures in brackets indicate how many experts have referred to a particular aspect.

Restaurant Descriptions. The collective restaurant descriptions are built up from the PMI points of many experts liking *and* disliking a restaurant (see figure 5). Storing expert and restaurant entities individually facilitates combination of entities in many different ways for many different purposes.

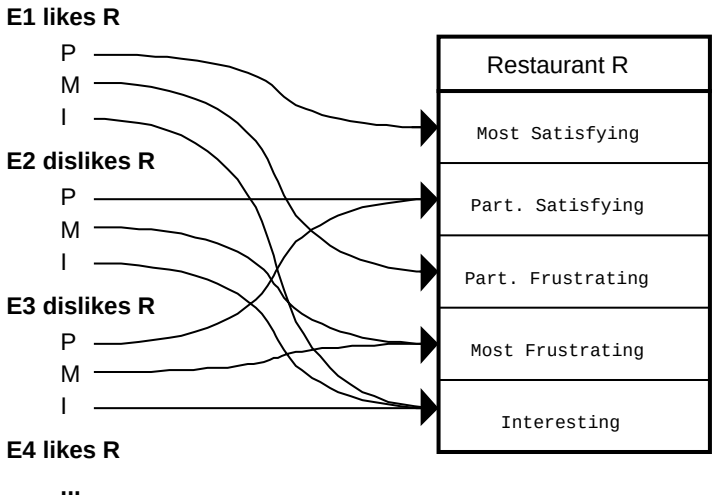

Fig. 5. A collective Description of a restaurant is build up from experts both liking and disliking the restaurant

Figures 6 gives an example of a restaurant description using the knowledge of ten experts. The aspects (in capitals) in this example will be part of the ontology, as described in

section 4.3. The values attached to these aspects are not part of the ontology, and it can be seen from the example that different experts might use different terms for roughly the same value. The assumption that these various terms might refer to the same value cannot be made before close scrutiny of the terms and their meaning. Hence, the terms have not forced to conform to a common vocabulary.

6 Conclusions

We have described a tool for elicitation of knowledge suitable for an advice giving system. The tool is based on de Bono's thinking tool PMI, and serves to elicit finer grain detail about an expert's opinion, especially where an expert would not normally have cause to question that opinion. The tool also provides a way of gathering information about the expert, so that generalisations about the experts can be used when applying their knowledge, for example, by comparing the experts' concerns with a client's. The elicitation tool has been used on a sample domain, that of choosing a restaurant, and we have shown how the tool can be used to analyse the domain, building up descriptions of both the experts and the restaurants they are discussing.

Simplicity of conception has been achieved: basically, the interview cycle is simply *one* expert giving information about *one* restaurant, whether or not that expert has been interviewed before, and whether or not that restaurant has been discussed by someone else before.

The tool has not been tested extensively yet, but from our initial observations with one study group, we have been able to deduce some characteristics of the expert from the information he has given. The information received has been found to be more balanced than it would have been with more directed questioning methods. Of course, for a realistic advice giving system, larger study groups, and many grouping criteria, would have to be used. This would require more time, and it is not clear whether the use of the PMI would, in the long run, save time and/or money.

Getting the experts to do the PMI was fairly easy and proved to be a relaxed way of getting the experts to talk. Once the subjects had understood the difference between "finding points through looking in the direction of Plus, Minus and Interesting" and "listing positive, negative and interesting points", and once they were told that what mattered was not the number of points found, but the concentration on *one* direction at a time, they relaxed and seemed to enjoy the session.

The tool provides a novel approach to Knowledge Acquisition. At the same time as building up information about the domain, we collect information about the expert and the expert's perspective on that domain. The more experts that contribute to the system, the richer a picture we achieve. The success of the system, i. e. the number of clients satisfied, increases proportionally with the number of experts consulted. This contrasts with other expert systems for which the difficulties in integrating knowledge increase if more experts are consulted.

The project stops just when things would start to become interesting: the evaluation and analysis of data collected. Further work is needed to develop the analysis phase, and to apply the knowledge elicited in a real advice-giving system. When we reach that stage, we will be ready to test whether our hypotheses are correct, by observing how the advice giving system performs. In particular, we would need to test whether the analysis we described above provides sufficient detail for advice to be matched to the client's needs, and also whether such a system can successfully deal with conflicts between the needs of a group of clients.

References

Boose, J. H. (1989) A survey of Knowledge Acquisition Techniques and Tools, *Knowledge Acquisition: An International Journal*, Vol 1, No 1, pp3-38.

De Bono, E. (1985) *De Bono's Thinking Course*, London: Ariel Books.

De Bono, E. (1988) *Letters to Thinkers: Further Thoughts on Lateral Thinking*, London: Penguin Books.

Du Boulay, B., and Ross, P. (1991) booknotes for MSc Intelligent Tutoring Systems, School of Cognitive and Computing Sciences, University of Sussex.

Brenner, M. (1982) "Response Effects of 'Role-restricted' Characteristics of the Interviewer", in W. Dijkstra and J. Van der Zouwen, eds., *Response Behaviour in the Survey Interview*, London: Academic Press.

DeLamater, J. (1982) "Response-effects of Question Content", in W. Dijkstra and J. van der Zouwen, eds., *Response Behaviour in the Survey Interview*, London: Academic Press.

Easterbrook, S. M., (1991) "Handling Conflict Between Domain Descriptions with Computer-Supported Negotiation", *Knowledge Acquisition: An International Journal*, Vol 3, pp 255-289.

Gammack, J. G. & Young, R. M. (1985) "Psychological Techniques for Eliciting Expert Knowledge", in M. A. Bramer, ed., *Research and Development in Expert Systems*, British Computer Society Workshop Series, Cambridge: Cambridge University Press.

Greenwell, M. (1990) "Knowledge Elicitation: principles and practice" in M. F. McTear and T. J. Anderson, eds., *Understanding Knowledge Engineering*, Chichester: Ellis Horwood.

Kelly, G.A. (1955) *The Psychology of Personal Constructs*, New York: Norton.

LaFrance, M. (1988) "The Knowledge Acquisition Grid: a method for training knowledge engineers", in B. Gaines and J. Boose, eds, *Knowledge Acquisition for Knowledge Based Systems, Vol. 1*, London: Academic Press.

Molenaar, N. J., (1982) "Response-effects of 'Formal' Characteristics of Questions", in W. Dijkstra and J. Van der Zouwen, *Response Behaviour in the Survey Interview*, London: Academic Press.

Regoczei, S. and Plantinga, E.P.O. (1988) "Ontology and Inventory", in B. Gaines and J. Boose, eds, *Knowledge Acquisition for Knowledge Based Systems, Vol. 2*, London: Academic Press.

Shaw, M. L. G. (1980) *On Becoming A Personal Scientist*, London: Academic Press.

Shaw, M. L. G. (1981) "Conversational Heuristics for Eliciting Shared Understanding", in M. L. G. Shaw, ed., *Recent Advances in Personal Construct Technology*, London: Academic Press.

Shaw, M. L. G., and Gaines, B. R. (1988) "A Methodology for Recognising Consensus, Correspondence, Conflict and Contrast in a Knowledge Acquisition System", Proceedings of the Workshop on Knowledge Acquisition for Knowledge-Based Systems, Banff, Nov. 7-11.

Wittgenstein, L., (1984) Ueber Gewissheit, Eds. G.E.M. Anscombe & G.H. von Wright, Suhrkamp Verlag.

Winograd, T., and Flores, F., (1986) *Understanding Computers and Cognition*, Reading MA: Addison Wesley.

Information relating to the Experts

name	- bystander
type	- vegetarian (8); caribbean (1); mexican (1);
interviewees	- 10
gender	- female (6); male (4);
age	- 28 (2); 25 (2); 27 (1); 26 (3); 23 (1); 22 (1);
occupation	- student (10);
bias	- likes (5); dislikes (5);
occasion	- evening_meal (10);

Categories of Satisfaction/Frustration

most_satisfying
PRIVACY has (1);
ATMOSPHERE relaxed (1); friendly (1); informal (1);
FOOD_CHOICE good (1); varied (1);
LIGHTING nice (1);
SERVICE friendly (2);
FOOD_CHOICE_DIET serve_vegan (1); vegetarian (1);
SETTING small (1);
FOOD_QUALITY good (3);
FOOD_PRICE good (1); reasonable (2); cheap (2);
STAFF_MANNERS friendly (1);
LOCATION easy_to_get_to (1);

partially_satisfying
FOOD_PRICE cheap (2); reasonable (1);
ATMOSPHERE good (1); friendly (1); cosy (1);
MUSIC play_your_own (1);
SETTING good_for_party (1);
STAFF_MANNERS friendly (3);
ATTITUDE tolerant (1); congenial (1);
LOCATION near_station (1); easy_to_reach (1);
WINE reasonable (1);

most_frustrating
FOOD_QUALITY bad (1); poor (1);
FOOD_CHOICE small (1); unimaginative (1);
LOCATION unpleasant (1);
SPACE cramped (3);
FOOD_QUANTITY small (1);
WINE expensive (1);
SERVICE_QUALITY abysmal (1);
SERVICE disorganised (1);
SERVICE_SPEED bad (1); slow (2);

partially_frustrating
ACCESSIBILITY no_disabled (2);
LIGHTING dark (1);
WINDOWS absent (1);
SERVICE_SPEED slow (4);
FOOD_QUALITY average (1);
SPACE cramped (2);

interesting
FOOD_TYPE interesting (1);
ATTITUDE casual (1);
FLAVOURING coconut (1);
FOOD_CHOICE interesting (1);
FOOD_VARIETY unusual (1); west_indian (1); interesting (1);
MUSIC play_your_own (3); varied (1);
DECOR unusual (1);
INTERIOR curious (1);

Fig. 6. Example of the contents of a restaurant description. The figures in brackets indicate how many experts have referred to a particular aspect.