

# Utility-based Agents for Hospital Scheduling: A Conjoint-Analytical Approach

Hans Czap, Marc Becker, Malte Poppensieker, Alexander Stotz

Department of Business Information Systems I  
University of Trier  
54296 Trier, Germany

{hans.czap, marc.becker, malte.poppensieker, stot4101}@uni-trier.de

**Abstract.** The scheduling of centralized operating theatres in large hospitals can be regarded as an archetypal cooperative decision problem. Multi-agent Systems (MAS) form an appealing paradigm to support different variants of Computer Supported Cooperative Work, like the situation presented in this paper. In an MAS-setting, each involved individual can be represented by an intelligent software agent that carries the specific constraints and the main preference-structures of his principal. The scheduling can then be done by inter-agent negotiations, resulting in a cooperative solution, which optimizes “social welfare” and medical and organizational resource allocation simultaneously. One of the most serious problems encountered in using such a utility-based approach to multi-agent scheduling is the actual measurement of human preferences and utility functions. In this paper, we propose using conjoint analysis as a solution to this problem and introduce a software component that can be easily integrated into existing agent-applications.<sup>1</sup>

## 1 Introduction

The scheduling of centralized operating theatres in large hospitals is a typical cooperative decision problem suited for delegation:

- It shows low involvement of the participants.
- It is time consuming for those who do the scheduling.
- It is characterized by a high degree of repetition, thus allowing for learning by feedback.

The process of scheduling operating room use involves different parties: surgeons, anesthetists and operating room nurses. Each one, in general, belongs to different – relatively autonomous – organizational units. Individuals, but also organizational units, have particular preference structures, which lead to conflicts of interests. These conflicts should be resolved in a fair manner. The problem of resolving different interests becomes worse by environmental influences: Hospitals operate in a highly

---

<sup>1</sup> The German National Research Foundation [DFG] has supported this paper as part of their priority research program “Multi-agent Systems and their Business Applications”.

dynamic and complex environment that generates the need to adapt very fast to changing environmental variables. In addition, at least in German hospitals, IT-systems will do the scheduling in a very simple manner, mainly showing tables with the assignments but not resolving any conflicts. Therefore, scheduling often takes place without IT-support.

Multi-agent Systems (MAS) are a well-known paradigm to support different variants of Computer Supported Co-operative Work, particular the situation presented here. Agent based scheduling allows a high grade of decentralization and flexibility. Besides conventional approaches, individual preferences can be taken into consideration. The work presented here assumes that each involved individual delegates the negotiation process to an intelligent software agent in order to arrive at a compromise schedule. Therefore, this agent must obey the hard constraints of hospital scheduling (which will be explained below) and he has to know the main preference-structures of his principal. These individual agents bargain about the resulting schedule. In doing so, they guarantee optimized social welfare and an optimized medical and organizational resource allocation.

In the next section, we present the details of scheduling operation theaters as we found it in a typical German hospital. In addition, we show how the scheduling by our agents will work. Focus of the paper lies on the problem of how to measure human preferences respectively utility functions, which is essential for delegation of any task. In order to understand the concepts, an introduction to general utility theory is necessary, which will be given in section 3. We show that although Expected-Utility-Theory (EU-Theory) still forms the backbone of game theoretic models, it is not very well suited for the application at hand. Physicians, nurses and so on, who are the decision makers in our hospital setting, are not specially trained in decision making necessary for expected-utility problems. Rather, we need a comparably simple decision scenario that is suited to elicit the preference structure of the people involved. In marketing theory, Conjoint-Analysis is a well-known technique, which aims at providing a simple way for the assessment of utility functions. In section 4, the basics of conjoint-analytical utility measurement are detailed. Finally, we present learning mechanisms that account for the inter-temporal validity of measurement.

## **2 The Policy-Agents-Research-Project**

The Policy-Agents-Project aims at providing a multi-agent system for hospital scheduling. In this section, first a general introduction to hospital scheduling is given. Afterwards a case study regarding typical operating theatre scheduling in large German hospitals is presented [8] that forms the application scenario of the Policy-Agents-Project. Based on this, the general project scope is further detailed.

### **2.1 General Aspects of Hospital Scheduling**

Scheduling in hospitals is done in two phases, in the large and the small. In the long run, patients get an approximate date of operation, which in some cases – if the operation is not vitally necessary – might be several months ahead. In Germany, as in

other European countries, the refunding of hospitals by public health insurances is highly regulated and limited on a yearly basis, yielding a backlog of needed operations. This backlog of several months essentially is caused by the limiting effects of the available yearly budget, which yields a steady number of operations per week independently of actual need. In general, one day before the fixed date of the operation, patients enter the hospital to be prepared. Public health insurances are very eager to shorten the time of patient stay at hospitals before operation and will not refund any overtime. Thus, after admission of a patient, there is a definite need to do the planned surgery operation as soon as possible. In the short run, scheduling of patients will happen from one day to the next.

In the following, we will concentrate on this short run scheduling of operation theatres. Short term scheduling is well known for being a process with an outcome that is highly dependent on situational variables:

- The duration of an operation in some cases cannot be determined beforehand. Thus, there is uncertainty of needed time.
- The specific tasks to be performed during an operation may depend on situations not observable during diagnosis (planning phase).
- The daily schedule often will be interrupted by incoming emergency cases. The frequency of emergencies depends on the medical department. For example, schedules of orthopedic departments seem to be very stable, whereas schedules of surgical or neurosurgical departments in general have to consider incoming emergencies by a high rate. If a patient is scheduled for operation today, it may happen that by incoming emergencies he suddenly will find himself being rescheduled for the next day.

In the analyzed hospital, rescheduling also results from a contract between nurses and hospital management that limits regular operating time from 8.00 am to 4.00 pm at the latest. Management had been forced to sign this contract, since dissatisfaction of nurses due to working overtime on a regular basis had been overwhelming. On the other hand, any postponing of an operation from one day to the next is a major source of dissatisfaction to patients, physicians, and management. It contributes essentially to the increase of the backlog and causes possible shortcuts in refunding since public insurances will not pay overtime in front of surgery operations.

## **2.2 The Conventional Process of Hospital Scheduling – Preferences and Constraints**

Traditionally, each medical department had its own operating room. This has been changed in hospital organizations where operating rooms are combined to form a centralized operating theatre. Thus, the operating rooms form an independent organizational complex, which is used by different medical departments. This leads to a better utilization of rooms, devices and personnel as well as to higher flexibility in planning and reacting to emergencies. On the other hand, transaction costs for manual scheduling rise dramatically, which causes the need for automated planning systems.

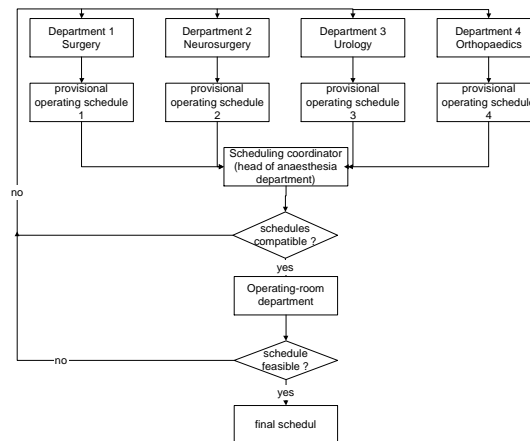
The conventional manual planning process, as we found, relies heavily on direct communication between different departments and follows in a sequential manner strict formal rules, reducing thus complexity.

In the first step, each medical department gets specific operating rooms at their daily disposal. In the second step, the department heads make a preliminary schedule of planned operations separately for each operating room. This preliminary schedule will obey department specific preferences (department policies). Some of them vary from department to department. Typical examples of these policies are:

- Patients to be scheduled are selected by urgency and availability of surgeries.
- Operations of children in general will be done before operations of adults. This reduces the probability of delays and reschedules for children.
- Performance of septical operations is scheduled at the end of day. This measure results from the time needed to disinfect the operating room afterwards, i.e. the room cannot be used immediately for a subsequent operation.

These preliminary schedules are transmitted to the operating theatre coordinator. In the conducted case study, the head of anaesthesia fulfilled this function. The operating theatre coordinator first checks compatibility of the preliminary plans with scarce resources. For example, in some cases the intensive care unit has shown to be a bottleneck. In addition, some operations need microscopes that are only available in a limited number. In the next step, the operating theatre coordinator decides about needed anaesthetists and makes an assignment. If incompatibilities occur, the coordinator has to contact the involved departments and negotiate a compromise solution.

Finally, the operating theatres plans are handed to the nurse personnel who have to assign the needed nurses. Again, incompatibilities respectively a shortcut of available operating room personnel has to be resolved by negotiation between the operating room department and the involved medical departments. The result of this negotiation process is the final operating room schedule.



**Fig. 1. Conventional process of scheduling**

This sequential planning procedure leads to measurable dissatisfaction of the lower working ranks, especially the operating room personnel, which can only react to the decisions made in earlier planning stages. In addition, the operating room nurses like

to have a definite time, when their regular work finishes in the evening. They are strongly dependent on the punctuality of the involved physicians, who in general are not very sensitive to the schedule. Therefore, a poor working climate results showing high fluctuation of operating room personnel. Figure 1 depicts the different planning steps.

### 2.3 The MAS-Solution

The scheduling problem presented here leads to different requirements for a multi-agent system:

- First, human-to-human interaction should be reduced to a minimum, thus reducing the needed time to resolve incompatibilities by phone.
- Second, sequential processes imply time-consuming feedback cycles, if any incompatibilities arise. Therefore, the planning procedure should happen in a simultaneous way, where any constraints are obeyed immediately.
- Third, a multi-agent system should take care of the organizational and individual interests of the involved parties whenever possible. This will give strong evidence for the acceptance of the planning system and will allow a better degree of staff satisfaction.

Based on the requirements an agent-based scheduling system for hospitals has been developed in the Policy-Agents-Project. In this paper, we give a short overview regarding the project scope as well as detailed references.

Scheduling of actions and resources in the Policy-Agent-MAS is done in two stages:

- In the first stage, the agent system creates a preliminary plan without respect to individual preferences. Only medical and organizational demands and constraints are taken into account. Therefore known scheduling approaches and algorithms can be used. The scheduler-component interacts with the planner via dialogs and offers him sub plans for modification, reordering or to place them into a Gantt-chart [5]. Sub plans consist of a set of actions, selected with respect to the constraints of the concepts of the used ontology OntHoS from the resource database [4].
- In the second stage, the preliminary plan is improved by agent negotiations. In this stage, the agents of the involved individuals, in the following called personal assistants, try to negotiate the best realizable working schedule for their principals. Using a negotiation approach based on the Nash-Bargaining-Solution the final schedule optimizes “social welfare”, i.e. respects the individual preferences of the involved staff, without scarifying medical or efficiency goals [7].

One of the most serious problems encountered in designing Policy-Agent MAS lies in the measurement of the preference structures of real human beings and their transformation into utility functions. For example, consider the simplified situation of table 1 comprising four medical departments, four anesthetists and four nurses.

Departments	Surgery	Orthopedics	Urology	Neurosurgery
Anesthetists	John	George	Jim	Andy
Nurses	Clara	Anne	Joyce	Hillary

**Table 1. Simple Assignment-Scenario**

The assignment-problem at hand, which is part of the scheduling problem, consists in finding a triple (Department, Anesthetist, Nurse), for example (Surgery, Jim, Anne), which will form a team to do the wished operation. Assume, the organization is indifferent to any concrete assignment, but imposes the constraint that all four anesthetists and nurses are assigned. Consequently, one has to find four triples, each consisting of an instance of  $(Department, Anesthetist, Nurse)$ . In addition, in our setting people involved will have specific preferences. For example, one might assume that John most preferably works together with Joyce in the surgery. He does not want to work with Anne in any department, whatever, and regards working with her as the worst case. In addition, working with Hillary in the surgery he likes as much as working with Clara in the urology.

Similar preferences exist for the other persons in our simple example. The problem in the following consists in finding for any person P a utility-function  $u_P(\cdot)$  that will measure by a real number any possible combinations of assignments and by this will express the preference of P. Thus, since:

$$\begin{aligned} (surgery, John, Clara) & \quad \tilde{u}_{John} & (surgery, John, Hillary), \\ (surgery, John, Hillary) & \quad \sim_{John} & (urology, John, Clara), \\ (urology, John, Clara) & \quad \tilde{u}_{John} & (surgery, John, Anne), \end{aligned}$$

the resulting utility function correctly must state:

$$\begin{aligned} u_{John}(surgery, John, Clara) & > u_{John}(surgery, John, Hillary), \\ u_{John}(surgery, John, Hillary) & = u_{John}(urology, John, Clara) \text{ and} \\ u_{John}(urology, John, Clara) & > u_{John}(surgery, John, Anne). \end{aligned}$$

In the following, we argue that the standard approach to preference and utility measurement, which is the Expected-Utility approach, is a rather theoretic concept not suited to practical application. Therefore, a different approach is necessary, which can be found in Conjoint-Analysis. We show how preferences, respectively utilities, are measured and how this fits to our agent application for hospital scheduling.

### 3 Preferences, Desires and Utility-Functions

British utilitarian philosophers and economists had introduced the concept of a utility during the eighteenth and nineteenth century. They understand utility as a measure of a person's wellbeing. In the view of these philosophers and economists, the concept of wellbeing relates directly to happiness. Thus, utility is the concept to explain a person's behavior as his attempt to maximize his wellbeing and happiness all at once [15].

Facing the difficulties of measuring a utility function, modern economists try to explain people's behavior in terms of preferences, i.e. ordinal comparisons of

different alternatives. Given the alternatives  $a_1$  and  $a_2$  they assume the individual to be able to express the relative preferences, i.e. to give an order  $a_1 \succ a_2$  or  $a_2 \succ a_1$  or  $a_1 \sim a_2$ . A utility function in this respect is understood as a real-valued function  $u(\cdot)$  that preserves the given preference order. In so far, the utility function of a person is nothing else than a convenient mathematical representation of his preferences. Such a utility-function is only valid on an ordinal scale of measurement. An ordinal approach suffices for the study of individual decision making where a single individual has to adapt optimally to a given situation, which cannot be altered in any way by this individual himself. This results in a simple maximization problem.

Negotiation problems are typical for many real life situations. Besides typical maximization problems, the solution of these kinds of problems depends not only on decisions taken by one party. Rather, the result depends on actions issued resp. decisions taken by any of the involved parties. Thus, all participants try to maximize their individual utility functions, of which they do not control all variables. Solving negotiation problems implies that concessions and compromises have to be made by each participant. This requires more knowledge than the information transferred in preference rankings. The individuals should not only know which alternative they prefer, they although have to have some notion about the strength preferring one alternative to another.

To solve this problem, von Neumann and Morgenstern used an axiomatic approach to prove the existence and the measurability of cardinal utility functions, based on the Bernoulli-Principle. They showed that by introducing probabilities and lotteries to the system of ordinal preference measurement, enough information is generated for deriving interval scaled utility functions valid up to linear transformations [21].

Given an individual with a complete preference structure expected-utility-theory assumes that the individual cannot only compare different alternatives but combinations of alternatives and stated probabilities. What this means for our hospital setting is best illustrated by the example given above regarding assignments of John.

First, complete preference structure means that John is able to express his preferences regarding all possible assignments (instances) of (Department, Anesthetist, Nurse), which counts to  $4^3 = 64$ , given - as in the example - 4 departments, 4 anesthetists and 4 nurses.

Second, John has to choose his best alternative,  $a_B = (\text{surgey, John, Clara})$ , and the worst,  $a_W = (\text{surgey, John, Anne})$ , and combines both by

$$l(p) = p*a_B + (1-p)*a_W, \quad 0 \leq p \leq 1,$$

to a lottery  $l(p)$ . This means, with probability  $p$  the best assignment  $a_B$  will happen and with probability  $(1-p)$  the worst one  $a_W$ . Since we have a complete ordering of all possible assignments, for any alternative  $a_i$ , which necessarily is ranked between  $a_W$  and  $a_B$  one can ask if there is a specific probability  $p_i$  which rates the sure alternative  $a_i$  and the lottery  $l(p_i)$  equally:

$$a_i \sim l(p_i), \quad l(p_i) = p_i*a_B + (1-p_i)*a_W, \quad 0 < p_i < 1.$$

Expected utility assumes that our individual John is able to provide these probabilities for each possible alternative  $a_i$ . Making the assignment  $u: a_i \rightarrow p_i$ , where the numerical value of  $p_i$  is meant, we arrive at a cardinal utility-function  $u(\cdot)$  [14].

The procedure shown, mainly aims at proving the existence of cardinal utility. It is sufficient for theoretical reasoning, but not intended to develop a reliable empirical method of utility measurement. Despite of this, expected utility is not suitable for most real-world applications, as will be shown below:

The first defect of the shown approach is that expected utility differs from both the classical concept, which is a numerical representation of wellbeing, and the modern one, which corresponds to the numerical representation of strength of preferences. Expected-utility inseparably mixes the individual's strength of preferences with the individual's attitude towards risk. As far as risky environments are considered, as is the case in financial analysis, this remains uncritical. In our settings, however, decisions typically are taken not under risk but under uncertainty. By decisions under risk the probability of the outcome is known (i.e. there is a probability and the value of this probability is known), whereas decisions under uncertainty specify situations, where probabilities of possible situations are completely unknown or even not meaningful. In other words, John might be able to rank the possible assignments, but definitely not to provide any probabilities combining possible outcomes of best and worst alternative. The differences between risk and uncertainty lead directly to the second defect of EU. It assumes that people are able to compare different probabilities, because they are trained to it from constantly making risky decisions in real life. Actually, they are only trained to make decisions in uncertain situations, which require using heuristics, best practices or feelings [9]. Most people are regularly untrained in evaluating probabilities, statements like: "if the probability of the good outcome is raised by 1%, then I am indifferent between the sure and the risky option" are seldom made in real life. Even trained decision theorists regularly exhibit irrational decision behavior in experiments about EU-measurement as was shown by Allais, resulting in the well-known Allais-Paradox [2].

The discussion above leads to the following results: Software agents that act as their principals' representatives need a utility function that emulates the principal's preference structure as exactly as possible. Utility functions based on EU cannot achieve this. In the following, we propose using Conjoint Analysis as an alternative method of utility measurement that is based on a measurement procedure resembling real life decision making more closely.

## **4 Conjoint Analysis as an Alternative Approach to Utility Measurement**

In comparison to expected utility, conjoint analysis (CA) offers an approach to utility measurement that has considerably lower cognitive demands on the individual. Instead of enhancing preference information by introducing lotteries and choice experiments, which adds a great deal of complexity to the measurement procedure, CA aims at statistically revealing additional information hidden in ordinal preference statements. This increases the effort for designing measurement interviews, but leads to significant reductions in complexity on the side of the respondent. In fact, he is only required to create a preference ranking over a selection of possible alternatives. Luce and Tukey laid out the foundations of CA in a seminal paper in 1964 [16].

During the last decades, CA has been successfully applied by psychologists and marketing researchers to a number of different problems, but was somewhat neglected in the decision sciences [10].

In CA, decision alternatives are described by a number of attributes, each attribute being made up of certain levels. In our example (table 1), attributes correspond to *Department*, *Anesthetist* and *Nurse*, and the level of each attribute to the possible instances as given by table 1. By decomposing the principal's ordinal preference evaluation (ranking) of the alternatives, CA assigns each of these levels a cardinal utility value, called part-worth-utility (short: part-worth). Relying on the fact, that through the ordering of multi-attributive objects (or alternatives) more information than a simple ordinal ranking is generated, the relative importance of each attribute level can be calculated and expressed as part-worth. This is done by analyzing the occurrence of the different attribute levels within the ranking. Combining the part-worths by means of an additive utility function an interval scaled total utility function – sufficient for agent negotiation applications – is generated.

Applying CA to agent applications can be done in three consecutive steps that will be detailed below:

- First, the application specific survey design must be created. In this step, the foundations of valid measurement are laid. Using CA as an interface between principal and agent, easiness of handling in combination with good validity of measurement is required.
- Second, the actual conjoint interview consisting of preference measurement and data analysis has to be conducted [3].
- In addition to traditional CA, in agent applications a third step has to be included. Utility functions of agents cannot be based on a static one-time evaluation. Rather, a learning mechanism should be established that accounts for the inter-temporal validity of measurement. Its main assignments consist of fixing measurement errors and adjusting to over-time changes in the principal's preference structure.

#### **4.1 Setting up the survey**

As stated earlier, CA demands some effort in constructing the survey; especially the correct determination of attributes and attribute levels for the specific application is crucial. In order to generate valid results, there are some constraints on the selection of attributes and attribute levels:

- All attributes relevant for the principal's decision have to be considered.
- The attributes must be independent of each other.
- There must be a compensatory relation between the attributes.

Though these constraints might appear to be very strict, most domains can be modeled according to them. Attributes in our application domain of scheduling of operation theaters relate to time-preferences (continuous working time, no overtime hours), to the kind of work (kind of surgery, specific task of nurses like theatre nurse

or background assistant) and to sympathy resp. antipathy of persons to collaborate within a team. Clearly, independence and compensatory effects can be assumed.

The alternatives, in our application the possible assignments, should be ranked by the individual. In general, since the number of all possible alternatives is far too large, only a subset will be presented for ranking. This subset is called the set of *stimuli*. The question, how to choose this subset leads to the distinction between:

- a full profile design,
- a random design and
- a systematic design.

In the full profile design, all possible alternatives are presented. Thus, every alternative becomes a *stimulus*. Because the amount of possible alternatives grows exponentially with the number of attributes, it can be overwhelming for any principal to rate all of them, even with a small number of attributes and levels. In order to make CA practically useful for agent applications, a method is required for reducing the quantity of stimuli, while still maintaining a good quality of the resulting utility function.

In a random design, a certain number of stimuli is chosen without respect to the distribution of the different attribute levels in the sample. It can be used for statistically estimating the utility over all alternatives. An advantage of this approach is that the size of the sample can be chosen arbitrarily. Still, as a major drawback, some attribute levels might not appear sufficiently often in the sample to allow an estimation of utility values with sufficient validity.

The systematic approach usually finds a better set of stimuli than the random design limiting the number of stimuli while still representing the set of all possible alternatives as close as possible. A systematic design can be used to guarantee uncorrelated estimation of all part-worths [13]. Establishing a systematic design is the method we propose for policy agents, since they should have a good understanding of their principal's preferences. Systematic designs, which allow uncorrelated estimation of the part-worths, relate to the work of Addelman, who has shown that the condition of "orthogonal frequencies" is sufficient to achieve this goal. It requires every attribute level to appear with all levels of the other attributes in proportional frequency to their number of appearances in the whole sample. Addelman calls designs that hold this condition *Orthogonal Main-Effect Plans (OMEPS)* [1].

In our approach, we chose an algorithm based on suggestions by Jacroux. His method guarantees computation of a minimal *OMEPS*, which consists of the smallest sample size still allowing uncorrelated estimation of part-worths [11].

## 4.2 Interview & Analysis

Having decided on the set of stimuli the next step is the analysis of the principal's preference structure. Therefore, the principal has to evaluate the stimuli. An intuitive method for evaluation is the ranking method. It requires the principal to assign a rank to each stimulus according to his preferences. Using CA in the context of agent systems, which relieve the principal from coordination tasks, implies that the

interaction with the principal should be as intuitive and easy as possible. We propose the following approach:

A limited number of stimuli is presented to the principal at once and has to be brought into the right order. Every additional stimulus is inserted into the existing order by pair wise choice, i.e. the principal repeatedly decides between the new stimulus and an already sorted one just by stating his preference between the two alternatives. After the ranking is finished, the part-worths for the different attribute levels are calculated based on the order of the stimuli revealed by the principal. Assuming an additive utility function, the principal's total utility for a multi-attribute  $a_0$  is represented by the sum of its part-worths [13]:

$$u(a_0) = \sum_{j=1}^J \sum_{m=1}^{M_j} \beta_{jm} * x_{jm}, \text{ where: } \begin{array}{l} u(a_0) = \text{total utility of alternative } a_0 \\ \beta_{jm} = \text{part-worth of level } m \text{ of attribute } j \\ x_{jm} = \begin{cases} 1 & \text{if level } m \text{ of attribute } j \text{ occurs in } a_0 \\ 0 & \text{else} \end{cases} \end{array}$$

Estimating the part-worths will be done by an ordinary least square regression (OLS). In doing so, to each ranked stimulus  $a_i$  a number  $z_i$  corresponding to the rank-level is assigned (most preferred stimuli get the highest number). By OLS the part worth utilities  $\mathbf{b}_{jm}$  are calculated, such that the sum of squared errors becomes minimal:

$$\min \sum_{i \in \text{set of stimuli}} (z_i - u(a_i))^2$$

In addition, the resulting  $\mathbf{b}_{jm}$  will undergo a process of normalization.

## 5 Learning

Considering the fact that individual preferences may change over time. A utility function that was determined once by CA cannot be regarded as statically valid forever. Instead, an agent system that is supposed to be in use for a longer period must be able to adjust dynamically to changes within the principal's preferences. That is, it needs to detect if the agents utility function still represents the principals preferences correctly and adjust it in case it does not.

To accomplish this task, some user interaction is required. As, obviously, intelligent agents are supposed to make their principals' lives easier, too much interaction is not beneficial.

Research has shown that while most users are willing to give some short feedback about the quality of the agent's work, they consider a longer procedure as frustrating and annoying [19]. Keeping that in mind, we have designed a procedure for permanently monitoring the quality of the agent's utility model, while reducing communication with the principal to a minimum.

## 5.1 Agent Preference Adaptation

The learning process proposed in this section is based on the main idea that in order to facilitate easy communication, the principal only needs to respond to a single question after selected negotiations done by the agent.

For this, he must evaluate the result of the negotiation ( $a_p$ ) ( $a_p = \text{primary assignment}$ ) together with the next possible alternative ( $a_s$ ), which might be the outcome of negotiation ( $a_s = \text{secondary assignment}$ ). Note that in the light of the agents' meaning  $a_p$  is the best possible result achievable during negotiation, whereas  $a_s$  is the second best one, that is the best possible assignment if  $a_p$  would not have been included in the set of possible alternatives. Consequently, the utility assignment expressed by the part-worths  $\beta_{jm}$  represents a preference order of the agent that states  $a_p \tilde{\mathbf{h}}_a a_s$  or  $a_p \sim_a a_s$ . (The symbols  $\tilde{\mathbf{h}}_a, \sim_a$  indicate preference ordering of agent, whereas  $\tilde{\mathbf{h}}_p, \sim_p$  that of principal.)

In order to get feedback, the agent presents both alternatives,  $a_p$  and  $a_s$  to his principal and asks for a ranking of these. Clearly, if the principal decides  $a_p \tilde{\mathbf{h}}_p a_s$ , or  $a_p \sim_p a_s$  and this coincides with preference order of agent on these two alternatives, there is no need to change anything. In all other cases, the preference order the agent had learned needs some adjustment.

This is a typical Credit-Assignment-Problem, where it is not evident which of the part-worths is to blame for the wrong decision [20]. In the following, we propose two methods for adaptation of the part-worths to count for the discrepancy in preference orders of agent and principal.

### Method 1

This method assumes only a "one-click" response of the principal. As the additional information-requirement is kept at a very low level, the improvement with respect to learning the correct preference (of principal) necessarily is very poor.

The method should be seen in the light of the arguments given under the heading "gradual learning". In general, one has to count for some *instability* of preference of principal implying the need for a learning procedure to *generalize of statistical noise* resulting of changing moods.

Given the need to adapt preferences of agent with respect to  $a_p$  and  $a_s$ , i.e.

$$(a_p \tilde{\mathbf{h}}_a a_s \text{ and } (a_p \hat{\mathbf{a}}_p a_s \text{ or } a_p \sim_p a_s)) \text{ or } (a_p \sim_a a_s \text{ and } (a_p \hat{\mathbf{a}}_p a_s \text{ or } a_p \tilde{\mathbf{h}}_p a_s))$$

At the first step the reduced design that led to the actual part-worths  $\beta_{jm}$  is augmented by the alternatives  $a_p$  and  $a_s$ , if these alternatives are not included in the reduced design, and the correct preference ordering is provided. Based on this changed reduced design new part-worths  $\beta'_{jm}$  are calculated.

In the case ( $a_p \sim_a a_s$  and ( $a_p \hat{\mathbf{a}}_p a_s$  or  $a_p \tilde{\mathbf{h}}_p a_s$ )) the resulting utility  $u'(\cdot)$  corresponding to the  $\beta'_{jm}$  necessarily expresses the correct preference with respect to  $a_p$  and  $a_s$ .

In the case ( $a_p \tilde{\mathbf{h}}_a a_s$  and ( $a_p \hat{\mathbf{a}}_p a_s$  or  $a_p \sim_p a_s$ )) part-worths  $\beta''_{jm}$  are calculated, such that the corresponding utilities  $u''(a_p)$  and  $u''(a_s)$  equal,  $u''(a_p) = u''(a_s)$ . Calculation

of  $\beta''_{jm}$  is as follows: If  $n = \sum_{j=1}^J \sum_{m=1}^{M_j} x_{jm}$  (symbols as above at the end of 4.2) designates

the number of all  $\beta_{jm}$ , then the tuple  $\beta = (\beta_{jm})$  and  $\beta' = (\beta'_{jm})$  are two points in  $\mathfrak{R}^n$ . On the line  $(\beta, \beta')$  a point  $\beta''$  exists yielding a utility-function  $u''(\cdot)$  having,  $u''(a_p) = u''(a_s)$ .

The resulting preference of agent consequently turns out to be  $a_p \sim_a a_s$ , which not in any case states the correct preference, but is a move into the right direction.

### Method 2

This method covers the case  $(a_p \not\sim a_s \text{ and } a_p \not\sim a)$  if at least one of the alternatives  $a_p$  or  $a_s$  is not included in the *set of stimuli*. By requiring additional information from the principal compared to method 1, it leads to a substantial improvement of the resulting utility function. The set of stimuli is augmented by the missing alternatives  $a_p$  or  $a_s$  (or both) and the principal is asked to rank these additional stimuli with respect to the others. Since there are at most two additional *stimuli* to be ranked, the necessary effort remains limited. The resulting utility correctly will state principal's preference.

## 5.2 Gradual learning in the case of Method 1

Psychology scholars distinguish two different forms of attitude change towards certain issues: Conditioning is associating the issue with a positive or negative mood created by another factor and thus changing the opinion towards it [17]. Persuasion is general communication that aims at altering decisions [6]. Both of these two forms are usually considered persistent.

Still, psychological research has shown that a great deal of a person's recorded attitudes depend on the person's actual mood. A learning algorithm must therefore try to recognize permanent changes of preference structure caused by either conditioning or persuasion but *should be invariant to temporal changes caused by certain moods*. In order to accomplish this, one needs a function for assessing the consistency of a perceived change in preferences with the user behaviour in the past. This consistency can be identified by the change of the n-tuples from  $\beta$  to  $\beta'$ . If  $\beta'$  is within a pre-defined vicinity of  $\beta$  changes are considered to correspond to a steady development, whereas in the other case a mood-based behaviour is assumed<sup>2</sup>.

## 6 Conclusion

This paper deals with a central problem of agent theory we encountered while designing an agent-based scheduling system suitable for hospital applications: that is how to align the actions of an agent to a principal's preference structure. Influenced by the long tradition of utility theory and its convincing theoretical results we propose

---

<sup>2</sup> While testing the algorithm, we got good results with a radius 0.2 as measure of vicinity.

the use of a utility-based decision component. This leads to the problem of how to determine the principal's utility function.

Expected Utility measurement, as the standard approach to utility measurement, turns out to show severe deficits proving not to be suited for applications of MAS. Conjoint analysis does serve this purpose and is well suited for determining human utility functions. To apply our ideas for real problems, we developed a specialized conjoint analysis tool for agent applications (*LACAM – Learning Agents and Conjoint Analytical Methods*), as part of our Policy-Agents Hospital-Scheduling MAS. Designed as a self-contained system component *LACAM* can be easily adapted to existing utility-based agent systems. Details may be found at: <http://www.wi.uni-trier.de/forschung/projekte/projekte/Agenten.htm>.

## References

1. Addelman, S.: Orthogonal Main-Effect Plans for Asymmetrical Factorial Experiments, in: *Technometrics*, Vol. 4(1), 1962, pp. 21-46.
2. Allais, M.: Le Comportement de l'Homme Rationnel devant le Risque: critique des Postulats et Axiomes de l'Ecole Américaine, in: *Econometrica*, Vol. 21, 1953, pp. 503-546.
3. Backhaus, K.; Erichson, B.; Plinke, W.; Weiber, R.: *Multivariate Analysemethoden – Eine anwendungsorientierte Einführung*, Berlin 2000.
4. Becker, M.; Heine, C.; Herrler, R.; Krempels, K.-H.: *OntHoS - an Ontology for Hospital Scenarios*, in: Moreno, A.; Nealon, J. (Hrsg.): *Applications of Software Agent Technology in the Health Care Domain*. Whitestein Series in Software Agent Technologies (WSSAT), Basel 2003, S.87-103.
5. Becker, M.; Krempels, K.-H.; Navarro, M.; Panchenko, A.: *Agent Based Scheduling of Operation Theaters*, Proceedings des EU-LAT E-Health-Workshop, Mexico 2003.
6. Bremsbeck, W.; Howell, W.: *Persuasion – A Means of Social Influence*, 2<sup>nd</sup> Ed., Englewood Cliffs, NJ 1976.
7. Czap, H.; Becker, M.: *Multi-Agent Systems and Microeconomic Theory: A Negotiation Approach to solve Scheduling Problems in High Dynamic Environments*, in Proceedings of 36th Annual Hawaii International Conference on System Sciences (CD-Rom), Hawaii 2003.
8. Czap, H.; Becker, M.; Schiffelholz, S.-D.: *Der Ablauf der kurzfristigen OP-Terminplanung am Beispiel des Krankenhauses der Barmherzigen Brüder, Trier e.V.*, Arbeitsbericht des Zentrums für Gesundheitsökonomie an der Universität Trier, Trier 2002.
9. Gigerenzer, G.; Selten, R. (eds.): *Bounded Rationality: The Adaptive Toolbox*, Cambridge 2001.
10. Green, P. E.; Krieger, A.M.; Wind, Y.: *Thirty Years of Conjoint Analysis: Reflections and Prospects*, in: *Interfaces*, Vol. 31/3, 2001, pp. 56-73.
11. Jacroux, M.: *A Note on the Determination and Construction of Minimal Orthogonal Main-Effect Plans*, in: *Technometrics*, Vol. 34(1), 1992, pp.92-96.
12. Kessler, M.; Poppensieker, M.; Porten, M.; Stotz, A.; Zub, D.: *Lernende Agenten & conjoint-analytische Verfahren - Entwicklung einer Conjoint-Analyse-Software zur Verwendung in FIPA-konformen Multiagentensystemen*, Studienprojekt am Lehrstuhl für Wirtschaftsinformatik I der Universität Trier, Trier 2004.

13. Klein, M.: Die Conjoint-Analyse: Eine Einführung in das Verfahren mit einem Ausblick auf mögliche sozialwissenschaftliche Anwendungen, in: ZA-Information No. 50, 2002, pp. 7-45.
14. Laux, H.: Entscheidungstheorie, Berlin et al 2003.
15. Luce, R. D.; Raiffa, H.: Games and Decisions – Introduction and Critical Survey, New York 1957.
16. Luce, R. D.; Tukey, J. W.: Simultaneous Conjoint Measurement: A New Type of Fundamental Measurement, in: Journal of Mathematical Psychology, Vol. 1, 1964, pp. 1-27.
17. Oskamp, S.: Attitudes and Opinions, 2<sup>nd</sup> Ed., Englewood Cliffs, NJ 1991.
18. Raghavarao, D.: Constructions and Combinatorial Problems in Design of Experiments, New York 1971.
19. Schiaffino, S.; Amandi, A.: User – interface agent interaction: personalization issues, in: International Journal of Human-Computer Studies, Vol. 60, 2004, pp. 129-148.
20. Sen, S.; Weiss, G.: Learning in Multiagent Systems, in: Weiss, G. (ed.): Multiagent Systems – A Modern Approach to Distributed Artificial Intelligence, Cambridge 1999.
21. Von Neumann, J.; Morgenstern O.: Theory of Games and Economic Behavior, Princeton 1942.