

Modelling Shared Decision Making in Medical Negotiations: Interactive Training with Cognitive Agents*

Volha Petukhova¹, Firuza Sharifullaeva² and Dietrich Klakow¹

¹Spoken Language Systems Group, Saarland University, Germany

²Saarland Informatics Campus, Saarbrücken, Germany

{v.petukhova,dietrich.klakow}@lsv.uni-saarland.de; s8fishar@stud.uni-saarland.de

October 11, 2019

Abstract

In the past decade, increasingly sophisticated models have been developed to determine which strategy explains human decision behaviour the best. In this paper, we model shared decision making in medical negotiations. Cognitive agents, who simulate various types of patients and are equipped with basic negotiation and decision making strategies, are tested in social learning setting. Human trainees were prompted to learn to make decisions analysing consequences of their own and partner's actions. Human-human and human-agent negotiations were evaluated in terms of the number of agreements reached and their Pareto efficiency, the number of the accepted negative deals and the cooperativeness of the negotiators' actions. The results show that agents can act as credible opponents to train efficient decision making strategies while improving negotiation performance. Agents with compensatory strategies integrate all available information and explore action-outcome connections the best. Agents that match and coordinate their decisions with their partners show convincing abilities for social mirroring and cooperative actions, skills that are important for human medical professionals to master. Simple non-compensatory heuristics are shown to be at least as accurate, and in complex scenarios even more effective, than the cognitive-intensive strategies. The designed baseline agents are proven to be useful in activation, training and assessment of doctor's abilities regarding social and cognitive adaptation for effective shared decision making. Implications for future research and extensions are discussed.

1 Introduction

Recently, the use of cognitive agents in interactive applications has gained lots of attentions. It has been proven that even very simple agents can exhibit complex emergent behavioural patterns Schelling (1958); Hegselmann et al. (2002). Advanced cognitive agents are able to produce detailed simulation of human learning, prediction, adaptation and decision making Marewski and Link (2014); Lee et al. (2015); Salvucci and Taatgen (2008). They are also perfectly capable to play the role of a believable character in various human-agent settings. Cognitive agents have been beneficially used for training various human skills, e.g. negotiation and coordination skills in job interviews and trading Lin et al. (2014); de Weerd et al. (2015), metacognitive skills in various learning settings and domains Rus et al. (2009); Van Helvert et al. (2016). Cognitive agents allow creating and manipulating specific situations in which human social learning and human interactive behaviour can be studied. It has been also demonstrated that the integration of cognitive agents into a dialogue system has important advantages for effective implementation of complex (multi-agent) dialogue models Malchanau (2018).

In this paper, we address the use of cognitive agents to train efficient decision making strategies in asymmetric medical negotiations. Here, learning occurs through the partner's interpretation of own and others successes and failures, and through reflection on the action consequences for the interactive outcomes Bandura and Walters (1977); Ajzen (1991). We designed baseline agents to investigate the effectiveness of various decision

*This is a pre-print version of the following paper: Volha Petukhova, Firuza Sharifullaeva and Dietrich Klakow. (2019) Modelling Shared Decision Making in Medical Negotiations: Interactive Training with Cognitive Agents. Matteo Baldoni, Mehdi Dastani, Beishui Liao, Yuko Sakurai, and Rym Zalila-Wenkstern (eds). The 22nd International Conference on Principles and Practice of Multi-Agent Systems (PRIMA2019). Lecture Notes in Computer Science, Vol.11873, Springer. DOI: https://doi.org/10.1007/978-3-030-33792-6_16




Figure 1: Decision Making and Social Learning Model. Adapted from [30].

making strategies for the patient’s therapy adherence behaviour. The agents are based on the recent developments in cognitive modelling. Cognitive models are developed producing detailed simulations of human decision making performance. We extend them to facilitate realistic social learning and interactive scenarios.

We first set a scene for our investigations discussing the important characteristics of medical interactions to be considered when designing social intelligent agents for the training of efficient decision making (Section 2). We review models of individual decision making and social learning aspects which are modelled to explain and predict changes in human behaviour in general. Subsequently, we present the concept of Shared Decision Making (SDM) in medical context. Section 3 discusses the design aspects related to the human-agent interaction giving a global outline of a set of negotiation tasks with increasing scenario complexity and performed interactive actions. Section 4 presents our baseline agents with details for decision making strategies selection and the agent’s feedback. In Section 5, we present the results of the experiments evaluating baseline agents simulating various decision making strategies. Finally, we summarize our findings and outlines directions for the future research and development.

2 Decision Making and Social Learning

International research has produced great deal of models describing how individuals make decisions. The problem is approached at many levels, e.g. concentrating on psychological processes, and on biological and environmental factors. According to the most widely applied theories which attempt to explain and predict human behaviour and behavioural changes, Theory of Planned Behaviour Ajzen (1991) and Social Cognitive Theory Bandura (2001), there are three key sets of decision-making determinants defined: (1) individual attitudes towards behaviour and its outcomes (A_p), (2) perceived social norms (PN) and (3) perceived behavioural control (PBC), see Figure 1. A_p beliefs are concerned with the individually perceived *importance* of the behaviour given the known benefits, risks and threats and the perceived level of *readiness* to perform (execute) certain behaviour. Individual decision making is influenced by the individual beliefs (confidence) about abilities to perform and control behaviour and its outcome, i.e. *self-efficacy*. If an individual has developed positive attitudes towards a particular behavioural change, e.g. ceasing smoking, however believes he is not capable to maintain this behaviour, this will not lead to intention to perform this action. Outcome expectations and self-efficacy are very important determinants of health behaviour and depend on features such as perceived difficulty of the behaviour and/or the perceived certainty of its benefits Strecher et al. (1986).

Perceptions about attitudes of others can influence the individual decision making. Learning that is facilitated by observation of, or in interaction with another individual or their products, is defined as social learning Hoppitt and Laland (2013). Individual experiences obtained through observation of successful or unsuccessful performance of others, *vicarious experiences*, may account for a major part of learning throughout life Ajzen (1991) and influence (self-)efficacy expectations. Social learning, however, depends on the ability of an individual to take another individual’s perspectives and use other people’s behaviour as a guide to their own. Thus, learning occurs through reflecting on experiences. The cognitive capacity to attribute mental states to self and others is considered as a key factor of social and cognitive adaptation and is known as Theory of Mind (ToM, Premack and Woodruff (1978)) skills. ToM abilities significantly influence decision-making processes, can enhance motivation and self-efficacy, and can be successfully trained in human-agent setting.

2.1 Asymmetries in doctor-patient communication

In medical encounters, certain asymmetries and an imbalance in the knowledge and relationship between interlocutors are observed. Doctors empowered with institutional authority may expand the distance to their non-expert patients. Differences in knowledge, inequity in social status and power may lead to miscommunication, have trust damaging effects and, after all, decrease patients' therapy adherence Rodriguez-Osorio and Dominguez-Cherit (2008).

Recently, Shared Decision-Making (SDM) models have been evaluated showing that patient's active participation in the decision-making process is fundamental for the therapy's success Mazur (2001). The shift from therapy 'compliance' to 'adherence' implies that patients have more autonomy in defining and following their medical treatments.

2.2 Shared Decision Making: monitoring attitudes and enhancing self-efficacy

In SDM, the form of the interaction such as negotiation plays an important role. Medical negotiations do not necessarily involve a conflict as in the case of distributive negotiations where any gain of one party is made at the expense of the other(s) Sandman (2009). The key characteristics of medical SDM are that (1) at least two participants - doctor and patient - are involved; (2) both parties share information; (3) both parties take steps to build a consensus about the preferred treatment; and (4) an agreement is reached on the treatment to be implemented Charles et al. (1997).

Medical negotiations can be accurately described in terms of a balancing of values like the patient's best interest, patient autonomy and patient adherence Wirtz et al. (2006). The patient's best interest is often modelled by taking the professional (doctor's) view on a patient's best interest. The patient's autonomy is respected based on an assessment of whether the patient will adhere to the treatment in question. To make it a shared decision, the patient will have to agree on and accept a compromise. The medical SDM can be best modelled as interest-based bargaining where parties reason about the interests of each other and, building a consensus, negotiate the best possible mutual agreement Makoul and Clayman (2006). Interests include the needs, desires, concerns, and fears important to each side. Parties' preferences regarding the possible agreements may be not completely in conflict and they may be adapted depending on the perception of the preferences and behaviour of others. Thus, our model accounts for (1) participant's beliefs about perceived importance and desires concerning the certain behaviour and its outcomes (attitudes); (2) participant's beliefs about his abilities to perform this behaviour (self-efficacy); (3) and the beliefs of the same kind about his partner.

3 Design

3.1 Use case and scenarios

Patient's therapy non-adherence have different forms, e.g. skipping the intake of the prescribed medicines, the failure to keep appointments, to follow recommended dietary, lifestyle changes, recommended preventive health practices Fawcett (1995); Peterson et al. (2007).

The best practice for health behaviour change remains face-to-face interaction with an expert counsellor Bickmore and Giorgino (2006). A medical professional dealing with a non-adherent patient should master negotiation skills in order to reach an efficient agreement by proposing regimen that are feasible to follow, showing an appropriate understanding of patient's condition and treatments, and exercising the right influence on patient's beliefs and attitudes taking patient's social, cognitive and economic constraints into account Burgoon et al. (1987).

The domain selected for our use case concerns the treatment of diabetes of Type 2. The patient-doctor negotiation scenario was designed based on the recommendations for patients according to the International Diabetes Federation (IDF, 2017) addressing four issues: (1) medication, (2) diet, (3) activity and (4) exercise recommendations. Each of these issues involves four important negotiation *options* with preferences assigned representing parties negotiation positions, i.e. preference profiles. Preferences are weighted in order of importance and defined as the participant's beliefs about *attitudes* towards certain behaviour and *abilities* to perform this behaviour. The goal of each partner is to find out the preference of each other and to search for the best possible mutual agreement.




Figure 2: Example of a participant's preference profile.

Five scenarios of various complexity based different preference profiles were designed. The preferences strength was communicated to the negotiators through colours, see Figure 2. The preferences values range from -4 for highly dispreferred (dark orange) options to 4 for highly preferred (dark blue) options. The preferences can be set by human participant and/or generated automatically by the system dependent what type of partner the human participant wants to negotiate with. A graphical user interface was designed where human trainees, including domain experts, can specify their preferences and select partner's preference profile, either conflicting, matching or overlapping. Three types of profiles are specified, see also Table 1:

- *Conflicting*: negotiators' preferences are completely the opposite to each other.
- *Matching*: preferences are of the same polarity, but different in strength.
- *Overlapping*: some preferences are of the same polarity and strength.

The human participant - doctor - negotiates with various agents who simulates various types of patients, selecting one option per issue. To create and manipulate various situations in which doctor's abilities regarding social adaptation in decision making are activated and assessed, simulated patients have different preferences and are equipped with a basic set of negotiation and decision-making strategies which are evaluated in human-agent setting and compared to the human-human performance.

Table 1: Preferences profiles examples.

Type	Human	Agent
Conflicting	[-2, 3, -1, 4]	[2, -3, 1, -4]
Matching	[2, -3, -1, 4]	[1, -2, -2, 3]
Overlapping	[-2, 3, 1, 4]	[-2, -1, 1, 4]

3.2 Negotiation space, actions and strategies

Negotiation partners state their positions and set the conditions for the further exchange. The obtained information serves to establish jointly possible values constituting the *negotiation space*. In medical negotiation, this happens when a doctor explores patient's attitudes and abilities: elicits description of preferable actions, encourages patient to share his experiences, and matches those with his professional expertise. The better all possible actions and parties experiences are explored and discussed, the better future agreements are reached. In our design, human trainees (doctors) negotiates therapeutic interventions which they believe are medically mandated and what they understand are desired by the patient. The success of the interactive 'claiming' and 'giving up' space depends not only on the medical competence of the doctor, but also on his social competences and ToM skills. Participants' tasks are to determine their own actions, to interpret partner's actions, and to adjust their behaviour accordingly. The agent achieves this by taking the perspective of its partner and using its own knowledge to evaluate the partner's strategy, i.e. apply ToM skills. The agent holds three sets of preference values: the agent's own preferences (zero ToM), the agent's beliefs about the partner's preferences (first-order ToM), and the agent's beliefs about the partner's beliefs about the agent's preferences (second-order ToM).

The successful medical negotiation involves adequate disclosure by both parties indicating their values as well as other relevant matters. We specified the set of actions based on the ISO 24617-2 dialogue act taxonomy Bunt et al. (2012) tailored to the medical counselling domain using the Roter Interaction Analysis System (RIAS, Roter and Larson (2002)). Table 2 provides an overview of actions modelled to be performed by the

Table 2: Taxonomy of the agent’s actions. Adapted from the ISO 24617-2 dialogue act taxonomy enriched with the RIAS categories proposed for future extensions(*).

Socio-emotional exchange(*)	Task-focused exchange	Semantic content			Global affect(*)
		Modality	Negotiation Move	Issue(options)	
show approval give compliment show empathy show concern/worry reassure/encourage ask for reassurance show understanding show compassion	(open-ended) set question (forced) choice question propositional questions check questions inform/answer (dis-)agreement advise suggest request/instruct offer promise	preference ability necessity acquiescence	(final) offer exchange concession deal withdraw	Figure 2 related to: therapeutic regimen lifestyle	uncertainty anxiety dominance attentiveness engagement friendliness anger

Table 3: Instance definition.

Slot	Value [range]	Explanation
Strategy	[<i>cooperative</i>] [<i>competitive</i>] [<i>neutral</i>]	the strategy associated with the instance
Agent-move-value-agent	[-4, 4]	the number of points the agent’s gets from his own move
Partner-move-value-agent	[-4, 4]	the number of points the partner’s move brings to the agent
Partner-move-greater	[<i>true</i> <i>false</i>]	true if the partner’s move brings at least as much as the agent’s one, otherwise - false
Next-move-value-agent	[-4, 4]	the number of points that the next best move can bring to the agent
Utility	[0, 17]	how valuable are the partner’s suggestions made by now, see equation 4
Shared utility	[0, 1]	how valuable are both partner’s suggestions for them, see equation 6
Agent-move	($M_1, \dots, M_n$)	the move that the agent should make in this context (2)
Partner-move	($M_1, \dots, M_n$)	the move that the agent believes the partner should make in this context (Table 2)
Compensation	[1, 4]	if the agent’s move is of the concession or exchange type, what is the minimum utility that the agent should look for choosing an alternative option

baseline agent and the categories proposed for future extensions (marked *). Semantic content of dialogue acts specifies (modalized) negotiation moves and their arguments expressing the importance, desires and abilities concerning the certain behaviour and its outcomes, i.e. patient’s attitudes and self-efficacy assessments.

Negotiation moves types, sequences and the expressed modality are used to compute negotiation strategies. We consider negotiators as *cooperative* if they share information about their preferences with their opponents. A cooperative negotiator prefers the options that have the highest collective value. If not enough information is available to make this determination, he will elicit this information from his opponent. A cooperative negotiator will not engage in positional bargaining holding on to a fixed set of preferences regardless of the interests of others, instead, he will attempt to find issues where a compromise is possible. *Competitive* negotiators prefer to assert their own preferred positions rather than exploring the space of possible agreements. A competitive negotiator will ignore partner’s interests and requests for information. Instead, he will find his own ideal position and insist upon it in the hope of making the opponent concede. He will threaten to end the negotiation. The competitive negotiator will accept an offer only if he can gain a significant number of points from it.

4 Baseline agents

4.1 Agent’s knowledge and memory

The baseline agents are designed using ACT-R cognitive architecture implemented in Java¹. Agent’s knowledge is encoded in instances. An instance consists of a representation of the current state of the world (what do I know, what do I know about others, what am I asked, what can I do, what has happened before), and an action to be taken in that situation (give information, run tests, examine something, reason about others, change attitude, etc.). An instance has a form of slot-value pairs representing context and actions. Table 3 depicts the structure of an instance.

Instances are stored in an ACT-R declarative memory which is represented as traces of instances used. At the beginning of the interaction, the agent may have no or weak assumptions about the partner’s preferences, thus instances may be empty or partially filled in. As the interaction proceeds the agent builds up more knowledge, i.e. learns by observing actions of others, storing those as instances, or by trying out actions itself and adjusting its instances based on feedback.

¹A Java Simulation and Development Environment for the ACT-R Cognitive Architecture - homepage <http://cog.cs.drexel.edu/act-r/about.php>

4.2 Agent's feedback

The agent's feedback actions are designed to assist the trainee (medical professional) to form a mental model of an agent (patient). By providing real-time feedback about the agent's cognitive state, the trainee should become aware of how his own actions influence patient's beliefs. This feedback comes in three forms: evaluation of agent's beliefs about trainee's preferences, evaluation of patient's self-efficacy beliefs (possible actions) and evaluation of trainee's negotiation strategy.

Every time the trainee makes a move expressing his preferences, the agent matches it to its preferences and available strategies and computes the most plausible action(-s), it also provides alternatives and plans possible outcomes. Since the agent knows why certain actions are performed, it can explain why its and partner's choices lead to the specific outcome. Evaluation of trainee strategy provides the trainee with feedback about how the agent views their overall performance. Both participants operate under constrain that negotiation outcome should be acceptable for both partners. Thus, the interactions were evaluated in terms of the percentage of reached agreements², percentage of negative outcomes³ and Pareto efficient outcomes⁴.

The summative feedback is generated when a negotiation round is over and includes the assessment of the overall cooperativeness level, percentage of reached agreements and negative outcomes, and the Pareto efficiency scores. All scores accumulate with each negotiation round indicating the learning progress. In the feedback on the best possible outcome is included.

4.3 Action selection decisions

The decision-making process can be simple when randomly picking options out of the available ones, or complex when systematically rating different aspects of the available choices. Human decision-making strategies may depend on various factors, including how much time they have to make the decision, the overall complexity of the decision, and the amount of ambiguity that is involved. According to Payne et al. (1993), decision makers choose strategies adaptively in response to different task demands, and often apply simplified shortcuts heuristics that allow fast decisions with acceptable losses in accuracy. Moreover, simple heuristics are often more or at least equally accurate in predicting new data compared to more complex strategies Czerlinski et al. (1999). Simple heuristics are more robust, extracting only the most important and reliable information from the data, while complex strategies that weight all pieces of evidence extract much noise, resulting in large accuracy losses when making predictions for new data Pitt and Myung (2002).

We implemented and assessed three decision-making strategies to simulate situations where different alternatives will be selected by the agent in a certain context in order to achieve acceptable outcomes: *recognition* or *activation*-based retrieval, *compensatory* models and *non-compensatory* heuristics.

When the agent has to make a decision, it activates his declarative knowledge and retrieves direct and vicarious experiences that are the most *active*, i.e. most recent and frequent. For every instance i in the set, the activation is computed as

$$A_i = B_i - MP \sum_{v,d} (1 - Sim(v, d)) \quad (1)$$

where MP , a mismatch penalty, reflects the amount of weighting given to the matching, i.e. the higher MP , the stronger the activation is affected by the similarity⁵. B_i , base level activation, is computed as

$$B_i = \ln \sum_{j=1}^n t_{ij}^{-d} \quad (2)$$

where t_{ij} is the time elapsed since the j_{th} presentation or creation of the instance i , and d is the memory decay rate⁶. The similarity $Sim(v, d)$ between the goal value v and the actual value d held in the retrieved instance is computed as

²We consider the agreement reached if parties agreed on all four issues.

³Negative deals are considered as flawed negotiation action, i.e. the sum of all reached agreements resulted in an overall negative value meaning that the partner made too many concessions and selected mostly dispreferred bright 'orange' options (see Figure 2).

⁴The negotiation is Pareto efficient if none of the negotiators could have achieved a higher score for themselves without a reduction in score of the other negotiator.

⁵We set MP constant high at 5, consistent with the value used in Lebiere et al. (2000). To disable MP , it can be set at 0.

⁶In the ACT-R community, 0.5 has emerged as the default value for the parameter d over a large range of applications, Anderson et al. (2004).

$$Sim(v, d) = \frac{1.0}{((v - d)^2 / 2.0 + 1.0)} \quad (3)$$

The agent makes its next move based on the value of ‘agent-move’ slot, see Table 3. The ACT-R mechanisms effectively account for both the effects of recency - more recent memory traces are more likely to be retrieved, and frequency - if a memory trace has been created or retrieved more often in the past it has a higher likelihood of being retrieved. By disabling *MP*, the agent will be able to retrieve past instances for reasoning even when a particular situation has not been encountered before. An instance does not have to be a perfect match to a retrieval request to be activated. ACT-R can reduce its activation to compute partial matching Lebiere et al. (2000). If the value is missing in an instance, ‘blending’ is proposed as a generalization of the retrieval mechanism, allowing to retrieve values from multiple instances Lovet et al. (1999).

To assess alternative decision making strategies, procedural knowledge was incorporated and condition-action rules were defined. For example, the weighted-additive rule (WADD) is computed by, first, weighting the dimensions (i.e. issues or criteria) on their relative importance by summing preference values of all attributes specified within this dimension divided by their number and then multiplying the preference values with their respective importance weights. To form an overall evaluation, the products are summed and the option with the highest value is chosen, *weighted sum model* (WSM). The total importance of an alternative A_i^{WSM} is computed as

$$A_i^{WSM} = \sum_{j=1}^n w_j u_{ij} \quad (4)$$

where w_j denotes the relative importance of the dimension D_j and u_{ij} is the utility value of alternative A_i when it is evaluated with relation to the dimension D_j . In case all dimensions are considered of equal importance as in our case, equal-weight rule (EQW) is applied. In compensatory decisions, a negative value of one attribute can be compensated by an equal or higher value of another attribute. Thus, compensatory strategy involves a systematic evaluation of multiple attributes and works well if all information is available modelling rational decision choices the best.

In contrast, non-compensatory strategies assume that decisions are often made based on the rejection of undesirable alternatives on the basis of one, or at most a few criteria. When faced with a more complex (multi-alternative) decision task, the subjects employed decision strategies designed to eliminate some of the available alternatives as quickly as possible and on the basis of a limited amount of information search and evaluation Billings and Scherer (1988). Here, values on the most salient dimension are processed first and alternatives that score lower are eliminated as unsatisfactory, also known as take-the-best (TTB) strategy. Values in other dimensions are not used for a compensation, thus a negotiator makes no trade-offs between attributes.

4.4 Negotiation strategy selection

In our first approach, the agent adjusts its negotiation strategy according to the perceived level of the opponent’s cooperativeness. The agent starts neutrally, requesting the partner’s preferences. If the agent believes the partner is behaving cooperatively, the agent will react with an cooperative negotiation move. If the agent experiences the partner as competitive, it will switch to a competitive mode. Such strategy is observed in human negotiation and coordination games Kelley and Stahelski (1970); Smith et al. (1982) and we call it *matching coordination* (MC). In interactions, interlocutors often mirror decision making behaviour of their opponents, in particular, where a clear division of roles and an asymmetric distribution of interactional power is observed. To simulate the mirroring decision making, activation of declarative knowledge within ACT-R’s declarative memory is used and the instance with the highest matching score M_{ip} is retrieved computed as

$$M_{ip} = A_i - MP \sum_{v,d} (1 - Sim(v, d)) \quad (5)$$

According to the compensatory decision-making model, in choice situations with multiple alternatives in multiple dimensions, if for a certain alternative in one dimensions scores are low and a higher score on another dimension can compensate for it, this alternative will be adopted. In other words, a high score on one dimension (e.g. medicine) can compensate a lower score on another dimension (e.g. diet). These are then combined to maximize a utility. For each dimensions, the overall scores are considered and alternatives with the highest

scores are chosen Einhorn and Hogarth (1981). Since, the goal of our decision-makers not to maximize their own utility but to achieve an acceptable, ideally Pareto efficient outcome, the negotiator will try to maximize the shared utility in one or multiple negotiation rounds across dimensions computed based on utility earned by both partners as a proportion of maximum utility possible

$$U_{shared} = \frac{u_{agent} + u_{trainee}}{maxU_{agent} + maxU_{trainee}} \quad (6)$$

Unlike the MC agent, the utility-based (UB) agent does not attempt to find the appropriate instance in the memory based on the similarity of the slots to the current context and the instance activation value. Decision will be made based on the accumulated agent's *utility* (u_{agent}) and the estimated *shared utility* (U_{shared}) values.

The agent starts by selecting options with the highest utility values, procedure which is regularly observed in human-human negotiations since it is always easier to bargain down than to bargain in. If the agent will continue to win points insisting on his preferences while ignoring the preferences of his partner's, after it collects enough points, e.g. it has already reached the amount of the half maximum score possible in the current scenario, it will switch to cooperative mode. Playing cooperatively, if the agent starts loosing too much, so that its utility score becomes lower than the half maximum score possible in the current scenario, it will switch to the competitive mode to compensate for its previous losses. We express the decisions to change a negotiation strategy as:

$$\text{SetStrategy}(u_{agent}, U_{shared}, \text{round}) = \begin{cases} \text{cooperative,} & \text{if } u_{agent} < \frac{1}{2} * \text{round} * (\max U_{agent}) \\ \text{competitive,} & \text{if } u_{agent} \geq \frac{1}{2} * \text{round} * (\max U_{agent}) \\ & \text{and } U_{shared} \geq \text{threshold} \\ \text{neutral,} & \text{otherwise} \end{cases}$$

Threshold for the shared utility value can be set via GUI; by default it set on 0.5, meaning that at least the half of the mutually acceptable agreements have to be reached.

To simulate repetitive negotiations, e.g. to analyse trainee's learning behaviour over time⁷, we model cross-rounds decision making strategies based on the negotiation history. Thus, our strategy changing policies account for successes and failures of previous rounds taking the current round number into consideration.

To simulate non-compensatory decision-making, the agent insists on the options beneficial for it, until an agreement on exactly one option in each dimension is reached. Here, the 'exchange' actions are impossible. The options with negative scores are eliminated by the agent. This strategy may result in agent's position bargaining when it sticks to the preferred options with the hope that partner concedes, and if not it breaks the negotiation proposing the final offer playing 'take-or-leave-it' strategy. The agent's behaviour can be described as follows:

$$\text{SetStrategy}(u_{agent}, U_{shared}, u_{min}) = \begin{cases} \text{cooperative,} & \text{if } U_{shared} < \text{threshold and} \\ & u_{min} < u_{agent} \\ \text{competitive,} & \text{if } U_{shared} \geq \text{threshold} \\ \text{take-or-leave-it,} & \text{if } u_{agent} = 0 \end{cases}$$

The time (number of moves) until the agent breaks the negotiation is configurable via GUI, as well as its $utility_{min}$.

5 Evaluation

We conducted a set of small-scaled evaluation experiments with ten participants involved in human-human negotiations and different five participants in human-agent negotiations. None of the participants was familiar with the topic. The age of the participants varied from 23 to 32 years old.

In human-human setting, one participant was randomly assigned the role of a doctor, the other participant the role of a patient. Each participant received his cover story and instructions, as well as the preference profile for each scenario, as shown in Figure 2. Participants were not allowed to share their preference information with each other. They were asked to negotiate an agreement with the highest possible value according to their preference information. Participants were allowed to break the negotiation if they feel that it is impossible to reach an agreement on the provided terms. No further rules on the negotiation process or time constraints were

⁷In real life, doctors and patient often do not meet only once, but share certain interaction history with each other.

Table 4: Comparison of human-human and human-agent negotiation performance.

Evaluation criteria	Human vs human	Human vs		
		MCS-agent	CS-agent	NCS-agent
Number of dialogues	25	25	25	25
Mean dialogue duration (in #turns)	23.0	22.4	17.2	23.8
Number of offers/per round	16.0	15.7	10.7	16.4
Agreements (in %)	78.0	95.6	87.0	76.2
Pareto efficient agreement (in %)	82.4	99.9	95.0	76.0
Negative deals (in %)	21.0	47.8	21.7	33.3
Cooperativeness rate (in %)	39.0	65.5	66.0	54.2

imposed. The interactions were recorded, transcribed and analysed. In total, we collected 25 human-human negotiations comprising about 575 speaking turns.

In human-agent negotiations, each human trainee in the doctor’s role negotiated with the simulated patient (agent) getting instructions similar to ones in the human-human setting. A trainee negotiated with an agent who uses (1) a **M**atched **C**oordination decision making **S**trategy (MCS); (2) a **C**ompensatory decision making **S**trategy (CS); and (3) a **N**on-**C**ompensatory decision making **S**trategy (NCS). A human trainee played five rounds with each agent within five randomly selected scenarios of different complexity. Totally, 75 human-agent negotiations were collected comprising 2049 turns.

5.1 Results

We compared the agents and human performance on the *number of agreements* reached, the ability to achieve *Pareto efficient* outcomes, maintain a reasonable *level of cooperativeness* while avoiding *negative deals*.

The obtained results are summarized in Table 4. We observed that trainees spent on average more time negotiating with a human than with a simulated patient (23 vs 21.1 turns). In human-human setting, actions other than related to the negotiation task were observed. Along with task-related offers, human participants performed frequent feedback, turn and time management, discourse structuring acts concerned with topic switches moving from one issue to another and decisions to continue, delay, reschedule or terminate the ongoing discussion and/or whole interaction. Agents, by contrast, were designed to produce actions concerned with the negotiation task. Humans reached on average a lower number of agreements when negotiating with agents than negotiating with each other, 78% vs 86.3%. Negotiations with humans as well as with the NCS-agents were often terminated or threaded to be terminated when partners were not willing to concede. Participants negotiating with the MCS- and CS-agents reached a similar number of Pareto efficient agreements (close to 100%). In negotiations with the NCS-agents, about 25% of outcomes were not Pareto efficient. Human participants showed a higher level of cooperativity when interacting with an agent, i.e. more than 50% of all actions are annotated as cooperative. We concluded that agents were useful for trainees to understand partner’s attitudes and abilities, to explore the negotiation space more optimally and to adapt their behaviour accordingly. A higher number of negative deals was observed for human-agent pairs, 21% vs 33%, mostly when interacting with the MCS-agents. Agents with compensatory strategies performed the best showing that they allow to efficiently explore action-outcome connections - the behaviour which leads to a limited number of negative deals.

We assessed different decision making strategies applied in scenarios of various complexity. The scenario complexity has been computed taking the difference between agent’s and trainee’s preference profiles, i.e. the higher the difference the more complex the scenario, see also Table 1. The results depicted in Figure 3 show that negotiations with all agents ended successfully, with a reasonable number of Pareto efficient agreements. When scenario was getting too complex, the NCS-agent was performing at least as accurate as the other two. Since non-compensatory decision-making does not require extensive cognitive efforts to evaluate attributes and reflect on the partner’s behaviour, it may be applied rather effectively under ‘unfavourable’ conditions like time pressure, distractions or physical and psychological exhaustion.

Asymmetries in preferences may not always yield the desired outcomes. This forces participants’ to explore the negotiation space more thoroughly, apply sophisticated ToM skills. Figure 4 illustrates examples of negotiation scenarios, a cooperative and a competitive ones. Our results showed that in cooperative settings, negotiations tend to resolve more quickly and are more likely to be Pareto efficient. In competitive settings, on the other hand, negotiations tend to take more turns, and are more likely to result in an outcome that is not Pareto efficient.




Figure 3: The impact of scenario complexity on (a) the number of agreements reached; (b) Pareto efficiency of these outcomes; and (c) the number of negative deals accepted.




Figure 4: Negotiation space for (a) a cooperative and (b) a competitive scenario. The top left panels show the possible negotiation outcomes in terms of the score for the patient (vertical) and the doctor (horizontal). Blue dots indicate outcomes that are acceptable to both negotiators, while red dots indicate unacceptable outcomes. Large dots show the possible outcomes for the patient (purple) and the doctor (yellow), as well as the final agreement (green). The sequence of offers made during the negotiation is indicated by connected dots. The panels on the right and the bottom show the same information from the perspective of the patient and the doctor, respectively.

Our general observations showed that with a more cooperative CS- or MCS-agent, the trainee learns to adapt his behaviour acting more cooperatively and gains scores close to his maximum utility. A competitive CS-agent often challenges the human to compensate for his previous losses, but it is punished by losing points. When a competitive human wants to take advantage of the cooperative MCS-agent, he starts to lose points and is unable to reach mutual agreements. In an NCS setting, the agent selects the competitive negotiation strategy more frequently and forces the human to act more competitively as well.

Our in-depth analysis of the logged interactions revealed that trainees used different negotiation tactics which resulted in different outcomes. Negotiators often delayed making complete agreements on the first discussed issue until the agreement on the next one is secured. They frequently revised their past offers. The order in which the issues are negotiated, i.e. negotiation agenda, might influence on the overall outcome. The most common strategy observed was issue-by-issue bargaining. Various negotiation tactics were concerned with the partners' alacrity to reveal or hide their preferences. When all preferences are brought on the negotiation table from the very beginning, agreements that are Pareto efficient were reached faster.

It has been noticed that not only asymmetries in preferences and participant's status may influence the decision making process, but participants of different gender and personality, and in different emotional state may adopt divergent strategies under identical conditions. To investigate relationships between participant's intrinsic characteristics and various dependent variables, a larger study needs to be conducted where personality traits will be assessed and specific emotions induced prior to the experiments.

Concerning the learning progress supported by the interaction with the agents, a follow up test-retest study with medical professionals and students will be performed. Our preliminary feedback indicates that most respondents think that the system presents an interesting form of skills training. The vast majority of users learned how to complete their tasks successfully in consecutive rounds.

6 Conclusions and future work

In this study, we investigated whether simple cognitive models can produce plausible simulations of human decision making performance acting as believable agents in human-agent interactive learning setting. Cognitive modelling of human intelligent behaviour not only enables better understanding of complex mental tasks, but also allow designing and controlling learning and interactive situations, scenarios and actors to assess human abilities for joint attention, social mirroring and cooperative actions - important in shared decision making process.

To facilitate realistic scenarios to study human social interactive behaviour, the agents who simulates different types of patients are equipped with different sets of preferences encoding desires and self-efficacy attitudes, and apply various negotiation and decision-making strategies. Human-human and human-agent negotiations were evaluated in terms of the number of agreements reached and their Pareto efficiency, the number of the accepted negative deals and the cooperativeness of the negotiators' actions. The results show that agents can act as credible opponents to train decision making for improved negotiation performance. Agents with compensatory strategies, due to their ability to explore action-outcome connections systematically using all available information, approximate rational decisions the best. Agents that match and coordinate their decisions with their partners show powerful abilities for social mirroring and cooperative actions, skills that are important for human medical professionals to master. Simple non-compensatory heuristics are shown to be at least as accurate, and in complex scenarios even more effective, than more cognitive-intensive strategies. The designed baseline agents are proven to be useful in activation, training and assessment of doctor's abilities regarding social and cognitive adaptation for effective shared decision making.

Past research has indicated that people select strategies adaptively depending on the situation they face Payne et al. (1993). Our baseline agents, while already showing convincing human-like decision making behaviour, are rather constrained in their abilities for adaptation. The strategies were programmed rather than learned. In the future, we plan to assess the impact of various (pragma-)linguistic and interactive strategies on the adaptive decision-making behaviour while accounting for the interwoven relationship between multi-modal language-specific schemes and emotional, social and cultural determinants. The design of sophisticated models of human decision-making behaviour have the potential to form the foundation for a new generation of interactive social systems.

References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2):179–211.
- Anderson, J., Bothell, D., Byrne, M., Douglass, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, 111(4):1036.
- Bandura, A. (2001). Social cognitive theory: An agentic perspective. *Annual review of psychology*, 52(1):1–26.
- Bandura, A. and Walters, R. (1977). *Social Learning Theory*. Upper Saddle River, N.J.: Prentice-Hall Inc.
- Bickmore, T. and Giorgino, T. (2006). Health dialog systems for patients and consumers. *Journal of biomedical informatics*, 39(5):556–571.
- Billings, R. S. and Scherer, L. L. (1988). The effects of response mode and importance on decision-making strategies: Judgment versus choice. *Organizational Behavior and Human Decision Processes*, 41(1):1–19.
- Bunt, H., Alexandersson, J., Choe, J.-W., Fang, A. C., Hasida, K., Petukhova, V., Popescu-Belis, A., and Traum, D. R. (2012). ISO 24617-2: A semantically-based standard for dialogue annotation. In *LREC*, pages 430–437. Citeseer.
- Burgoon, J. K., Pfau, M., Parrott, R., Birk, T., Coker, R., and Burgoon, M. (1987). Relational communication, satisfaction, compliance-gaining strategies, and compliance in communication between physicians and patients. *Communications Monographs*, 54(3):307–324.
- Charles, C., Gafni, A., and Whelan, T. (1997). Shared decision-making in the medical encounter: what does it mean?(or it takes at least two to tango). *Social science & medicine*, 44(5):681–692.
- Czerlinski, J., Gigerenzer, G., and Goldstein, D. G. (1999). How good are simple heuristics? In *Simple heuristics that make us smart*, pages 97–118. Oxford University Press.
- de Weerd, H., Broers, E., and Verbrugge, R. (2015). Savvy software agents can encourage the use of second-order theory of mind by negotiators. In *CogSci*.
- Einhorn, H. J. and Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgement and choice. *Annual review of psychology*, 32(1):53–88.
- Fawcett, J. (1995). Compliance: definitions and key issues. *The Journal of clinical psychiatry*.
- Hegselmann, R., Krause, U., et al. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of artificial societies and social simulation*, 5(3).
- Hoppitt, W. and Laland, K. N. (2013). *Social learning: an introduction to mechanisms, methods, and models*. Princeton University Press.
- Kelley, H. and Stahelski, A. (1970). Social interaction basis of cooperators’ and competitors’ beliefs about others. *Journal of personality and social psychology*, 16(1):66.
- Lebiere, C., Wallach, D., and West, R. (2000). A memory-based account of the prisoner’s dilemma and other 2x2 games. In *Proceedings of International Conference on Cognitive Modeling*, pages 185–193.
- Lee, H., Betts, S., and Anderson, J. (2015). Learning problem-solving rules as search through a hypothesis space. *Cognitive science*.
- Lin, R., Kraus, S., Mazliah, Y., et al. (2014). Training with automated agents improves people’s behavior in negotiation and coordination tasks. *Decision Support Systems*, 60:1–9.
- Lovett, M., Reder, L. M., and Lebiere, C. (1999). Modeling working memory in a unified architecture. *Models of working memory: mechanisms of active maintenance and executive control*, page 135.

- Makoul, G. and Clayman, M. L. (2006). An integrative model of shared decision making in medical encounters. *Patient education and counseling*, 60(3):301–312.
- Malchanau, A. (2018). *Cognitive architecture of multimodal multidimensional dialogue management*. PhD thesis, Ph. D. Thesis, University of Saarland, Saarbrücken.
- Marewski, J. and Link, D. (2014). Strategy selection: An introduction to the modeling challenge. *Wiley Interdisciplinary Reviews: Cognitive Science*, 5(1):39–59.
- Mazur, D. J. (2001). *Shared Decision Making in the Patient-Physician Relationship: Challenges Facing Patients, Physicians, and Medical Institutions*. Tampa, FL: American College of Physician Executives.
- Nowak, S. A., Matthews, L. J., and Parker, A. M. (2017). A general agent-based model of social learning. *Rand health quarterly*, 7(1).
- Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge university press.
- Peterson, A. M., Nau, D. P., Cramer, J. A., Benner, J., Gwadry-Sridhar, F., and Nichol, M. (2007). A checklist for medication compliance and persistence studies using retrospective databases. *Value in Health*, 10(1):3–12.
- Pitt, M. A. and Myung, I. J. (2002). When a good fit can be bad. *Trends in cognitive sciences*, 6(10):421–425.
- Premack, D. and Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain sciences*, 1(04):515–526.
- Rodriguez-Osorio, C. A. and Dominguez-Cherit, G. (2008). Medical decision making: paternalism versus patient-centered (autonomous) care. *Current opinion in critical care*, 14(6):708–713.
- Roter, D. and Larson, S. (2002). The roter interaction analysis system (rias): utility and flexibility for analysis of medical interactions. *Patient education and counseling*, 46(4):243–251.
- Rus, V., Lintean, M., and Azevedo, R. (2009). Automatic detection of student mental models during prior knowledge activation in metatutor. *International Working Group on Educational Data Mining*.
- Salvucci, D. and Taatgen, N. (2008). Threaded cognition: an integrated theory of concurrent multitasking. *Psychological review*, 115(1):101.
- Sandman, L. (2009). The concept of negotiation in shared decision making. *Health Care Analysis*, 17(3):236–243.
- Schelling, T. C. (1958). The strategy of conflict. prospectus for a reorientation of game theory. *Journal of Conflict Resolution*, 2(3):203–264.
- Smith, L., Pruitt, D., and Carnevale, P. (1982). Matching and mismatching: The effect of own limit, other’s toughness, and time pressure on concession rate in negotiation. *Journal of Personality and Social Psychology*, 42(5):876.
- Strecher, V. J., McEvoy DeVellis, B., Becker, M. H., and Rosenstock, I. M. (1986). The role of self-efficacy in achieving health behavior change. *Health education quarterly*, 13(1):73–92.
- Van Helvert, J., Petukhova, V., Stevens, C., de Weerd, H., Börner, D., Van Rosmalen, P., Alexandersson, J., and Taatgen, N. (2016). Observing, coaching and reflecting: Metalogue - a multi-modal tutoring system with metacognitive abilities. *EAI Endorsed Transactions on Future Intelligent Educational Environments*, 16(6).
- Wirtz, V., Cribb, A., and Barber, N. (2006). Patient–doctor decision-making about treatment within the consultation: a critical analysis of models. *Social science & medicine*, 62(1):116–124.